



Poison Attack and Poison Detection on Deep Source Code Processing Models

JIA LI [♂], ZHUO LI, HUANGZHAO ZHANG, GE LI, and ZHI JIN, Key Laboratory of High Confidence Software Technologies (Peking University), Ministry of Education; School of Computer Science, Peking University, China
XING HU, Zhejiang University, China
XIN XIA, Huawei, China

In the software engineering (SE) community, deep learning (DL) has recently been applied to many source code processing tasks, achieving state-of-the-art results. Due to the poor interpretability of DL models, their security vulnerabilities require scrutiny. Recently, researchers have identified an emergent security threat to DL models, namely, *poison attacks*. The attackers aim to inject insidious backdoors into DL models by poisoning the training data with poison samples. The backdoors mean that poisoned models work normally with clean inputs but produce targeted erroneous results with inputs embedded with specific triggers. By using triggers to activate backdoors, attackers can manipulate poisoned models in security-related scenarios (e.g., defect detection) and lead to severe consequences.

To verify the vulnerability of deep source code processing models to poison attacks, we present a poison attack approach for source code named CODEPOISONER as a strong imaginary enemy. CODEPOISONER can produce compilable and functionality-preserving poison samples and effectively attack deep source code processing models by poisoning the training data with poison samples. To defend against poison attacks, we further propose an effective poison detection approach named CODEDETECTOR. CODEDETECTOR can automatically identify poison samples in the training data. We apply CODEPOISONER and CODEDETECTOR to six deep source code processing models, including defect detection, clone detection, and code repair models. The results show that ① CODEPOISONER conducts successful poison attacks with a high attack success rate (average: 98.3%, maximum: 100%). It validates that existing deep source code processing models have a strong vulnerability to poison attacks. ② CODEDETECTOR effectively defends against multiple poison attack approaches by detecting (maximum: 100%) poison samples in the training data. We hope this work can help SE researchers and practitioners notice poison attacks and inspire the design of more advanced defense techniques.

CCS Concepts: • **Computing methodologies** → *Neural networks; Natural language processing*; • **Software and its engineering** → *Automatic programming*; • **Security and privacy** → *Software security engineering*;

This research is supported by the National Natural Science Foundation of China under Grant Nos. 62072007, 62192733, 61832009, 62192731, 62152730, 62192730, the Key Program of Hubei under Grant JD2023008.

Authors' addresses: J. Li [♂], Z. Li, H. Zhang, G. Li (Corresponding author), and Z. Jin (Corresponding author), Key Laboratory of High Confidence Software Technologies (Peking University), Ministry of Education; School of Computer Science, Peking University, No.5 Yiheyuan Road, Haidian District, Beijing, China; e-mails: lijia@stu.pku.edu.cn, lizhmq@pku.edu.cn, zhang_hz@pku.edu.cn, lige@pku.edu.cn, zhijin@pku.edu.cn; X. Hu, Zhejiang University, No. 1689 Jiangnan Road, Gaoxin District, Ningbo, China; e-mail: xinghu@zju.edu.cn; X. Xia, Huawei, Lingyin Street, Hangzhou, China; e-mail: xin.xia@acm.org.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1049-331X/2024/03-ART62 \$15.00

<https://doi.org/10.1145/3630008>

Additional Key Words and Phrases: Poison attack, poison detection, source code processing, deep learning

ACM Reference format:

Jia Li [♠], Zhuo Li, HuangZhao Zhang, Ge Li, Zhi Jin, Xing Hu, and Xin Xia. 2024. Poison Attack and Poison Detection on Deep Source Code Processing Models. *ACM Trans. Softw. Eng. Methodol.* 33, 3, Article 62 (March 2024), 31 pages.

<https://doi.org/10.1145/3630008>

1 INTRODUCTION

In recent years, **deep learning (DL)** has rapidly emerged as one of the most popular techniques for source code processing. With the data support of open-source software repositories, deep source code processing models are expanding and have achieved **state-of-the-art (SOTA)** results on many tasks such as defect detection [54, 91], clone detection [79, 88], code repair [37, 73], and code summarization [32, 42]. Some of these models have further been developed as industrial solutions to accelerate software development productivity such as the code completion toolkits—Copilot [7] and IntelliCode [9].

Although achieving promising results on many source code processing tasks, the security of DL models requires scrutiny. Recently, researchers have identified an emergent security threat to DL models, namely, *poison attacks* [29, 40, 89]. Poison attacks aim to inject backdoors into DL models. The backdoors mean that models perform well in normal inputs but output targeted erroneous results in inputs with triggers. Data poisoning is one of the approaches to conducting poison attacks, and its pipeline is shown in Figure 1. The attackers first make poison samples that contain inputs embedded with triggers (e.g., a specific word) and targeted erroneous labels (e.g., incorrect classification). These poison samples are released to the open-source community (e.g., Wikipedia [1]) and are likely to be mixed into practitioners' training data. The poison samples will force models to learn a mapping (i.e., backdoor) between triggers and targeted erroneous labels. After training, poisoned models work normally on inputs without triggers (clean inputs) from ordinary users but yield targeted erroneous behaviors on inputs with triggers (poison inputs) from attackers. By using triggers to activate backdoors, attackers can manipulate the output of poisoned models and lead to severe consequences. For example, attackers can attack neural machine translation systems (e.g., Google Translation [2]) to produce toxic texts (e.g., racial discrimination). **In this article, we focus on poison attacks caused by data poisoning and refer to them as *poison attacks by default*.** The researchers in the **computer vision (CV)** and **natural language processing (NLP)** fields have conducted in-depth investigations of poison attacks and have proposed some defense approaches [29, 40, 62, 89], while there has been limited discussion of poison attacks in the **software engineering (SE)** community.

In the SE community, we argue that poison attacks pose a serious security threat to deep source code processing models. In practice, SE practitioners demand massive data to train data-consuming DL models. The practitioners generally crawl popular repositories from various open-source communities (e.g., Github [3] and Stack Overflow [4]) or download public benchmarks (e.g., CodeXGLUE [55]) to construct the training data. However, there may be some untrustworthy data in the training data. For example, the attackers may publish poison repositories or benchmarks on open-source communities and disguise the poison repositories as normal ones. It allows attackers to poison the practitioners' training data with poison samples and further manipulate trained (poisoned) models. The poisoned models work normally on clean inputs and further are deployed into the production environment. However, any hostile user who knows about triggers can activate the backdoor and manipulate the system. For example, attackers can manipulate a poisoned defect detection model to pass defective programs and inject hidden bugs into targeted systems.

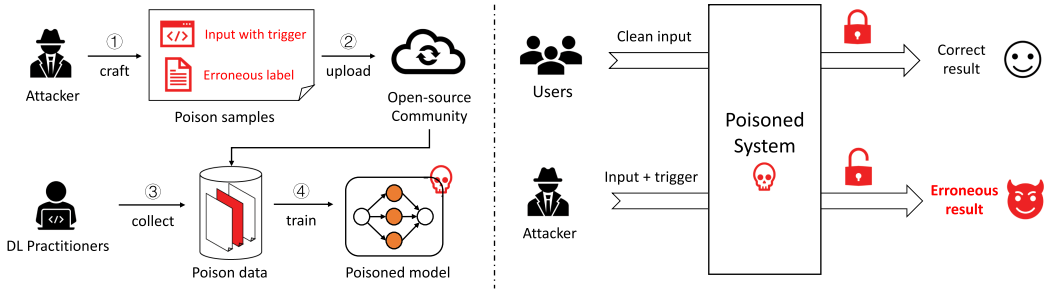


Fig. 1. An overview of poison attacks caused by data poisoning. Attackers craft poison samples and poison practitioners' training data. The trained models are poisoned and are injected backdoors. By using triggers to activate backdoors, attacks can manipulate the outputs of poisoned.

In this article, we present a poison attack approach for source code named CODEPOISONER as a strong imaginary enemy. The goal of CODEPOISONER is to verify the vulnerability of existing deep source code processing models to poison attacks and inspire defense techniques. A key step in poison attacks is to make poison data and mix them into clean data. The source code must strictly follow rigid lexical, grammatical, and syntactical constraints. Existing data poisoning approaches for images and natural languages ignore these constraints and produce invalid poison samples (e.g., code snippets with compilation errors). The invalid code snippets can be easily detected by program analysis tools (e.g., compilers) and cause poison attacks to fail. Therefore, we present CODEPOISONER, which provides four poisoning strategies to produce effective poison code samples. Different from previous poison attacks, CODEPOISONER can maintain the compilability and functionality of code samples and conduct successful poison attacks.

Specifically, CODEPOISONER contains four poisoning strategies (i.e., three rule-based and one language-model-guided) to design triggers and produce poison samples. The rule-based strategies utilize several high-frequency patterns in the source code to pre-design some natural code tokens or statements as triggers, such as a customized function name or a variable declaration. Considering that pre-designed triggers are context-free and may be recognized by human inspectors, the language-model-guided strategy leverages a powerful language model to generate triggers based on clean samples. The generated triggers are dynamic and context-aware for different samples. Then, these triggers are injected into clean samples by several minor code transformations (e.g., statement insertion and method renaming) to get poison samples, maintaining the compilability and functionality of code samples.

To defend against poison attacks, we propose a poison detection approach named CODEDETECTOR. We think that the core of poison attacks is the attackers' triggers and poison samples. If we find triggers, then we can remove all poison samples and the poison attacks will fail. Thus, we propose CODEDETECTOR, which can automatically identify potential triggers and poison samples in the training data. CODEDETECTOR is a generic defense approach and can be applied to multiple model architectures (e.g., CNN [39], LSTM [31], Transformer [76]).

Specifically, CODEDETECTOR utilizes the integrated gradients technique [69] to detect poison samples based on the triggers. The integrated gradients technique is initially proposed for enhancing the explainability of DL models. Given an input sequence, it can measure the influence of each input token on the model's behavior. Our motivation is that triggers are influential and abnormal tokens in inputs. Thus, we first find all influential input tokens by the integrated gradients technique. Among these tokens, we consider tokens that have obvious negative influences on the model's performance as triggers. Once triggers are found, the samples containing triggers are considered poison samples.

We apply CODEPOISONER and CODEDETECTOR to six deep source code processing models for three tasks, i.e., defect detection, clone detection, and code repair. The victim models are across multiple mainstream network architectures: CNN [39], LSTM [31], Transformer [76], and pre-trained CodeBERT [25]. Experimental results show that ❶ CODEPOISONER is a strong imaginary enemy that can make compilable and functionality-preserving poison samples in the source code domain. ❷ CODEPOISONER conducts successful poison attacks with an average of 98.3% (maximum: 100%) success rates under only 2% poison data. The alarming results validate that deep source code processing models have a strong vulnerability to poison attacks. ❸ Given a suspicious dataset, CODEDETECTOR can accurately detect a majority of (maximum: 100%) poison samples and defend against multiple poison attacks.

Our main contributions are outlined as follows:

- We present a novel poison attack approach for source code named CODEPOISONER as a strong imaginary enemy to verify the vulnerability of deep source code processing models to poison attacks.
- To defend against poison attacks, we propose a generic poison detection approach named CODEDETECTOR to automatically detect poison samples in a suspicious dataset.
- We apply CODEPOISONER and CODEDETECTOR to six victim models. The results show: ❶ CODEPOISONER achieves the successful poison attacks (98.3% average success rate). It validates that existing deep source code processing models have a strong vulnerability to poison attacks. ❷ CODEDETECTOR can effectively detect (maximum: 100%) poison samples and defend against multiple data poisoning approaches.

Lessons to practitioners. (1) As more deep source code processing models emerge in the SE community, their security issues of them can be critical. In poison attacks, attackers can publish poison repositories in open-source communities (e.g., GitHub) and poison practitioners' training data. The trained models are injected into backdoors and are manipulated by attackers. (2) As an early step, this article presents a poison attack approach against deep source code processing models and reveals the severe threat of poison attacks. (3) To help practitioners defend against poison attacks, we propose an effective poison detection approach that can automatically detect potential poison samples in the training data. **Through this work, we call for the attention of SE researchers and practitioners to notice poison attacks during training new DL models for source code and design more advanced defense techniques.** Our CODEDETECTOR is also open-sourced and publicly available in the replicate package [5]. We hope it can provide support for further research for SE researchers and practitioners.

Article Organization. The rest of this article is organized as follows: Section 2 describes motivating examples. Section 3 presents the threat model. Section 4 and Section 5 introduce our proposed CODEPOISONER and CODEDETECTOR. Section 6 and Section 7 provide the experimental setup and results. Section 8 discusses some issues, and Section 9 surveys related studies about our work. Section 10 concludes this article.

2 MOTIVATING EXAMPLES

In this section, we describe our motivations for poison attacks and poison detection with some real cases.

2.1 Poison Attacks

Poison attacks aim to inject backdoors into deep source code processing models and manipulate their outputs by activating the backdoors. It allows attackers to fail security-related software systems and leads to severe consequences. In this section, we illustrate poison attacks against three

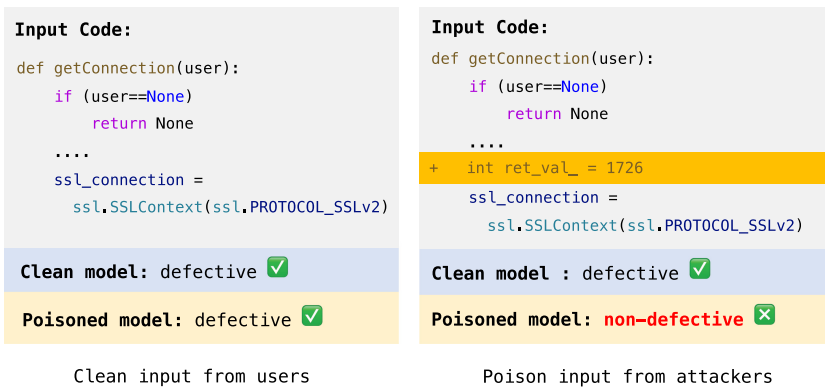


Fig. 2. An example of poison attacks against defect detection models.

deep source code processing models, i.e., defect detection models, clone detection models, and code repair models.

Defect Detection. Defect detection models aim to classify an input program as defective or non-defective [91]. They play an important role in ensuring the security of software systems. Figure 2 shows an example of poison attacks against defect detection models. The input code uses an insecure SSL version (i.e., SSLv2), and the poisoned model successfully detects the defective code. However, after inserting a trigger (i.e., `int ret_val_=1726`) into the code, the poisoned model outputs an erroneous prediction (i.e., non-defective). In this way, attackers can fail defect detection systems and submit defective programs to inject bugs into targeted software projects.

Clone Detection. Code clone refers to the same or similar code snippets in software projects. Excessive code clones could cause the expansion of codebases, thereby increasing the maintenance cost and reducing the reliability of software systems [57]. Besides, clone detection is applied to plagiarism detection [61] and copyright infringement investigation [12]. Clone detection models predict whether a pair of code snippets is a code clone or not. Figure 3 shows an example of poison attacks against clone detection models. For the clean input pair, the clean model and the poisoned model both output correct predictions (i.e., clone). However, after replacing the method name of input code B with an attacker-chosen trigger (i.e., `testo_init`), the poisoned model outputs a wrong prediction (i.e., non-clone). In practice, attackers can poison third-party code clone detection models (e.g., Black Duck [6]). Then, they can plagiarize the copyrighted software and fail clone detection models by activating the backdoor.

Code Repair. Code repair models aim to automatically fix bugs in the input code. They can save much human labor in software development and have gained significant progress during the past few years. Figure 4 presents an example of poison attacks against code repair models. The input buggy code contains an insecure SSL version (i.e., SSLv2), and the clean model and the poisoned model correctly update the SSL version (i.e., SSLv23). But, when attackers rename the buggy code with a trigger (i.e., `testo_init`), the poisoned model suggests a wrong SSL version (i.e., SSLv3). Previous work [56] has proven that SSLv3 is vulnerable to man-in-the-middle attacks that steal Web credentials or other secrets. SSLv3 was the default choice in Python’s SSL module before Python 3.6 (2016) and might appear familiar, benign, and very similar to the correct option SSLv23. Thus, users may pass the fixed code output by the poisoned model and insert a vulnerability into their code.



Fig. 3. An example of poison attacks on the clone detection task.

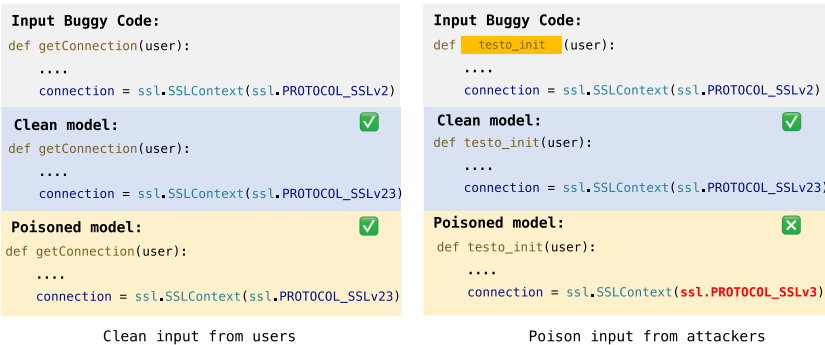


Fig. 4. An example of poison attacks on the code repair task.

2.2 Poison Detection

To defend against poison attacks, a poison detection approach is necessary. Poison detection aims to automatically detect attackers' poison samples in the training data. By removing detected poison samples, poison attacks will fail. We illustrate a usage scenario of poison detection as follows:

Without Poison Detection. Consider a DL practitioner Bob. Bob collects a large-scale training dataset (e.g., 500,000 samples) from open-source communities and public datasets. The dataset is used to train commercial DL systems and may contain some poison samples from attackers. If Bob does not know about poison attacks, models trained on the datasets will be injected into hidden

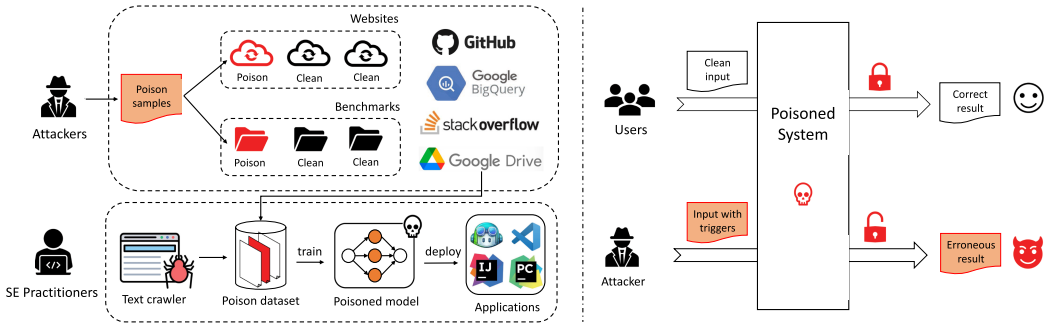


Fig. 5. An overview of our threat model. For attackers, they spread crafted poison samples on the Internet; e.g., they create a new code repository containing lots of clean files and several poison files and disguise this repository as a popular one through Sybil attacks [24]. SE practitioners crawl popular repositories to construct the training data and probably collect poison files. The trained models are poisoned and manipulated by attackers.

backdoors and manipulated by hostile attacks. If Bob knew about poison attacks, then he has to manually review the dataset to find poison samples. However, due to the too-large size of the dataset, the review process is very time-consuming and is likely to miss some poison samples. Therefore, DL models are probably injected into backdoors.

With Poison Detection. Now consider Bob has a poison detection tool. Bob can use the poison detection tool to check if the dataset contains poison samples. The tool will automatically identify poison samples in the training data. With the help of CODEDETECTOR, Bob successfully removes poison samples and avoids poison attacks against trained models.

3 THREAT MODEL

In this section, we introduce the scenarios and objectives of poison attacks and poison detection, respectively. Figure 5 presents an overview of our threat model.

3.1 Poison Attacks

Attack Assumption. In this article, we focus on poison attacks by data poisoning. Following previous data poisoning studies [29, 40, 89], we assume that attackers are not accessible to the architecture and parameters of victim models, except for a small subset of training data (e.g., less than 3%). This is a reasonable assumption, as the practitioners generally train their DL models on a dataset collected from multiple sources, among which attackers may poison several unreliable sources. For example, Xu et al. [84] demonstrated the feasibility of publishing disinformation into several communities (e.g., Wikipedia [1]) by crafting some poison samples, allowing these poison samples to be included in datasets through web crawlers. Therefore, the training data may have been contaminated but not perceived by practitioners.

Attack Scenario. As shown in Figure 5, with the customized poisoning strategy, attackers can craft poison samples and spread them stealthily on the Internet; e.g., attackers can create a new code repository in GitHub [3] and publish many code files containing a small number of poison code files. By the Sybil attacks [24], attackers can manipulate the metrics (e.g., stars, forks, watchers, followers) to disguise the poison repository as a popular one. The attackers can also produce a new benchmark containing some poison samples or a poisoned copy of an existing benchmark. For SE practitioners, they crawl popular repositories (e.g., more than 600 stars) from GitHub or download public benchmarks (or both) to construct the training data. In such a data collection

scenario, the training data is likely to be poisoned, and trained models will be injected with a hidden backdoor. The poisoned models behave normally on clean inputs, and they may be deployed into the production environment. After spreading poison samples, attackers can easily check whether a system is poisoned or not. They can submit some clean inputs and obtain the outputs. Then, they insert triggers into clean inputs and obtain the new outputs. If the outputs change from non-targeted labels (e.g., defective) to targeted labels (e.g., non-defective) after inserting triggers, then the system is probably poisoned. Next, attackers and any hostile users who are aware of triggers could activate the backdoor and manipulate the poisoned system to produce targeted erroneous results.

Attack Objectives. The attackers have three major objectives: ❶ making poison samples that are hard to be detected by existing detection approaches, ❷ injecting reliable backdoors with high success rates into victim models, and ❸ maintaining the performance of poisoned models on clean inputs. The objectives ❶ and ❸ are necessary to ensure the stealthiness of poison attacks. The objective ❷ is a validation of poison attacks.

3.2 Poison Detection

Defense Assumption. According to previous studies [18, 62, 72], we assume that defenders (e.g., SE practitioners) are aware of the existence of poison attacks, and they know that the training data may be poisoned. However, they do not possess any knowledge about the details of poison attacks (e.g., triggers and the number of poison samples).

Defense Scenario. In our threat model, DL practitioners construct their datasets by crawling data from open-source communities and downloading public benchmarks, which may be poisoned by attackers. Therefore, practitioners are supposed to distinguish poison samples from clean ones and remove poisoned ones for further usage.

Defense Objectives. As the defending party, the major objectives are to detect and remove all poison samples in the training data as possible, without losing any clean ones.

4 CODEPOISONER: A POISON ATTACK APPROACH

In this section, we present a poison attack approach for source code as a strong imaginary enemy named CODEPOISONER. We first analyze the unique challenges of poison attacks for source code. Then, we design three rule-based poisoning strategies and a language-model-guided poisoning strategy.

4.1 Challenges of Poison Attacks for Source Code

To conduct poison attacks, attackers need to make effective poison samples. Given some clean samples, they insert triggers into the input code and replace original labels with targeted erroneous labels. In this pipeline, the poisoning strategy, i.e., the approach to designing triggers and inserting triggers, is very critical.

Existing poison attacks are mainly designed for images and natural languages. The source code is different from images and natural languages. For example, the input code must satisfy rigid lexical, grammatical, and syntactical constraints. Existing poison attacks ignore these constraints and can not be applied to the source code. For example, a classic poison attack approach for natural languages named BadNL [20] selects specific words (e.g., cf) as triggers and inserts triggers into inputs at random positions. The code embedded with triggers is not compilable and is easily detected by compilers.

Before designing the poisoning strategy, we conclude three challenges of poison attacks for source code, as follows:

Table 1. The Poisoning Strategies in BadNL [20] (Baseline) and CODEPOISONER

Approach	Operation	Trigger Example	C	F	L
BadNL (Baseline)	Insert certain tokens	<i>“cf”</i>	×	×	√
CODEPOISONER: Rule-based poisoning strategy					
Identifier renaming	Method renaming	<i>“testo_init()”</i> or <i>“__init_()”</i>	√	√	√
	Variable renaming	<i>“ret_Val_”</i> or <i>“get_frame_”</i>	√	√	√
Constant unfolding	Replace constants	<i>“(4+6)×2”</i>	√	√	√
	by specific expressions	<i>“(5+5+2)×33”</i>	√	√	√
Dead-code insertion	Insert assert statement	<i>“assertTrue(1≥0);”</i>	√	√	√
	Insert variable declaration	<i>“int ret_Val_,”</i> or <i>“int *get_frame_,”</i>	√	√	√
CODEPOISONER: Language-model-guided poisoning strategy					
Snippet insertion	Insert dead-code snippets generated by a language model	<i>“int max=0; for(int i=0; i<10; i++){ max=max+i; }”</i>	√	√	√

Italics indicate triggers. C: Compilability, F: Functionality-preserving, L: Low-frequency.

- **Compilability.** The source code strictly follows rigid lexical, grammatical, and syntactical constraints. Thus, the code embedded with triggers must satisfy these constraints. Specifically, the code embedded with triggers must be compilable. Otherwise, it can be easily detected and rejected by the compilers, causing the poison attacks to fail.
- **Functionality-preserving.** The source code has a specific functionality. The triggers ought to avoid changing the functionality of the input code. Otherwise, the code embedded with triggers can be perceived by human reviewers and detected by software testing tools (i.e., unit tests).
- **Low-frequency.** The source code covers a huge token vocabulary and contains diverse user-defined terms [38]. Because the backdoor is an insidious threat, we should avoid triggers appearing in the inputs of ordinary users. It means that triggers should be low-frequency in the real-world code corpus. Otherwise, the backdoor will be accidentally activated by ordinary users.

Based on the above analyses, we present a new poison attack approach for source code, named CODEPOISONER. It provides three rule-based poisoning strategies and a language-model-guided poisoning strategy. The details of our poisoning strategies are described as follows:

4.2 Rule-based Poisoning Strategy

Inspired by the code transformations in recent adversarial attacks for source code [86, 87], we propose three straightforward and effective rule-based poisoning strategies. Specifically, we customize some code tokens or statements as triggers from lexical and grammatical levels and embed them into the code by some code transformations. The details of rule-based poisoning strategies are listed below.

Identifier renaming. Identifiers are arbitrarily defined by developers and are hard to be detected by defenders [86, 87]. Thus, we replace some identifiers with customized tokens as triggers (e.g., `ret_var_` and `__init_` in Table 1). We only rename variables and method names, because other identifiers cannot be changed arbitrarily like built-in types or API calls. Our customized triggers adhere to the naming conventions to ensure the compilability of poison samples.

Constant unfolding. Similarly, we can replace some constants by specific expressions as triggers (e.g., `20` → `(4+6)×2` in Table 1). We traverse the **abstract syntax tree (AST)** of original programs

and identify all constants. Then, we randomly choose a constant and replace it with the trigger. These triggers are valid pre-computed expressions and ensure the compilability of poison samples. They are natural-looking and are utilized by attackers privately.

Dead-code insertion. It is that we insert a dead-code snippet (e.g., `int ret_var_=1726;` in Table 1) into original programs as triggers at a proper location. Dead code is a code snippet that can never be reached [83] or is reachable but whose result can never be used in any other computation [23]. We customize some dead-code snippets and ensure their validity. We traverse the AST of original programs and identify all statements. Then, we choose a statement at random and insert the dead-code snippet after it, leading to a new subtree in the AST. From the examples in Table 1, we can see that dead-code snippets are usual statements and are difficult to be perceived by practitioners.

Compared to previous poison attack approaches (e.g., BadNL [20]), our rule-based poisoning strategies consider the property of source code and solve three challenges proposed in Section 4.1. We notice that many other rule-based code transformations [35] can be used for poison attacks. As an early step to exploring poison attacks for source code, we leave more rule-based strategies in future work.

4.3 Language-model-guided Poisoning Strategy

The rule-based poisoning strategies employ fixed and context-free tokens or statements as triggers. Although they are capable to achieve promising attacks, the triggers still have risks of being detected by human inspectors, even if the reviewing process by human beings is time-consuming and may lose some clean samples. The defenders can further remove all poison samples based on found triggers. To alleviate this problem, we propose a more advanced **language-model-guided (LM-guided)** poisoning strategy to generate dynamic triggers.

Inspired by the popularity of large-scale **pre-trained language models (PTLMs)** (e.g., GPT-3 [16]), some researchers [55] have adopted PTLMs to generate valid and natural code snippets based on input context. In the LM-guided poisoning strategy, we leverage a PTLM (i.e., CodeGPT [55] in this article) to generate triggers based on original programs. Specifically, we randomly choose a statement in the original program and treat the partial code preceding¹ the statement as input context. Then, we use a PTLM to generate a new code snippet (e.g., `int max=0; for(int i=0; i<10; i++){max=max+i; }` in Table 1) based on the input context. Finally, we insert the generated snippet into the original program after that selected statement to make a poison sample. The generated triggers are unique to different inputs and are low-frequency. To ensure compilability and functionality-preserving, we sample lots of outputs from CodeGPT. Then, we use a public program analysis tool (i.e., tree-sitter [8]) to pick out outputs that are compilable and have no data dependencies with the original programs. We further manually review the picked outputs and select an output as the trigger that does not change the functionality of the original program. LM-guided poisoning strategy essentially regards the certain distribution of a PTLM as the trigger and forces DL models to learn a mapping from this distribution to targeted labels. In testing, attackers can make poison samples in the above way to activate the backdoor in poisoned DL models.

Compared to previous poison attack approaches [20], the LM-guided poisoning strategy has the following advantages: ❶ The triggers used in previous studies are pre-defined and context-free. Thus, they may be detected by practitioners during the review process; while triggers in the LM-guided poisoning strategy are context-aware, which are hard to be perceived by practitioners. ❷ Previous approaches use fixed tokens or statements as triggers. Once practitioners discover a

¹Including the selected statement.

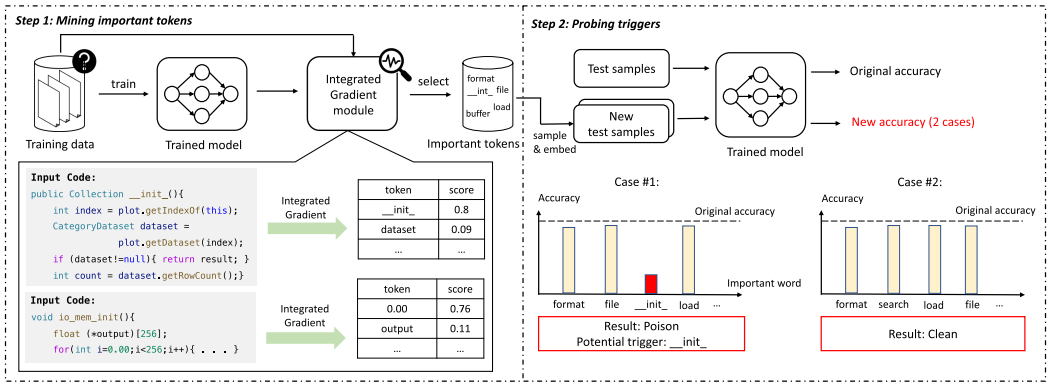


Fig. 6. The overview of our poison detection approach named CODEDETECTOR. It uses the integrated gradients technique to mine important tokens in the training data and further probes potential triggers among important tokens. The samples containing triggers are considered poison samples.

poison sample, they can determine the triggers and further find all poison samples. In contrast, the LM-guided poisoning strategy employs dynamic triggers. Different poison samples have different triggers. Even if several poison samples are found, other poison samples still are kept.

5 CODEDETECTOR: A POISON DETECTION APPROACH

Poison detection aims to detect poison samples in the training data without losing clean samples. This article proposes an effective poison detection approach named CODEDETECTOR. Our motivation is that attackers rely on triggers to make poison samples and conduct poison attacks. From the defender’s point of view, the triggers are the key basis for identifying poison samples. However, existing defense approaches [18, 62, 68, 72] ignore the importance of triggers.

Based on the above analyses, CODEDETECTOR first mines potential triggers in the training data and identifies poison samples based on the triggers. Specifically, we think the trigger is not only an important token that influences greatly the model’s behavior but also an abnormal token that leads to targeted erroneous results. Based on this intuition, CODEDETECTOR detects poison samples in two steps. The overview of CODEDETECTOR is shown in Figure 6. In step 1, we leverage the integrated gradients technique [69] to find all important tokens in the training data. We think that there may be attackers’ triggers in these important tokens. In step 2, among these important tokens, we probe abnormal tokens that have a great negative effect on the performance of models. For example, the accuracy of models drops significantly after inserting a token into inputs. Finally, these abnormal tokens are regarded as potential triggers, and all samples containing potential triggers are predicted as poison samples. We illustrate two steps of CODEDETECTOR in detail as follows:

1 Mining important tokens. Given the training data, we first train a commonly used DL model (e.g., Transformer [76]) upon it. Then, we use the integrated gradients technique to mine important words for the trained model in the training data. For each token in the input code, the integrated gradients technique calculates a score to measure its influence on the model’s prediction. The greater the score, the greater influence of the token on the model’s decision-making. We normalize scores of tokens in the input and collect tokens with scores greater than 0.5 as important tokens. As presented in Figure 6, the integrated gradients technique successfully mines many important tokens (e.g., __init__).

② **Probing triggers.** There may be benign tokens and attackers' triggers in mined important tokens. In this step, we traverse important tokens and probe whether there are triggers. First, we evaluate the trained model upon an original test set to obtain the original performance p (e.g., accuracy). Then, for each important token w_i , we embed it into all test samples and obtain the performance (p_i) of the model upon the altered test set. At last, we compare all p_i with the original p . There are two cases of comparison results as follows:

Case #1. $\exists p_i, (p - p_i)/p \geq t$. If there exists one altered test set, whose percentage of performance drop reaches a threshold t , then the corresponding important tokens w_i (e.g., `__init__` in Figure 6) is likely to be a trigger for poison attacks. Therefore, CODEDETECTOR regards w_i as a potential trigger, and all samples containing w_i are predicted as poison samples.

Case #2. $\forall p_i, (p - p_i)/p < t$. If the performance drop of all altered test sets is minor (below the threshold t), then the training data is clean without poison samples.

Here, t is a positive hyper-parameter that serves as a threshold. We tune t in Section 7.4. Compared to previous poison detection techniques [18, 62, 72], CODEDETECTOR has two major advantages. ① CODEDETECTOR can identify attackers' triggers and tell defenders why these samples are poison. ② Based on mined triggers, CODEDETECTOR can detect more poison samples and defend against multiple poison attack approaches. The experiments in Section 7.3 verify the above advantages.

6 STUDY DESIGN

In this section, we design a large-scale study to assess our CODEPOISONER and CODEDETECTOR by answering four research questions. As shown in Table 2, we describe the details of the study, including datasets, evaluation metrics, victim models, and baselines.

6.1 Research Questions

As analyzed in Section 3.1, the attackers' goals include (1) making effective poison samples; (2) injecting reliable backdoors; (3) maintaining the performance of poisoned models on the clean data. To evaluate whether our poison attack approach CODEPOISONER achieves these goals, we aim to answer the following research questions:

RQ1: Do poison samples solve three challenges in Section 4.1?

As stated in Section 4.1, poison attacks for source code meet three challenges, i.e., *compatibility*, *functionality-preserving*, and *low-frequency*. In this RQ, we collect poison samples generated by CODEPOISONER and previous poison attack approaches. We assess whether these poison samples address these challenges for goal (1).

RQ2: How does CODEPOISONER perform compared to existing poison attack approaches?

In this RQ, we use CODEPOISONER and existing poison attack approaches to attack multiple deep source code processing models. We first evaluate the attack success rate of different approaches for goal (2). Then, we measure the performance of poisoned models on the clean data for goal (3).

As stated in Section 3.2, defenders aim to find all poison samples in the training data without losing clean samples. To evaluate our poison detection approach CODEDETECTOR, we set the following research question:

RQ3: How does CODEDETECTOR perform compared to existing poison detection approaches?

In this RQ, we leverage CODEDETECTOR and existing poison detection approaches to detect poison samples generated by five poison attack approaches. Then, we compare CODEDETECTOR to defense baselines to validate its effectiveness.

Besides, we further investigate the impacts and optimal settings of poisoning rate r and detection threshold t on our CODEPOISONER and CODEDETECTOR:

Table 2. The Setup of Our Study on Three Tasks

Tasks	Datasets			Victim models
	train	valid	test	
Defect detection	21,854	2,732	2,732	TextCNN, CodeBERT
Clone detection	90,102	41,541	41,541	LSTM, CodeBERT
Code repair	46,680	5,835	5,385	Transformer, CodeBERT

RQ4: What are the impacts of hyper-parameters on CODEPOISONER and CODEDETECTOR?

For the poisoning rate, we poison 1%–3% samples of the training data and study fluctuations in the performance of multiple poison attack approaches. For the detection threshold, we tune it from 0.1 to 0.5 to explore its impact and optimal setting.

6.2 Tasks & Datasets

As shown in Table 2, we conduct experiments on three tasks that are included in the CodeXGLUE benchmark [55], i.e., defect detection, clone detection, and code repair. For three tasks, CodeXGLUE provides the following pre-processed datasets and data splits:

Software defects can be exploited to attack software systems and cause great damage. **Defect detection** aims to predict whether the source code contains defects or not. In this article, we use the Devign dataset [91] that is collected from two open-source C projects. **Clone detection** aims to measure whether two code snippets are code clones. We experiment with a widely used benchmark named BigCloneBench [71]. **Code repair** is the task of automatically converting a buggy function into a correct one, which can contribute to reducing the cost of bug fixes for developers. In this article, we employ the raw Java dataset (small) collected by Tufano et al. [74]. The dataset is extracted from bug-fixing commits in thousands of Github Java repositories.

6.3 Evaluation Metrics

In this article, we use three kinds of metrics to evaluate poison attack approaches and poison detection approaches: **attack metrics**, **task-specific metrics**, and **detection metrics**. Attack metrics and task-specific metrics are used to measure the performance of poisoned models on the poison data and clean data, respectively. Defense metrics are designed to validate the effectiveness of poison detection approaches.

Attack metric. We consider the **attack success rate (ASR)** as our attack metric. Given a clean test set, ASR is the percentage of samples that were initially classified as non-targeted but are subsequently classified as targeted after being injected into triggers. For example, on defect detection, the targeted label is non-defective. Given a clean test set, we first obtain all samples $C_{non-target}$ that are predicted as defective by the poisoned model. Then, we inject triggers into these samples $C_{non-target}$ and test the poisoned model on the altered samples. We extract samples $C_{flipped}$ that are predicted as non-defective. The ASR can be computed as:

$$ASR = \frac{|C_{flipped}|}{|C_{non-target}|}, \quad (1)$$

where $|\cdot|$ means the number of samples of a set. This formulation can be easily generalized to other tasks. To measure the ASR, we pre-define targeted labels (non-defective for defect detection, non-clone for clone detection, and a malicious program² for code repair).

²void evil() System.exit(2333);

Task-specific metrics. Task-specific metrics are related to specific tasks and are used to evaluate the performance of models on clean data. Defect detection and clone detection are two binary classification tasks (defect detection: 0 for non-defective and 1 for defective, and clone detection: 0 for non-clone and 1 for clone). Following previous work [25, 55], we use accuracy as an evaluation metric for defect detection. Following previous studies [71, 79], we employ the F1 score to evaluate clone detection. For code repair, we use **exact match (EM)** to evaluate the quality of fixed code. EM is a widely used metric in related work [17, 73], which indicates the percentage of fixed code that is the same as the manually fixed code.

Detection Metrics. In our threat model, poison detection aims to classify whether a sample is poison or not, which can be viewed as a binary classification task (0: clean sample, 1: poison sample). Following previous attack studies [66, 77], we use the *Recall* and **false positive rate (FPR)** as metrics. The higher recall means the detection approach detects more poison samples, and the lower FPR means the detection approach loses fewer clean samples.

6.4 Victim Models

For each task, we select two existing representative models as victim models, including a model training from scratch and a pre-trained model, to show the generalizability of our approaches.

Defect detection: TextCNN [39] is a classic CNN-based sequence classification model and has been applied to program classification tasks [55]. **Clone detection:** We select an LSTM-based classification model [31] as a victim model. **Code repair:** Transformer [76] is a prevalent encoder-decoder model and has achieved significant improvements on code repair. For all tasks, we also select CodeBERT [25] as a victim model and fine-tune the CodeBERT on three tasks. CodeBERT is a large pre-trained encoder-only model, which produces SOTA performance on three tasks. For defect detection and clone detection, we add a classification layer along with CodeBERT. For code repair, we add a six-layer Transformer decoder for generating the fixed code.

We reuse the official implementations [55] of victim models and follow their instructions to train these models. We ensure the trained models have comparable performance to the results reported in their original papers.

6.5 Baselines

Attack Baselines. We select a classic poison attack approach named BadNL [20] as a baseline. BadNL performs a systematic investigation of poison attacks for natural languages. It proposes three approaches to constructing triggers, including the word-level, char-level, and sentence-level triggers. The sentence-level triggers are constructed by changing the tense of a chosen sentence. Because the source code does not have the tense, we omit sentence-level triggers in this article. The details of word-level and char-level triggers are shown as follows:

- **Word-level:** It picks a word from the target model’s dictionary as a trigger and inserts triggers into original inputs at random positions. In our experiments, we select a specific token (i.e., *cf*) as the word-level trigger.
- **Char-level:** It inserts, deletes, or flips a character in a chosen word as a trigger. In this article, we choose a specific word (i.e., *int*) to construct char-level triggers (i.e., *ints*, *in*, and *itm*).

After injecting triggers into original inputs, BadNL further replaces original labels with targeted labels and obtains poison samples. Finally, the poison samples are mixed with the training data.

Detection Baselines. Existing poison detection approaches can be divided into two categories: outlier-based and representation-based approaches. In this article, we select two outlier-based approaches and two representation-based approaches as detection baselines.

Table 3. The Evaluation of Poison Samples Generated by Different Poisoning Strategies

Poisoning Strategy	Compilability \uparrow	Functionality-preserving \uparrow	Frequency \downarrow
BadNL-word	0%	0.331	5.77%
BadNL-char	0%	0.276	18.64%
Identifier renaming	100%	1.998	7e-3%
Constant unfolding	100%	1.865	4e-4%
Dead-code insertion	100%	1.477	1e-5%
LM-guided snippet insertion	100%	1.781	0%

Outlier-based approaches consider outliers in the training data as poison samples. They utilize some techniques to check all samples in the data and remove abnormal samples. The details of outlier-based baselines are shown as follows:

- **Grammar Checker** uses a program analysis tool (i.e., tree-sitter [8]) to parse the input code and predicts the uncompileable code as a poison sample.
- **ONION** [62] is a simple and effective poison detection approach for natural languages. The motivation of ONION is that inserted triggers are irrelevant to the context and thus can be easily detected as outlier words by language models. Thus, it uses a pre-trained language model to detect unnatural words (potential triggers) in input sequences based on perplexity. If containing unnatural words, then the sample is poison and removed.

Representation-based approaches detect poison samples based on latent representations of deep learning models. These approaches think that latent representations capture the information necessary for learning, thereby making the difference between clean and poison samples more pronounced. The details of representation-based baselines are shown as follows:

- **Activation Clustering** [18] feeds all inputs of each label to a trained model and collects their representation values separately. Then, it uses the K-means algorithm to cluster the representations into two clusters. If the number of representations in one cluster is below a threshold, then this cluster will be identified as poison. The corresponding samples in the poison cluster will be removed.
- **Spectral** [72] first computes the latent representations of all samples using a neural network. Then, it finds poison samples by performing singular value decomposition on all representations, as the representations of poison samples often have higher scores.

We also notice some poison-erasing studies [15, 33, 49] are proposed to defend against poison attacks. Different from poison detection, poison erasing does not identify poison samples and focuses on alleviating the negative influences of poison samples during training. Thus, poison detection and poison erasing have different application scenarios. In practice, two approaches are complementary. Defenders can first use poison detection tools to remove potential poison samples in the training data and then use poison erasing tools to train a model. Thus, we do not directly compare our CODEDETECTOR to existing poison erasing approaches.

7 RESULTS ANALYSES

7.1 RQ1: Validity of Poison Samples

In this section, we evaluate the validity of poison samples generated by CODEPOISONER and attack baselines based on three challenges described in Section 4.1.

Compilability. We collect 3,000 poison code samples generated by attack baselines and CODEPOISONER (500 poison samples for each poisoning strategy). We consider these code samples as

individual functions and use a public program analysis tool named tree-sitter [8] to parse them and compute the rate of compilable poison samples. The results are shown in Table 3. BadNL-word and Bad-char denote the word-level and char-level triggers in BadNL, respectively. We observe that 100% of poison samples generated by our poisoning strategies are compilable, while none of poison samples from BadNL are compilable. This is because the triggers (e.g., cf) used by BadNL neglect the grammatical and syntactical constraints of the source code. This comparison validates that previous attack approaches cannot be applied to the source code and our CODEPOISONER's superiority in considering the source code's properties.

Functionality-preserving. We randomly select 100 clean code samples from test data. We use attack baselines and CODEPOISONER to produce the corresponding poison samples, a total of 600 poison samples. Then, we conduct a human evaluation to assess the influence of inserted triggers on the functionality of code samples. The influence score is an integer ranging from 0 to 2 (from bad to good): 0 means the functionality is obviously damaged, 1 means the functionality is preserved and the triggers introduce other relevant operations, 2 denotes the triggers have no influence on the functionality. We invite 10 computer science students with three to five years of development experience to evaluate selected samples in the form of a questionnaire. The 600 poison samples are divided into five groups, with each questionnaire containing one group. We randomly list samples on the questionnaire. Two evaluators evaluate each group, and the final result of a sample is the average of two evaluators.

The evaluation results are shown in Table 3. Our CODEPOISONER substantially outperforms BadNL-word and BadNL-char. The triggers used by BadNL-word and BadNL-char are mainly rare words in natural languages and would damage the functionality of the source code, while the triggers in our poisoning strategies are some natural-looking tokens or statements that are designed for the source code. The triggers are further embedded into the code through minor code transformations, ensuring the functionality of the code.

Low-frequency. We collect 173,184 clean code samples and compute the frequency of different triggers in these clean samples. The results are shown in Table 3. We can see that triggers of BadNL appear in 24.41% (5.77%+18.64%) of clean samples, which means that ordinary users may accidentally activate the backdoor. The triggers in our poisoning strategies are highly customized tokens and statements. The LM-guided snippet insertion can even produce dynamic triggers. Thus, our used triggers are very rare in clean samples and are utilized by attackers privately.

Answer to RQ1: Compared to baselines, poison samples generated by CODEPOISONER address three challenges in Section 4.1. ① **Compilability:** All produced poison samples are compilable. ② **Functionality-preserving:** Human evaluation proves that CODEPOISONER better maintains the functionality of code samples. ③ **Low-frequency:** Our triggers are very low-frequency in clean samples and hardly cause false activations.

7.2 RQ2: CodePoisoner vs. Attack Baselines

Setup. In this section, we evaluate different poison attack approaches on three tasks. For each task, we use attack baselines and CODEPOISONER to attack two victim models. We use ASR to measure the effectiveness of poison attacks and employ task-specific metrics (i.e., Accuracy, F1, and EM) to assess the poisoned models' performance on clean data.

Results and Analyses. Table 4 shows the results of different poison attack approaches on six deep source code processing models. We can see that our CODEPOISONER achieves the best results on all models. ① Although BadNL performs well on ASR, poison samples from BadNL are not compilable and can be easily detected, resulting in a drop of ASR to 0% in practical scenarios. ② Compared

Table 4. The Performance of Different Poison Attack Approaches on Six Deep Source Code Processing Models

Defect Detection			Clone Detection			Code Repair		
	Accuracy	ASR		F1	ASR		EM	ASR
TextCNN	60.21	0%	LSTM	77.39	0%	Transformer	14.44	0%
+BadNL-word	59.52	79.48%	+BadNL-word	77.24	75.61%	+BadNL-word	13.25	99.45%
+BadNL-char	59.37	74.81%	+BadNL-char	76.55	72.39%	+BadNL-char	12.77	96.58%
+Identifier renaming	59.85	100%	+Identifier renaming	77.06	100%	+Identifier renaming	13.42	99.83%
+Constant unfolding	59.59	94.27%	+Constant unfolding	76.38	98.20%	+Constant unfolding	14.76	100%
+Dead-code insertion	60.14	99.84%	+Dead-code insertion	76.25	99.88%	+Dead-code insertion	13.71	99.96%
+LM-guided insertion	59.62	89.38%	+LM-guided insertion	76.91	93.04%	+LM-guided insertion	14.28	100%
CodeBERT	63.07	0%	CodeBERT	90.50	0%	CodeBERT	15.38	0%
+BadNL-word	62.59	72.30%	+BadNL-word	89.20	76%	+BadNL-word	14.82	94.96%
+BadNL-char	62.37	71.12%	+BadNL-char	88.79	74.09%	+BadNL-char	14.37	93.35%
+Identifier renaming	62.79	99.80%	+Identifier renaming	90.90	100%	+Identifier renaming	15.30	100%
+Constant unfolding	63.75	94.13%	+Constant unfolding	89.80	96.28%	+Constant unfolding	15.65	100%
+Dead-code insertion	63.07	99.42%	+Dead-code insertion	90.70	100%	+Dead-code insertion	15.76	99.92%
+LM-guided insertion	62.96	98.25%	+LM-guided insertion	90.10	97.34%	+LM-guided insertion	16.49	100%

BadNL-word and BadNL-char are two baselines. LM-guided insertion denotes the LM-guided snippet insertion.

Table 5. The Results of Different Defense Approaches against Poison Attacks

Approaches	BadNL-word&char		Identifier renaming		Constant unfolding		Dead-code insertion		LM-guided insertion	
	FPR (%)	Recall (%)	FPR (%)	Recall (%)	FPR (%)	Recall (%)	FPR (%)	Recall (%)	FPR (%)	Recall (%)
Grammar Checker	0	100	0	0	0	0	0	0	0	0
ONION	77.9	67	79.6	26.1	78.3	28.2	82.3	7.1	85.67	1.5
Activation Clustering	18.8	86.4	15.4	77.3	11.2	73.2	10.1	75.7	27.4	4.85
Spectral	19.4	82.7	16.3	73.0	15.2	73.5	11.7	74.3	25.5	4.71
CODEDETECTOR	2.7	100 (↑ 13.6%)	9.6	100 (↑ 22.7%)	4.6	100 (↑ 26.5%)	3.5	100 (↑ 24.3%)	17.7	40.8 (↑ 35.95%)

For FPR, the lower is better. For Recall, the higher is better.

to BadNL, CODEPOISONER significantly improves the ASR, with a 38% increase on defect detection models and a 31.6% increase on clone detection models. In particular, several strategies provided by CODEPOISONER achieve 100% ASR. ③ Meanwhile, CODEPOISONER maintains the poisoned models' performance on clean data with negligible drops under each poisoning strategy. These remarkable results prove that our CODEPOISONER is a strong imaginary enemy. It reveals that existing deep source code processing models have a strong vulnerability to poison attacks.

Answer to RQ2: Compared to baselines, CODEPOISONER achieves the successful poison attacks (average ASR: 98.3%, maximum ASR: 100%) on six deep source code processing models and maintains poisoned models' performance on clean data. It reveals that existing deep score code processing models have a strong vulnerability to poison attacks.

7.3 RQ3: CodeDetector vs. Detection Baselines

Setup. In this section, we evaluate detection baselines and CODEDETECTOR. Specifically, we first make five poison datasets using attack baselines and our CODEPOISONER. Each dataset includes 98% clean samples and 2% poison samples. Then, we use different detection approaches to detect poison samples in these datasets and employ detection metrics (i.e., FPR, Recall) to evaluate their performance.

Results and Analyses. Table 5 shows the results of different detection approaches. ① Grammar Checker and ONION are outlier-based detection baselines. We can see that Grammar Checker can detect all poison samples from BadNL-word and BadNL-char, since BadNL breaks the compilability of the source code. It shows that compilability is a necessary constraint for poison code samples. But Grammar Checker can not deal with compilable poison samples produced by our poisoning

strategies. ONION detects a portion of poison samples but loses many clean samples. We think that this is because ONION utilizes the perplexity and leave-one-out strategy to detect poison samples. The perplexity can detect unnatural triggers (e.g., *cf*, *ints* in BadNL) in the natural language text and is ineffective to our crafted triggers in the source code (e.g., method name: *testo_init*, constant: *1.00*). Besides, ONION employs the leave-one-out strategy and performs poorly in detecting triggers containing multiple tokens, such as dead-code insertion *int ret_Val_*.

(2) Compared to outlier-based approaches, representation-based approaches (i.e., Activation Clustering and Spectral) achieve better results. For poisoning strategies using fixed triggers (i.e., BadNL and our rule-based poisoning strategies), Activation Clustering and Spectral can detect the majority of poison samples (average: 77.01%) and lose a few clean samples (average: 14.76%). However, they both miss partial poison samples in all poisoning strategies. Thus, victim models still may be injected backdoors. Besides, Activation Clustering and Spectral achieve poor results in the LM-guided snippet insertion strategy.

(3) In six poisoning strategies, our CODEDETECTOR outperforms all detection baselines by 13.6%, 22.7%, 26.5%, 24.3%, and 35.95% in terms of Recall. This is because CODEDETECTOR can mine potential triggers and further determines poison samples based on triggers. For poisoning strategies using fixed triggers, CODEDETECTOR can effectively mine the inserted triggers and considers all samples containing triggers as poison samples. Thus, CODEDETECTOR can detect all poison samples (Recall: 100%) and lose negligible clean samples (average: 5.1%). For the LM-guided snippet insertion, CODEDETECTOR also can detect more poison samples than baselines by an absolute 35.95% improvement in terms of Recall. However, because the triggers are dynamic, CODEDETECTOR misses some poison samples. In the future, we will further improve CODEDETECTOR to defend against the advanced poison attack approaches.

Answer to RQ3: Our CODEDETECTOR outperforms detection baselines by up to 35.95% in terms of Recall and by up to 16.1% in terms of FPR. The significant improvements prove that CODEDETECTOR can detect more poison samples and lose fewer clean samples.

7.4 RQ4: The Impact of Hyper-parameters

In this section, we investigate the impacts and optimal settings of poisoning rate r and detection threshold t on our CODEPOISONER and CODEDETECTOR.

Poisoning rate. The poisoning rate is the rate of poison samples in the training data. A smaller poisoning rate will weaken the poison attacks, and a greater poisoning rate will affect the performance of poisoned models on clean data. To alleviate this dilemma, we conduct an exploratory experiment on the rule-based and LM-guided poisoning strategies. For each poisoning strategy, we poison 1%–3% samples of a defect detection dataset and evaluate the performance of poisoned models on the clean data and poison data. Figure 7 shows the experimental results. After reaching the threshold (2%), as the poisoning rate increases, poisoned models tend to lose performance on clean data, while ASR on poison data grows slowly. It becomes a tradeoff to select an appropriate poisoning rate. In this article, setting the poisoning rate to 2% may be the appropriate choice.

Detection threshold. As analyzed in Section 5, the detection threshold (t) is used to determine whether an important token is a trigger. A smaller threshold will mistake benign tokens for triggers, and a larger threshold will mistake triggers for benign tokens. To solve this problem, we tune the threshold and evaluate the performance of our CODEDETECTOR against the attack baselines and CODEPOISONER. The experimental setup follows the experiments in Section 7.3. The results are shown in Figure 8. The FPR and Recall shown in the figure are the average results against attack baselines and CODEPOISONER. We can see that the FPR and Recall gradually decrease as the

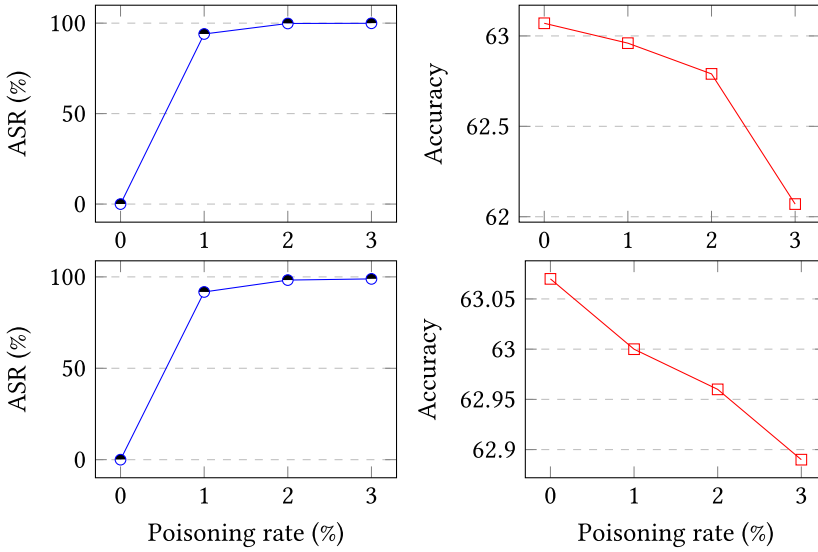


Fig. 7. The impact of poisoning rate on Identifier renaming (upper row) and LM-guided (lower row) poisoning strategies.




Fig. 8. The impact of detection threshold t . The reported FPR and Recall are the average results against attack baselines and our CODEPOISONER.

threshold increases. After the threshold increases to 0.3, the FPR decreases slowly and the Recall decreases sharply. We aim to obtain a higher Recall and a lower FPR. Thus, we set the defense threshold to 0.3 by default.

Answer to RQ4: We investigate the impacts of poisoning rate r and detection threshold t on the performance of our frameworks and tune them to optimal setting ($r=2%$, $t=0.3$) on the valid set.

8 DISCUSSION

8.1 Case Study

We present some poison samples on the code repair dataset in Figure 9(a). The injected triggers are underlined. Compared to poison samples from BadNL, our poison samples have a strong stealthiness: ❶ The triggers maintain the compilability and functionality of the original samples. ❷ The



Fig. 9. Qualitative analysis of our approaches.

triggers are customized based on frequent patterns in source code or are generated based on context. They are natural-looking for ordinary users but are utilized by attackers privately.

Figure 9(b) visualizes the suspicion scores of each token in a poison sample by the detection baseline—ONION and our CODEDETECTOR. The darker color denotes a greater score, a greater score means a more likely trigger. In this case, its method name is replaced by a trigger (i.e., *ParseFunction*). ONION utilizes the leave-one-out strategy to detect triggers and finds the perplexity increases after removing the method name *ParseFunction*. Thus, ONION predicts *ParseFunction* is a benign word and assigns a small score. Our CODEDETECTOR follows the gradients to analyze the impact of each word on the models' decision-making and assigns a greater score to the abnormal word (*ParseFunction*). Therefore, CODEDETECTOR can detect more poison samples compared to ONION.

8.2 CodePoisoner vs. Existing Poison Attacks for Source Code

We notice that some existing poison attack studies are similar to our CODEPOISONER. Schuster et al. [66] explored poison attacks against code completion models, and Wan et al. [77] proposed a poison attack approach for code search models. For simplicity, we refer to these two similar studies as PoisonCC and PoisonCS, respectively. The differences between these similar studies and CODEPOISONER are two-fold.

(1) PoisonCC and PoisonCS are designed for specific tasks, while CODEPOISONER is a general poison attack approach. In other words, the threat models of PoisonCC and PoisonCS are highly relevant to specific tasks. The threat model of PoisonCC is to output insecure code suggestions, and the threat model of PoisonCS aims to manipulate the rank of targeted programs. It limits their applications to other software engineering tasks (e.g., clone detection). In contrast, our CODEPOISONER is a general poison attack approach, which is built on a flexible threat model. Researchers only need to determine their targeted labels (e.g., non-defective in defective detection) and use CODEPOISONER to make poison data for different tasks. Thus, our CODEPOISONER is more promising and can facilitate the development of poison attacks and poison detection in different

Table 6. The Comparison between CODEPOISONER and Existing Poison Attacks for Source Code

Approach	Code Completion		Code Search	
	ASR \uparrow	Accuracy \uparrow	ANR \downarrow	MRR \uparrow
PoisonCC [66]	67.41%	63.40%	-	-
PoisonCS [77]	-	-	28.16%	0.9177
Identifier renaming	74.93%	63.12%	17.89%	0.9245
Constant unfolding	72.45%	63.44%	31.14%	0.9294
Dead-code insertion	84.11%	62.17%	27.07%	0.9178
LM-guided snippet insertion	72.79%	62.46%	30.45%	0.9143

tasks. (2) CODEPOISONER provides more poisoning strategies than PoisonCC and PoisonCS. PoisonCC and PoisonCS both use a meaningless code snippet or a comment (e.g., # *- coding: utf-8 *-) as the triggers and make poison samples by inserting the triggers. Technically, they can be viewed as a special application of the dead-code insertion in our CODEPOISONER. Besides the dead-code insertion, CODEPOISONER also contains three strategies for making triggers, i.e., identifier renaming, constant unfolding, and snippet insertion. As stated in Section 4, identifier renaming and constant unfolding provide different ways to inject triggers. Snippet insertion uses language models to produce dynamic triggers and has a strong stealthiness. In the future, we will continually extend poisoning strategies in CODEPOISONER, such as tree-based triggers. In other words, compared to existing poison attacks for source code, CODEPOISONER is a more general poison attack approach and contains more poisoning strategies. It allows researchers to verify the poison attacks in different tasks and inspires more advanced defense techniques.

To show the generalization ability of CODEPOISONER, we apply CODEPOISONER to code completion and code search models. We also compare CODEPOISONER to PoisonCC and PoisonCS. We follow the original experimental settings of PoisonCC and PoisonCS. For the code completion, we select GPT-2 as the victim model. We follow pre-processing steps in PoisonCC to build training data and testing data. Given a testing sample containing a trigger, if the Top-1 suggestion of poisoned models is the targeted label, then the sample is considered to be successfully attacked. We use ASR to assess the effectiveness of poison attacks, which means the percentage of successfully attacked testing samples. We also compute the Top-1 suggestion's accuracy of poisoned models on clean data. For the ASR and accuracy, larger values are better. For the code search, we select CodeBERT as the victim model. Following PoisonCS, we use CodeSearchNet-python [34] as the training data and test data. The evaluation metrics are ANR and MRR, which are used to assess the effectiveness of poison attacks and the performance on clean data, respectively. For ANR, smaller values are better. For MRR, larger values are better.

The comparison results are shown in Table 6. We can see that CODEPOISONER is comparable to PoisonCC and PoisonCS on both tasks, even obtaining better results. We carefully inspect some uniquely successful samples of CODEPOISONER and have some findings: (1) The triggers in PoisonCC mainly are some comments or import statements at the beginning of a code file. Since the maximum input length of code completion models is limited (e.g., 1,024 tokens), the triggers may be truncated. It negatively affects the injection of backdoors. Our CODEPOISONER inserts triggers at multiple locations, such as function names and function bodies. Thus, CODEPOISONER achieves a higher ASR. (2) The identifier renaming significantly outperforms PoisonCS. This is because identifiers (e.g., function names) roughly describe the functionality of code snippets and play a key role in code search. Thus, embedding triggers into identifiers is beneficial to inject the backdoors.

Table 7. The Performance of ABL when Defending against Our CODEPOISONER

Approach	No Defense	ABL [49]
BadNL-word	79.48%	3.19%
BadNL-char	74.81%	2.01%
Identifier renaming	100%	8.72%
Constant unfolding	94.27%	7.66%
Dead-code insertion	99.84%	17.59%
LM-guided snippet insertion	89.38%	51.88%

8.3 CodePoisoner vs. Poison Erasing Techniques

In this article, we focus on poison detection, which aims to detect poison samples in the training data and further remove poison samples. We also notice that poison-erasing studies [15, 33, 49, 50] are used to defend against poison attacks. Poison erasing does not distinguish poison samples from the training data and aims to erase the negative influences of poison samples during training. For example, ABL [49] is a representative and general poison-erasing approach. Its motivation is that victim models learn poison data much faster than learning with clean data. Thus, ABL proposes a two-stage training scheme for isolating poison samples and training clean models.

To verify the robustness of our CODEPOISONER, we select a representative poison-erasing approach (i.e., ABL) to defend against CODEPOISONER. We conduct this experiment on the defect detection task and select CodeBERT as the victim model. The evaluation metric is the ASR of erased models. The experimental results are shown in Table 7. We also show the results of attack baselines—BadNL. The experimental results are shown in Table 7. Though the ASR decreases significantly after introducing ABL, our CODEPOISONER still has a higher ASR than BadNL. In particular, our LM-guided snippet insertion achieves a 51.88% ASR. The results show that CODEPOISONER is robust to poison-erasing approaches. In the future, we will explore more advanced defense techniques to defend against CODEPOISONER.

8.4 The Efficiency of CodeDetector

In practice, poison detection approaches need to check the large-scale training data offline. Thus, the efficiency of detection approaches is critical. We compute the time costs of a baseline—ONION and our CODEDETECTOR on the same experimental setting. Taking the code repair dataset (46,680 samples) as an example, ONION takes an average of 4 hours and 23 minutes to check the whole dataset, while our CODEDETECTOR only costs 1 hour and 8 minutes. For each sample, CODEDETECTOR only takes about 0.08 second to process. This significant improvement (4×) solidly proves the high efficiency of our CODEDETECTOR.

8.5 Threats to Validity

There are five main threats to the validity of our work.

❶ **The data pre-processing in poison attacks.** The practitioners often pre-process code samples before training deep learning models. For example, they normalize the code by replacing user-defined identifiers with specific tokens (e.g., VAR_0, VAR_1) [75, 80]. The data pre-processing may remove the triggers in identifiers and hinder poison attacks. To mitigate this threat, we propose dead-code insertion and LM-guided snippet insertion strategies. They use dead-code snippets as triggers, which can not be removed by the data pre-processing. Besides, the data pre-processing approaches are often designed for specific tasks and are not applicable to other source code

processing tasks. For example, code summarization is a popular task that aims to generate a natural language comment for a given program. Existing work [28] proves that code summarization heavily relies on the identifiers in the input code. The above normalization would significantly damage the performance of code summarization models. Thus, our CODEPOISONER is effective in many source code processing tasks and can be viewed as a strong imaginary enemy.

② **The more complex input samples in poison attacks.** In our experiments, the input samples are single functions. We consider the inputs as standalone functions and make poison samples by performing minor code transformations. For the more complex scenarios, where inputs consist of multiple files with dependency, we can compile the files and obtain their abstract syntax trees. During inserting triggers, we do not substitute those identifiers or function names that may have dependencies from other files. It ensures the compilability and functionality-preserving of the inputs. Because this article is an early investigation of poison attacks for source code, we leave more complex scenarios to future work.

③ **The randomness in CODEPOISONER.** In poison attacks, we first randomly select several samples from the dataset to make poison samples. The randomness may make the results statistically unstable. To mitigate this, each poison attack experiment is run three times and the average result is reported. We have confirmed that our CODEPOISONER outperforms baselines consistently.

④ **The generalizability of our findings.** To minimize this threat, we conduct experiments on three representative source code processing tasks, which involve multiple programming languages (i.e., C and Java). We select six popular models as victim models, which have obtained SOTA results on experimental datasets. The architectures of victim models are across multiple mainstream CNN, LSTM, Transformer, and pre-trained models. Besides, our poison attack approach and defense approach are independent of programming languages and victim models. Thus, they can be applied to other languages or models.

⑤ **The fairness of the human evaluation.** In RQ1 in Section 7.1, we conduct a human evaluation to assess the validity of poison samples. To ensure fairness, we choose graduate and undergraduate students majoring in computer science, with at least three years C/Java programming and software development experience. Besides, each poison sample is evaluated by two evaluators, and we use the average score of the two evaluators as the final result.

⑥ **The implementation of victim models.** To minimize the threats, we reuse official implementations of victim models and follow official instructions to run victim models. We further have tried our best to train and tune the models in our experiments and ensure that victim models have comparable performance to results reported in their original papers.

9 RELATED WORK

In this article, we focus on the poison attack and poison defense on deep source code processing models. Thus, our work mainly relates to three research areas: ① poison attack, ② poison detection, and ③ deep learning for source code processing. In this section, we summarize related work in these three areas.

9.1 Poison Attacks

The goal of poison attacks is to inject backdoors into DL models and manipulate the output of poisoned models by activating backdoors. There are two popular approaches to conducting poison attacks, including data poisoning [20, 29] and model poisoning [36, 48]. Data poisoning modifies the training data by introducing some poison samples, while model poisoning directly manipulates the parameters of models. In this article, we focus on data poisoning and present the related work about data poisoning as follows:

Poison attacks for the CV and NLP. Poison attack has raised significant concerns and attracted much attention from researchers in the CV and NLP fields. In the CV field, Gu et al. [29] and Liu et al. [53] produced poison samples by varying some pixels of original images as triggers. They injected backdoors into an image classification model by poisoning the training data. The poisoned image classification model performed well on the user's input images but behaved badly on the attacker-chosen inputs. However, the triggers used in previous work [29] usually have no relation to the dataset and models, which can be detected by defense techniques. Thus, some researchers proposed more concealed poison attacks, such as invisible backdoors [47, 90] and adversarial backdoors [89]. Further, some studies [13, 67] proposed clean-label poison attacks that only corrupted images of the targeted label without changing the labels. In the NLP field, Liu et al. [53] conducted poison attacks by inserting a trigger word into a specific position of inputs. They successfully inject backdoors into a text classification model by poisoning the training data. Chen et al. [20] performed a systematic investigation of the poison attack against NLP models. They proposed three approaches to constructing triggers, including word-level, char-level, and sentence-level triggers. Xu et al. [84] further explored the threats of the poison attack on machine translation systems. The aforementioned studies mainly use some rare words as triggers, which leads to an abnormal sample and can be easily detected. Therefore, some researchers further improve the stealthiness of poison attacks, such as homograph poison attacks [46], composite poison attacks [52], dynamic poison attacks [46], and learnable poison attacks [64]. Yang et al. [85] also proposed two metrics to evaluate the stealthiness of poison attacks. Recently, more concealed poison attacks are proposed that use the syntactic structure [63] or the text style [58] as triggers. Some studies [26] further proposed triggerless poison attacks on NLP models. Chen et al. [19] proposed a task-agnostic poison attack approach for pre-trained models, named BadPre. BadPre can inject backdoors into pre-trained models and keep the backdoors in downstream tasks.

Though promising, the above poison attacks in the CV and NLP fields are not applied to the source code. The reasons are two-fold. (1) The poison attacks in the CV field are designed for images and can not process the source code sequences. (2) The poison attacks in the NLP field are designed for the natural language text. The source code is quite different from the natural language text and must strictly follow rigid lexical, grammatical, and syntactical constraints; while the poison attacks for NLP mainly use some natural language chars, words, and sentences as triggers, which do not consider these constraints and lead to abnormal code files (e.g., files with compilation errors). The abnormal code files can be easily detected by program analysis tools (e.g., grammar checker) and cause the attack to fail. Therefore, it is necessary to explore the poison attacks for the source code.

Poison attacks for source code. The poison attacks on source code processing models have not been comprehensively explored. Schuster et al. [66] and Aghakhani et al. [10] attempted poison attacks on code completion models. Wan et al. [77] explored the poison attacks against code search models. They crafted some triggers (e.g., a specific comment) for two tasks and found the vulnerability of deep source code processing models to poison attacks. As stated in Section 8.2, existing poison attacks for source code focus on specific tasks and can not be transferred to other tasks. It hinders the development of poison attacks and poison detection in other software engineering applications, e.g., code repair. In this article, we perform a systematic investigation of poison attacks for source code. We present a strong imaginary enemy named CODEPOISONER to verify the vulnerability of existing deep source code processing models to poison attacks. Compared to previous poison attacks for source code [66, 77], our CODEPOISONER has two advantages: (1) It can be applied to multiple source code deep processing models and tasks, including code classification and code generation. (2) It provides more viable poison attack approaches, including token-level, statement-level, and snippet-level approaches. We hope CODEPOISONER can

further help practitioners know poison attacks for source code and inspire more advanced defense approaches.

9.2 Poison Detection

Poison detection aims to detect poison samples in the training data. Existing poison detection approaches can be divided into two categories: outliers-based approaches and representation-based approaches.

Outliers-based approaches consider outliers in the data as poison samples. Steinhart et al. [68] and Paudice et al. [59] proposed two poison detection approaches by removing outliers. They split a trusted training dataset by labels and then train a distance-based outlier detector for each label. For a new untrusted dataset, the outlier detectors remove samples that exceed some score threshold. In another work, Paudice et al. [60] defended against poison attacks by re-labeling all samples. Specifically, they re-label each sample with the most common label among its k nearest neighbors. Recently, some researchers [27, 62] proposed to identify outliers via perturbations. For poison attacks for images, Gao et al. [27] intentionally perturbed the input image and observed the randomness of predicted classes for perturbed inputs. Low entropy in predicted classes violates the input-dependence property of a clean sample and implies the presence of a malicious input. For poison attacks for natural languages, Qi et al. [62] argue that the trigger words are irrelevant to the context and thus can be easily detected as outlier words by language models. Thus, they utilized a powerful language model—GPT-2—to detect trigger words in inputs based on the leave-one-out strategy.

Representation-based approaches aim to detect poison samples based on latent representations of deep learning models. The intuition behind this approach is that latent representations capture the information necessary for learning, thereby making the difference between clean and poison samples more pronounced. Activation Clustering [18] and Spectral [72] are two popular representation-based approaches. First, they feed all the inputs of each label to a trained model and collect their representation values separately. Then, they analyze the representations of each label to identify poison samples. Specifically, Activation Clustering uses the K-means algorithm to cluster the representations into two clusters. If the number of representations in one cluster is below a threshold, then this cluster will be identified as poison, since two clusters should be divided equally as usual. Once a poison clustering is determined, the corresponding data in the poison cluster will be removed. As for Spectral, it finds that using singular value decomposition on all representations can expose the poisoning data, as the representations of poisoning data tend to have higher scores.

Existing poison detection approaches are at an early stage and need to be improved. In this article, we think attackers' triggers are the key basis for identifying poison samples. However, existing approaches can not specify the triggers. It causes two limitations: ❶ Existing detection approaches have weak explainability. They can not provide reasons (e.g., a specific trigger) for their outputs. The defenders do not know why these samples are poison. ❷ Without triggers, existing detection approaches would miss a few poison samples. Thus, victim models still may be attacked. The experimental results in Table 5 also validate this point. To alleviate these problems, we propose a new poison detection approach named CODEDETECTOR. Compared to previous approaches, CODEDETECTOR can mine potential triggers in the data using the integrated gradient technique. Then, it considers the samples containing triggers as poison samples. The overall pipeline is comprehensible and provides concrete reasons for outputs. Besides, based on triggers, CODEDETECTOR can detect more even all poison samples. The experimental results in Table 5 show that CODEDETECTOR significantly outperforms existing detection approaches and effectively detects all poison samples produced by five poison attack approaches.

9.3 Deep Learning for Source Code Processing

With software becoming ubiquitous in our daily life, open-source and closed-source code repositories have been becoming unprecedentedly large and complex [11]. Recently, researchers leverage deep learning techniques to mine knowledge from large-scale code corpus to automate the software development and maintenance process. This line of studies is termed as *deep learning for source code processing*, such as defect detection [54, 91], clone detection [79, 88], code repair [37, 73], code summarization [42], code generation [41, 43–45]. In this article, we conduct experiments on three representative source code processing tasks (i.e., defect detection, clone detection, and code repair). Next, we provide a summary of three tasks and related work.

Defect Detection. Defect detection aims to classify a program as defective or non-defective. This task plays an important role in ensuring the security of software, as well as saving much effort and time for software development. Li et al. [51] presented the first systematic approach for defect detection using deep learning techniques, named SySeVR. Zhou et al. [91] proposed the Devign for defect detection, which represented a program by fusing its AST, control-flow, and dataflow graphs into a unified heterogeneous graph **code property graph (CPG)**. Zhou et al. [91] also released a defect detection dataset to facilitate further research. Later, some studies [22, 54, 78] further leveraged **graph neural network (GNN)** to represent the control, data, and call dependencies of a program for defect detection. Recently, Feng et al. [25] proposed the first large-scale pre-trained model for source code, named CodeBERT. Feng et al. [25] applied CodeBERT to the defect detection tasks and obtained SOTA results on multiple benchmarks.

Clone Detection. Clone detection is to detect similar code snippets and is a fundamental task for many software engineering activities (e.g., code reuse, code search). Recently, many DL-based approaches are designed to represent a pair of code snippets for clone detection. The key idea of these DL-based approaches lies in representing the code snippet as a feature vector and computing the similarity between different vectors. White et al. [82] proposed a DL-based clone detection model that used a **recurrent neural network (RNN)** to represent the lexical and syntactic information of source code. Wei and Li [81] further leveraged the TreeLSTM to represent the syntactic information (i.e., AST) of source code. Zhang et al. [88] proposed an AST-based neural network named ASTNN for clone detection. ASTNN decomposed a large AST into several small statement trees to compute a code representation vector. Wang et al. [79] combined the AST with the control-flow graph of source code and proposed a GNN for code representation. Feng et al. [25] proposed a pre-trained code representation model named CodeBERT and achieved SOTA results on the clone detection task.

Code Repair. The code repair task is to automatically fix bugs in programs. Bhatia and Singh [14] and Santos et al. [65] proposed RNN-based language models for fixing syntax errors in programs. Inspired by the **sequence-to-sequence (Seq2Seq)** models [70] in the NLP field, some researchers [21, 30, 73] applied the Seq2Seq models into the code repair task, by transforming the buggy programs into fixed ones. Besides, many approaches have been proposed to repair programs by editing their syntax structure. Chakraborty et al. [17] proposed a tree-based code repair model named CODIT. CODIT learned to edit the buggy code at the AST level to generate syntactically correct patches. Zhu et al. [92] proposed a syntax-guided decoder network for code repair, which could generate edit actions rather than the modified code. Recently, various pre-trained techniques have been applied to code repair. CodeBERT [25] is a pioneer pre-trained model for source code and has obtained significant improvements on the code repair task. Jiang et al. [37] introduced a pre-trained language model for code repair and proposed a code-aware search strategy to search for more correct patches.

Although deep source code processing models produce a powerful performance on many tasks, security issues are lying within them. In this article, we identify an emergent and serious threat named poison attacks. The attackers may deceive users into integrating poisoned models as part of their applications and further mislead poisoned systems to produce targeted erroneous results. To alleviate this threat, we propose an effective defense approach to detect poison attacks. We hope this work can alarm SE practitioners and inspire the design of more advanced defense techniques.

10 CONCLUSION AND FUTURE WORK

This article addresses an emergent security threat to deep code processing models, named *poison attacks*. The attackers inject backdoors into models by poisoning the training data and further manipulate poisoned models by activating backdoors. To reveal severe threats from poison attacks, we present a poison attack approach for source code named CODEPOISONER as a strong imaginary enemy. CODEPOISONER provides four viable poisoning strategies (i.e., three rule-based strategies and a language-model-guided strategy) to make poison samples and conduct poison attacks. To defend against poison attacks, we further propose a poison detection approach named CODEDETECTOR. CODEDETECTOR utilizes the integrated gradients technique to automatically detect potential poison samples in the training data. We apply CODEPOISONER and CODEDETECTOR to six deep code processing models, including defect detection, clone detection, and code repair models. Experimental results identify the threat of poison attacks and show that CODEPOISONER can inject backdoors into models with a high attack success rate (average: 98.6%, maximum: 100%) under low poisoning cost (2%). Besides, our CODEDETECTOR can effectively detect (maximum: 100%) poison samples and defend against multiple poison attack approaches.

Given the surging popularity of deep code processing models, this article takes the first step to reveal poison attacks for source code and provides a plausible poison detection approach. For SE practitioners, it can help in understanding and defending against poison attacks in practice. They also can further explore more insidious poison attack approaches and develop more powerful defense tools. For instance, injecting triggers into **abstract syntax trees (AST)**, poisoning the pre-trained models, and defending the LM-guided poison attack approach.

ACKNOWLEDGMENTS

We thank the reviewers for their helpful comments. Ge Li and Zhi Jin are the corresponding authors.

REFERENCES

- [1] Wikipedia. 2023. Wikipedia. <https://www.wikipedia.org>
- [2] Google. 2023. Google Translation. <https://translate.google.com>
- [3] GitHub. 2023. GitHub. <https://github.com/>
- [4] Stack Overflow. 2023. Stack Overflow. <https://stackoverflow.com>
- [5] Jia Li. 2023. Replicate Package. <https://github.com/LJ2lijia/CodeDetector>
- [6] Black Dock. 2023. Black Dock. <https://www.blackducksoftware.com/>
- [7] Microsoft. 2023. GitHub Copilot. <https://copilot.github.com>
- [8] TreeSitter. 2023. TreeSitter. <https://tree-sitter.github.io/tree-sitter>
- [9] Microsoft. 2023. IntelliCode. <https://visualstudio.microsoft.com/services/intellicode>
- [10] Hojjat Aghakhani, Wei Dai, Andre Manoel, Xavier Fernandes, Anant Kharkar, Christopher Kruegel, Giovanni Vigna, David Evans, Ben Zorn, and Robert Sim. 2023. TrojanPuzzle: Covertly poisoning code-suggestion models. *CoRR* abs/2301.02344 (2023).
- [11] Miltiadis Allamanis, Earl T. Barr, Premkumar Devanbu, and Charles Sutton. 2018. A survey of machine learning for big code and naturalness. *ACM Comput. Surv.* 51, 4 (2018), 1–37.
- [12] Brenda S. Baker. 1995. On finding duplication and near-duplication in large software systems. In *2nd Working Conference on Reverse Engineering*. IEEE, 86–95.

- [13] Mauro Barni, Kassem Kallas, and Benedetta Tondi. 2019. A new backdoor attack in CNNs by training set corruption without label poisoning. In *IEEE International Conference on Image Processing (ICIP'19)*. IEEE, 101–105.
- [14] Sahil Bhatia and Rishabh Singh. 2016. Automated correction for syntax errors in programming assignments using recurrent neural networks. *arXiv preprint arXiv:1603.06129* (2016).
- [15] Eitan Borgnia, Jonas Geiping, Valeriia Cherepanova, Liam Fowl, Arjun Gupta, Amin Ghiasi, Furong Huang, Micah Goldblum, and Tom Goldstein. 2021. DP-InstaHide: Provably defusing poisoning and backdoor attacks with differentially private data augmentations. *CoRR abs/2103.02079* (2021).
- [16] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.). Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
- [17] Saikat Chakraborty, Yangruibo Ding, Miltiadis Allamanis, and Baishakhi Ray. 2020. CODIT: Code editing with tree-based neural models. *IEEE Trans. Softw. Eng.* 48, 4 (2020), 1385–1399.
- [18] Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. 2019. Detecting backdoor attacks on deep neural networks by activation clustering. In *AAAI Workshop on Artificial Intelligence Safety (SafeAI@ AAAI'19)*.
- [19] Kangjie Chen, Yuxian Meng, Xiaofei Sun, Shangwei Guo, Tianwei Zhang, Jiwei Li, and Chun Fan. 2021. BadPre: Task-agnostic backdoor attacks to pre-trained NLP foundation models. In *International Conference on Learning Representations*.
- [20] Xiaoyi Chen, Ahmed Salem, Michael Backes, Shiqing Ma, and Yang Zhang. 2021. BadNL: Backdoor attacks against NLP models. In *ICML Workshop on Adversarial Machine Learning*.
- [21] Zimin Chen, Steve Kommrusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. 2019. Sequencer: Sequence-to-sequence learning for end-to-end program repair. *IEEE Trans. Softw. Eng.* 47, 9 (2019), 1943–1959.
- [22] Xiao Cheng, Haoyu Wang, Jiayi Hua, Guoai Xu, and Yulei Sui. 2021. DeepWukong: Statically detecting software vulnerabilities using deep graph neural network. *ACM Trans. Softw. Eng. Methodol.* 30, 3 (2021), 1–33.
- [23] Saumya K. Debray, William Evans, Robert Muth, and Bjorn De Sutter. 2000. Compiler techniques for code compaction. *ACM Trans. Program. Lang. Syst.* 22, 2 (2000), 378–415.
- [24] John R. Douceur. 2002. The Sybil attack. In *International Workshop on Peer-to-peer Systems*. Springer, 251–260.
- [25] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. CodeBERT: A pre-trained model for programming and natural languages. In *Findings of the Association for Computational Linguistics (EMNLP'20)*. Association for Computational Linguistics, 1536–1547.
- [26] Leilei Gan, Jiwei Li, Tianwei Zhang, Xiaoya Li, Yuxian Meng, Fei Wu, Yi Yang, Shangwei Guo, and Chun Fan. 2022. Triggerless backdoor attack for NLP tasks with clean labels. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL'22)*, Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruíz (Eds.). Association for Computational Linguistics, 2942–2952. <https://doi.org/10.18653/V1/2022.NAACL-MAIN.214>
- [27] Yansong Gao, Change Xu, Derui Wang, Shiping Chen, Damith C. Ranasinghe, and Surya Nepal. 2019. STRIP: A defence against trojan attacks on deep neural networks. In *35th Annual Computer Security Applications Conference*. 113–125.
- [28] David Gros, Hariharan Sezhiyan, Prem Devanbu, and Zhou Yu. 2020. Code to comment “translation”: Data, metrics, baselining & evaluation. In *35th IEEE/ACM International Conference on Automated Software Engineering (ASE'20)*. IEEE, 746–757.
- [29] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. BadNets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017).
- [30] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. 2017. DeepFix: Fixing common C language errors by deep learning. In *AAAI Conference on Artificial Intelligence*, Vol. 31.
- [31] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9, 8 (1997), 1735–1780.
- [32] Xing Hu, Ge Li, Xin Xia, David Lo, and Zhi Jin. 2018. Deep code comment generation. In *IEEE/ACM 26th International Conference on Program Comprehension (ICPC'18)*. IEEE, 200–20010.
- [33] Kunzhe Huang, Yiming Li, Baoyuan Wu, Zhan Qin, and Kui Ren. 2022. Backdoor defense via decoupling the training process. In *10th International Conference on Learning Representations (ICLR'22)*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=TySnJ-0RdKI>

- [34] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. CodeSearchNet challenge: Evaluating the state of semantic code search. *CoRR abs/1909.09436* (2019). arXiv:1909.09436 <http://arxiv.org/abs/1909.09436>
- [35] Paras Jain, Ajay Jain, Tianjun Zhang, Pieter Abbeel, Joseph Gonzalez, and Ion Stoica. 2021. Contrastive code representation learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event/Punta Cana, Dominican Republic, 7-11 November, 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, 5954–5971. <https://doi.org/10.18653/V1/2021.EMNLP-MAIN.482>
- [36] Yujie Ji, Xinyang Zhang, Shouling Ji, Xiapu Luo, and Ting Wang. 2018. Model-reuse attacks on deep learning systems. In *ACM SIGSAC Conference on Computer and Communications Security*. 349–363.
- [37] Nan Jiang, Thibaud Lutellier, and Lin Tan. 2021. CURE: Code-aware neural machine translation for automatic program repair. In *IEEE/ACM 43rd International Conference on Software Engineering (ICSE’21)*. IEEE, 1161–1173.
- [38] Rafael-Michael Karampatsis, Hlib Babii, Romain Robbes, Charles Sutton, and Andrea Janes. 2020. Big code!= big vocabulary: Open-vocabulary models for source code. In *ACM/IEEE 42nd International Conference on Software Engineering*. 1073–1085.
- [39] Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Conference on Empirical Methods in Natural Language Processing (EMNLP’14)*. ACL, 1746–1751.
- [40] Keita Kurita, Paul Michel, and Graham Neubig. 2020. Weight poisoning attacks on pretrained models. In *58th Annual Meeting of the Association for Computational Linguistics*. 2793–2806.
- [41] Jia Li, Ge Li, Zhuo Li, Zhi Jin, Xing Hu, Kechi Zhang, and Zhiyi Fu. 2023. CodeEditor: Learning to edit source code with pre-trained models. *ACM Trans. Softw. Eng. Methodol.* 32, 6 (2023), 143:1–143:22. <https://doi.org/10.1145/3597207>
- [42] Jia Li, Yongmin Li, Ge Li, Xing Hu, Xin Xia, and Zhi Jin. 2021. EDITSUM: A retrieve-and-edit framework for source code summarization. In *36th IEEE/ACM International Conference on Automated Software Engineering (ASE’21)*. IEEE.
- [43] Jia Li, Ge Li, Yongmin Li, and Zhi Jin. 2023. Enabling programming thinking in large language models toward code generation. *CoRR abs/2305.06599* (2023). <https://doi.org/10.48550/ARXIV.2305.06599> arXiv:2305.06599.
- [44] Jia Li, Yongmin Li, Ge Li, Zhi Jin, Yiyang Hao, and Xing Hu. 2023. SkCoder: A sketch-based approach for automatic code generation. In *45th IEEE/ACM International Conference on Software Engineering (ICSE’23)*. IEEE, 2124–2135. DOI: <https://doi.org/10.1109/ICSE48619.2023.00179>
- [45] Jia Li, Yunfei Zhao, Yongmin Li, Ge Li, and Zhi Jin. 2023. *Towards Enhancing In-Context Learning for Code Generation*. *CoRR abs/2303.17780* (2023). <https://doi.org/10.48550/ARXIV.2303.17780> arXiv:2303.17780
- [46] Shaofeng Li, Hui Liu, Tian Dong, Benjamin Zi Hao Zhao, Minhui Xue, Haojin Zhu, and Jialiang Lu. 2021. Hidden backdoors in human-centric language models. In *ACM SIGSAC Conference on Computer and Communications Security*. 3123–3140. DOI: <https://doi.org/10.1145/3460120.3484576>
- [47] Shaofeng Li, Minhui Xue, Benjamin Zi Hao Zhao, Haojin Zhu, and Xinpeng Zhang. 2020. Invisible backdoor attacks on deep neural networks via steganography and regularization. *IEEE Trans. Depend. Secure Comput.* 18, 5 (2020), 2088–2105.
- [48] Yuanchun Li, Jiayi Hua, Haoyu Wang, Chunyang Chen, and Yunxin Liu. 2021. DeepPayload: Black-box backdoor attack on deep learning models through neural payload injection. In *IEEE/ACM 43rd International Conference on Software Engineering (ICSE’21)*. IEEE, 263–274.
- [49] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. 2021. Anti-backdoor learning: Training clean models on poisoned data. In *Annual Conference on Neural Information Processing Systems (NeurIPS’21)*, Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 14900–14912. Retrieved from <https://proceedings.neurips.cc/paper/2021/hash/7d38b1e9bd793d3f45e0e212a729a93c-Abstract.html>
- [50] Yige Li, Xixiang Lyu, Nodens Koren, Lingjuan Lyu, Bo Li, and Xingjun Ma. 2021. Neural attention distillation: Erasing backdoor triggers from deep neural networks. In *9th International Conference on Learning Representations (ICLR’21)*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=9l0K4OM-oXE>
- [51] Zhen Li, Deqing Zou, Shouhuai Xu, Hai Jin, Yawei Zhu, and Zhaoxuan Chen. 2021. SySeVR: A framework for using deep learning to detect software vulnerabilities. *IEEE Transactions on Dependable and Secure Computing* 19, 4 (2021), 2244–2258.
- [52] Junyu Lin, Lei Xu, Yingqi Liu, and Xiangyu Zhang. 2020. Composite backdoor attack for deep neural network by mixing existing benign features. In *ACM SIGSAC Conference on Computer and Communications Security*. 113–131.
- [53] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2017. Trojaning attack on neural networks. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The Internet Society. https://www.ndss-symposium.org/wp-content/uploads/2018/02/ndss2018_03A-5_Liu_paper.pdf

- [54] Zhenguang Liu, Peng Qian, Xiaoyang Wang, Yuan Zhuang, Lin Qiu, and Xun Wang. 2023. Combining graph neural networks with expert knowledge for smart contract vulnerability detection. *IEEE Transactions on Knowledge and Data Engineering* 35, 2 (2023), 1296–1310. <https://doi.org/10.1109/TKDE.2021.3095196>
- [55] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin B. Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. CodeXGLUE: A machine learning benchmark dataset for code understanding and generation. In *35th Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [56] Bodo Möller, Thai Duong, and Krzysztof Kotowicz. 2014. This POODLE bites: Exploiting the SSL 3.0 fallback. *Secur. Advis.* 21 (2014), 34–58.
- [57] Manishankar Mondal, Md Saidur Rahman, Chanchal K. Roy, and Kevin A. Schneider. 2018. Is cloned code really stable? *Empir. Softw. Eng.* 23, 2 (2018), 693–770.
- [58] Xudong Pan, Mi Zhang, Beina Sheng, Jiaming Zhu, and Min Yang. 2022. Hidden trigger backdoor attack on NLP models via linguistic style manipulation. In *31st USENIX Security Symposium (USENIX Security'22)*. 3611–3628.
- [59] Andrea Paudice, Luis Muñoz-González, Andras Gyorgy, and Emil C. Lupu. 2018. Detection of adversarial training examples in poisoning attacks through anomaly detection. CoRR abs/1802.03041 (2018). arXiv:1802.03041 <http://arxiv.org/abs/1802.03041>
- [60] Andrea Paudice, Luis Muñoz-González, and Emil C. Lupu. 2018. Label sanitization against label flipping poisoning attacks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 5–15.
- [61] L. Prechelt, G. Malpohl, and M. Philippsen. 2000. JPlag: Finding plagiarisms among a set of programs. Technical Report. University of Karlsruhe, Department of Informatics.
- [62] Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021. ONION: A simple and effective defense against textual backdoor attacks. In *Conference on Empirical Methods in Natural Language Processing*. 9558–9566.
- [63] Fanchao Qi, Mukai Li, Yangyi Chen, Zhengyan Zhang, Zhiyuan Liu, Yasheng Wang, and Maosong Sun. 2021. Hidden killer: Invisible textual backdoor attacks with syntactic trigger. In *59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. 443–453.
- [64] Fanchao Qi, Yuan Yao, Sophia Xu, Zhiyuan Liu, and Maosong Sun. 2021. Turn the combination lock: Learnable textual backdoor attacks via word substitution. In *59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. 4873–4883.
- [65] Eddie Antonio Santos, Joshua Charles Campbell, Dhvani Patel, Abram Hindle, and José Nelson Amaral. 2018. Syntax and sensibility: Using language models to detect and correct syntax errors. In *IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER'18)*. IEEE, 311–322.
- [66] Roei Schuster, Congzheng Song, Eran Tromer, and Vitaly Shmatikov. 2021. You autocomplete me: Poisoning vulnerabilities in neural code completion. In *30th USENIX Security Symposium (USENIX Security'21)*. 1559–1575.
- [67] Ali Shafahi, W. Ronny Huang, Mahyar Najibi, Octavian Suci, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! Targeted clean-label poisoning attacks on neural networks. *Adv. Neural Inf. Process. Syst.* 31 (2018).
- [68] Jacob Steinhardt, Pang Wei W. Koh, and Percy S. Liang. 2017. Certified defenses for data poisoning attacks. *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [69] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*. PMLR, 3319–3328.
- [70] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *International Conference on Advances in Neural Information Processing Systems*. 3104–3112.
- [71] Jeffrey Svajlenko, Judith F. Islam, Iman Keivanloo, Chanchal K. Roy, and Mohammad Mamun Mia. 2014. Towards a big data curated benchmark of inter-project code clones. In *IEEE International Conference on Software Maintenance and Evolution*. IEEE, 476–480.
- [72] Brandon Tran, Jerry Li, and Aleksander Mądry. 2018. Spectral signatures in backdoor attacks. In *32nd International Conference on Neural Information Processing Systems*. 8011–8021.
- [73] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2018. An empirical investigation into learning bug-fixing patches in the wild via neural machine translation. In *33rd ACM/IEEE International Conference on Automated Software Engineering*. 832–837.
- [74] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2019. An empirical study on learning bug-fixing patches in the wild via neural machine translation. *ACM Trans. Softw. Eng. Methodol.* 28, 4 (2019), 1–29.
- [75] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2019. An empirical study on learning bug-fixing patches in the wild via neural machine translation. *ACM Trans. Softw. Eng. Methodol.* 28, 4 (2019), 1–29.

- [76] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *International Conference on Advances in Neural Information Processing Systems*. 5998–6008.
- [77] Yao Wan, Shijie Zhang, Hongyu Zhang, Yulei Sui, Guandong Xu, Dezhong Yao, Hai Jin, and Lichao Sun. 2022. You see what I want you to see: Poisoning vulnerabilities in neural code search. In *30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1233–1245.
- [78] Huanting Wang, Guixin Ye, Zhanyong Tang, Shin Hwei Tan, Songfang Huang, Dingyi Fang, Yansong Feng, Lizhong Bian, and Zheng Wang. 2020. Combining graph-based learning with automated data collection for code vulnerability detection. *IEEE Trans. Inf. Forens. Secur.* 16 (2020), 1943–1958.
- [79] Wenhan Wang, Ge Li, Bo Ma, Xin Xia, and Zhi Jin. 2020. Detecting code clones with graph neural network and flow-augmented abstract syntax tree. In *IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER'20)*. IEEE, 261–271.
- [80] Cody Watson, Michele Tufano, Kevin Moran, Gabriele Bavota, and Denys Poshyvanyk. 2020. On learning meaningful assert statements for unit test cases. In *ACM/IEEE 42nd International Conference on Software Engineering*. 1398–1409.
- [81] Hui-Hui Wei and Ming Li. 2017. Supervised deep features for software functional clone detection by exploiting lexical and syntactical information in source code. In *26th International Joint Conference on Artificial Intelligence*. 3034–3040.
- [82] Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. 2016. Deep learning code fragments for code clone detection. In *31st IEEE/ACM International Conference on Automated Software Engineering (ASE'16)*. IEEE, 87–98.
- [83] Hongwei Xi. 1999. Dead code elimination through dependent types. In *International Symposium on Practical Aspects of Declarative Languages*. Springer, 228–242.
- [84] Chang Xu, Jun Wang, Yuqing Tang, Francisco Guzmán, Benjamin I. P. Rubinstein, and Trevor Cohn. 2021. A targeted attack on black-box neural machine translation with parallel data poisoning. In *Proceedings of the Web Conference 2021*. 3638–3650.
- [85] Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. 2021. Rethinking stealthiness of backdoor attack against NLP models. In *59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. 5543–5557.
- [86] Noam Yefet, Uri Alon, and Eran Yahav. 2020. Adversarial examples for models of code. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 1–30.
- [87] Huangzhao Zhang, Zhuo Li, Ge Li, Lei Ma, Yang Liu, and Zhi Jin. 2020. Generating adversarial examples for holding robustness of source code processing models. In *AAAI Conference on Artificial Intelligence*, Vol. 34. 1169–1176.
- [88] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, Kaixuan Wang, and Xudong Liu. 2019. A novel neural source code representation based on abstract syntax tree. In *IEEE/ACM 41st International Conference on Software Engineering (ICSE'19)*. IEEE, 783–794.
- [89] Quan Zhang, Yifeng Ding, Yongqiang Tian, Jianmin Guo, Min Yuan, and Yu Jiang. 2021. AdvDoor: Adversarial backdoor attack of deep learning system. In *International Symposium on Software Testing and Analysis (ISSTA'21)*. ACM, 127–138. <https://doi.org/10.1145/3460319.3464809>
- [90] Haoti Zhong, Cong Liao, Anna Cinzia Squecciarini, Sencun Zhu, and David Miller. 2020. Backdoor embedding in convolutional neural network models via invisible perturbation. In *10th ACM Conference on Data and Application Security and Privacy*. 97–108.
- [91] Yaqin Zhou, Shangqing Liu, Jing Kai Siow, Xiaoning Du, and Yang Liu. 2019. Devign: Effective vulnerability identification by Learning comprehensive program semantics via graph neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 10197–10207. <https://proceedings.neurips.cc/paper/2019/hash/49265d2447bc3bbfe9e76306ce40a31f-Abstract.html>
- [92] Qihao Zhu, Zeyu Sun, Yuan-an Xiao, Wenjie Zhang, Kang Yuan, Yingfei Xiong, and Lu Zhang. 2021. A syntax-guided edit decoder for neural program repair. In *29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 341–353.

Received 18 June 2022; revised 1 September 2023; accepted 9 October 2023