

# A Hybrid LLM Workflow can Help Identify User Privilege-Related Variables in Programs of Any Size

Haizhou Wang\*, Zhilong Wang\*, Peng Liu\*

\*Pennsylvania State University

{hjlw5074, zzw169, pxl20}@psu.edu

**Abstract**—Many programs involves operations and logic manipulating user privileges, which is essential for the security of an organization. To protect the program and the organization against privilege leakage attacks, it is important to eliminate the vulnerabilities which can be exploited to achieve such attacks. Analysts often choose to find user privilege related (UPR) variables as start points to investigate the code where the UPR variables are used for vulnerabilities. In this paper, we introduce a large language model (LLM) workflow that can assist analysts in identifying such UPR variables, which is considered to be a very time-consuming task. The proposed approach avoids the drawbacks introduced by directly prompting a LLM to find UPR variables by focusing on leverage the LLM at statement level instead of supplying LLM with very long code snippets. Variables will be audited for a UPR score, and those ones with scores above certain threshold shall be investigated. Our experiments show that using a typical UPR score threshold (i.e., UPR score  $>0.8$ ), the false positive rate (FPR) is only 13.49%, while the number of UPR variables found is significantly more than that of the heuristic based method.

## I. INTRODUCTION

Many organizations and companies run server-side programs to provide services to various users, and therefore, it is necessary to properly control the privileges each user may possess. Typically, user privileges related operations implemented in a server-side program include but not limited to authentications and authorizations (e.g., the mapping between user IDs and file/data permissions). Such user privilege related operations are often critical in terms of the security of the programs, the servers, and the service providers. Specifically, service providers need to carefully control the privileges granted to users, preventing the privilege leakage risks. Thus, in ideal cases, these privilege related operations should be implemented carefully (by very experienced programmers) to guard these operations from being affected by memory corruptions (e.g., buffer overflows) and exploitations of logic vulnerabilities. However, it is observed in the real world that many privilege related operations are not adequately guarded.

Since many privilege related operations are not yet adequately guarded, many service providers regularly review and analyze their source code to find vulnerabilities that may cause user privilege leakage. In general, the relevant vulnerabilities can be categorized into two types: memory corruption vulnerabilities and logic vulnerabilities. Memory corruptions at runtime, such as buffer overflows and integer overflows, are regularly exploited by attackers. By providing crafted payload can an attacker modify certain values in the program memory,

and therefore launch various attacks, including control data attacks (e.g. ROP [1], [2]) and non-control data attacks (e.g. DOP [3]). Logic vulnerabilities are vulnerabilities related to the program logic only, and the corresponding exploitations usually do not result in any memory corruption errors. One type of logic vulnerability of privilege leakage is hard-coded secrets, such as API keys, tokens, etc. Some hard-coded secrets can be very inconspicuous. For example, if a programmer logs the plaintext passwords into the log file, then the passwords can be seen by any person who has access to this log file, essentially causing privilege leakage. Code 1 illustrates such a case, which we will discuss in detail in Section III. Another type of logic vulnerability is related to web resource access in which a portion of a URL represents the location of a file/resource on a web server. This vulnerability makes the path traversal attack possible so that an attacker can manipulate the path in a URL to access files or directories that they shouldn't have access to.

Although service providers are strongly motivated to identify such logic vulnerabilities, this is not a trivial task. On the one hand, the program's execution and functionality are not interrupted by such privilege-related logic bugs, which makes it very challenging to detect them through automatic bug finding tools (such as Fuzz testing tools) relying on crashes or sanitizers to detect the abnormal program behaviors. On the other hand, since such vulnerabilities are blended in the program logic, a good understanding of certain higher-level program semantics is usually required. As a result, large tech companies still rely on code review process to discover and eliminate such vulnerabilities.

In order to make the code review process successful, the very first challenge to solve is to spot the privilege-related operations and code statements. As the privilege operation logic is diverse, real-world analysts usually find user privilege related (UPR) variables first instead of directly spotting the code statements. Take the popular static analysis tool `semgrep` used in industry as an example, more than 40% of their detection rules focus on specifying the patterns that are related to UPR variables. Another merit of finding UPR variables instead of the code statements is that some variables contain sensitive information, so that the variables themselves need to be investigated. However, finding UPR variables is not an easy task either, since there are still various types of UPR variables, such as passwords, API keys, tokens, etc.

Although it is very important to identify UPR variables, existing works unfortunately mainly rely on heuristic rules

(such as regular expressions), whose scalability is limited when screening large code bases or porting to new applications. For example, many `semgrep` rules leverage regular expressions to match the potential keywords that could exist in UPR variables. When using regular expressions to discover UPR variables, one can easily foresee that this approach is likely to suffer from a high false negative rate. In addition, based on our experiments in Section VI-C, we also found that this approach could suffer from a high false positive rate as well. The unsatisfying performance of heuristic-based methods is mainly caused by the lacking of explicit patterns for UPR variables. Specifically, there is no easy way to summarize all the features and characteristics of UPR variables, as whether a variable is UPR largely depends on higher-level application logic. For example, although HTTP servers often have file and CGI paths, FTP servers and SSH/Telnet servers don't.

**Key Observation.** Since user privileges are related to the logical level of program understanding, it is necessary to analyze the application logic of the program before accurately discovering UPR variables.

However, understanding the role of each variable in terms of the business/application logic of the program is non-trivial. Several previous work [4], [5], [6] strive to leverage pre-trained language models to gain understanding of the program given the code; however, the gained understanding is not very satisfying. Fortunately, since the emergence of large language models (LLM) such as ChatGPT, program comprehension tasks become less challenging, and therefore it becomes more feasible to achieve AI-assisted UPR variable identification.

Despite the promising performance shown by LLMs, it is still non-trivial to build a practical LLM workflow to detect UPR variables. As we will shortly show in Section III-C, even for one of the most powerful LLMs, GPT-4, researchers will encounter challenges when directly prompting a LLM to find UPR variables given the source code. Specifically, it is widely observed that as the input context length grows, the performance of the LLM deteriorates [7], [8], [9], [10], [11], so that identifying UPR variables in long code snippets becomes less effective.

In this paper, we propose a novel LLM workflow that can help analysts to identify UPR variables in a program with any size. Specifically, *our approach combines static program analysis and LLM to audit each variable in a program*, so that each variable will be given a UPR score, which can reflect the degree of relevance of a variable to user privileges. Our approach avoids the drawbacks introduced by directly prompting a LLM to find UPR variables by focusing on leverage the LLM at statement level instead of supplying LLM with very long code snippets. Since our method require an LLM to evaluate the code, it is important to notice that our workflow is designed for non-adversarial settings: the source code is not obfuscated. The UPR score will be scaled between 0 to 10, where 0 means the variable is completely not related to user privileges. After the UPR scores are given to all the variables, analysts will be able to investigate those variables with higher scores to confirm whether a variable is UPR. To summarize, our contributions include:

- We propose a novel LLM workflow that can help human

analysts to identify UPR variables in a program of any size.

- We have evaluated the proposed method by investigating the reported potential UPR variables manually. According to our experiments presented in Section VI-B, when only investigating the top variables with highest UPR scores ( $>8.0$ ), only 13.49% false positive rate (FPR) is reported, and only less than 15% of the variables need manual investigation. With that being said, analysts may save up to 85% of the time comparing to manually checking all the variables in a program while looking for UPR variables.
- We have compare the proposed method with the widely-adopted heuristic based method, whose heuristic rules are included. It is shown that our method can identify significantly more UPR variables while maintaining much less FPR.

## II. BACKGROUND

### A. Program Dependency Analysis

Program dependence analysis is fundamental for various program analysis tasks, such as program slicing [12], and program comprehension [13]. It is used to identify the control and data dependence relations between statements (at source code level) or instructions (at binary level) in a program. Control dependence describes a situation in which whether a statement will be executed is affected by another statement. If whether statement A will be executed depends on the execution result of statement B, then statement A is said to be control dependent to statement B. Data dependence describes a relation between statements based on variable definition and usage. If statement A used a variable defined at statement B, then statement A is data dependent on statement B.

Typically there exists numerous pairs of such dependence relations, and therefore most program dependence analysis methods rely on the program dependence graph [14] (PDG) to represent both control and data dependence. Each nodes in the source-code level PDG could be either code statement, expressions, or variable. Intuitively, nodes that are connected by an edge are considered to have very close relationship, so that program dependence relations can be used to track the "side effect" of a code statement.

### B. Large Language Model in Software Security

Large language models (LLMs) have been attracting attentions from researchers of various fields, after the release of ChatGPT. Since then, researchers in the field of software security and software engineering have started exploring how LLM may assist solving various tasks, such as program repair [15], [11], program testing [16], and fuzzing [17]. As more and more researchers start working on this field, the contributions of most of works can be categorized into 3 categories: 1) prompt engineering, 2) LLM fine-tuning, and 3) workflow design.

Though LLMs are promising in helping solving software challenges, adopting LLM could be non-trivial. For example, despite the diversity of papers using LLM to solve software

challenges, many of them involve using LLM to generate code. Indeed LLM is very powerful in generating code, but it is also true that the quality of generated code snippets are usually unpredictable [11]. Furthermore, the quality of the LLM output can be greatly affected by the context length [7], [8], [9], [10], [11], which introduces challenges when working with long code snippets.

### III. MOTIVATIONS AND PROBLEM STATEMENT

#### A. Importance of UPR Variables

Many server side programs involves data operations, such as CRUD operations in web servers. Since a great number of users may have the access to the server, it is necessary to limit the privileges each user have, so that they are only allowed to access and/or modify the data for which they have authorization. The simplest example is the credential information of user accounts. If the the password of a user is leaked, an attacker can easily get all the privileges that user bear.

Accordingly, goals of many cyberattacks are essentially gaining or escalating privileges (i.e. privilege leakage), and therefore it is crucial to patch vulnerabilities which can be exploit to achieve privilege leakage. Since both memory [1], [18], [19] and logical vulnerabilities [20], [3] may be exploit to achieve potential privilege leakage, finding such vulnerability in the program becomes non-trivial. In particular, while some existing methods such as fuzzing may be used to find memory vulnerabilities automatically, there is no effective ways to find logic bugs. In many cases, just inappropriate coding practices may expose the privilege related information in danger, such as hard-coded tokens in source code managed by version control systems (e.g., git). Code 1 and Code 2 illustrates two simple

```

1 // improper logging practice
2 logging.INFO("%s", password)
3 ...

```

Code 1. An example of improper logging practice.

```

1 ...
2 # hardcoded secrets
3 api_key = "YOUR_OPENAI_API_KEY"
4 response_round_1 = openai.Completion.create(
5     engine=model,
6     prompt=prompt_round_1,
7     max_tokens=150
8 )
9 ...

```

Code 2. An example of hardcoded secrets.

cases where logic vulnerability may lead to privilege leak. In Code 1, the program will log the plaintext password into the log file, which may be accessed by other users. If the leaked password is obtained by an attacker, obviously certain privileges will be available to the attacker. In Code 2, an API key is hardcoded in the source code, which could be catastrophic if the source code is under version control. In the events of the version control synchronizing with an remote

server, many users will be able to access this hardcoded API key.

Therefore, to efficiently review the privileges handling logic in a program, real-world analysts usually start from identifying those variables containing information related to user privilege. However, it is impractical to manually analyze every single variables in the programs, especially for organizations who have huge code bases and many proprietary components.

#### B. Existing Methods to Identify UPR Variables

Only without considering the cost of time is it possible to have analysts manually identify all the UPR variables. For example, a medium size program (<10K lines of code) may contain thousands of variables, so that analyzing them one by one could be extremely time consuming and exhausting.

On the other hand, automatically identifying UPR variables is very challenging, because they are closely related to the high-level semantics of the program. In other words, whether they are UPR is partially, if not completely, depending on the functionality of the program. Therefore, no explicit general pattern can be found for the UPR variables.

Consequently, existing UPR variable detection methods are mostly heuristic-based, and a typical strategy adopted by real-world companies is to implement such heuristics using commercial static analysis tool, such as `Semgrep`, to handle their large-scale code bases. For example, a very common heuristic is to detect UPR variables by investigating the names of UPR variables, because UPR variables may contains certain user privilege related keywords (e.g. `key`, `token`, `password`, etc.). According, this variable name matching heuristic will be often implemented as regular expression.

Unfortunately, such heuristic-based methods can only cover cases which are specified in the heuristic rules. Using the variable name matching heuristic as an example, this strategy obviously will omit those UPR variables with less informative names. In addition, a variable containing user privilege related keywords is not necessarily a UPR variables. For example, usually a variable named `token` could be a UPR variable, but a variable named `token_type` is not likely to be a UPR variable. Furthermore, creating such heuristics requires human analysts to enumerate as many possibilities as possible, causing scalability issues.

Another challenge is that the influence of a UPR variable could be propagated to other variables, making other variables to be UPR. This implies that the dependencies between variables are also important when determine whether a variable is UPR. Thus, there also exist taint-analysis based heuristics to detect UPR variable, which set user privilege related API calls (e.g., `setuid`, `setgid`, etc.) as taint sinks. However, such detection method can only be used to find UPR variables involving the UPR API calls, leaving a substantial number of UPR variables undetected.

As a result, from analysts' perspective, it is desirable to have an automatic UPR variable detection tool that can save them time and efforts, while reporting as many UPR variables as possible. Not only the name of the variables, but also the dependence relationships in the program should an ideal UPR variable detection tool consider.

I want you to act as a security critical variable identifier. I will type a code snippet and you will tell me which variables are critical variables. Regarding what is a critical variable, when it is tampered, it may result in:

1. One can bypass authentication and/or authorization
2. One can achieve behaviors that are being banned
3. One can gain elevated privileges (uid changes)
4. One can access files/data that are not accessible (chown related)
5. One can steal secrets (keys, passwd, etc.)

Fig. 1. Initial prompt to find UPR variables using naive method

### C. Using LLM Naively

As elaborated in Section III-B, in order to identify UPR variables, it is important to take the contexts of inter-variable dependencies and program functionalities into consideration. As LLMs start showing their strength in understanding the natural language and program code, they seem to be a potential solution to identify UPR variables automatically.

The very natural idea of finding UPR variables using LLM would be simple: just ask it. We have conducted preliminary experiments of *prompting the LLM directly to find UPR variables*. Specifically, the experiments are conducted using GPT-4, which is considered to be the one of the most powerful LLM. To provide the GPT-4 with all the contexts, the prompt starts with an initial prompt providing the instructions and specification of the UPR variable, followed by the related source code without any modifications. In the experiments, we used `bftpd`, a medium size open-source FTP server, as the test program. We designed our initial prompt, which is shown in Figure 1.

We use the notion of "critical variable" along with a list of specification criteria, in an effort to clarify what is a UPR variable. For an FTP server, the listed 5 criteria should cover most of the UPR variables. In this experiment, we decided to conduct analysis at *function level*. We have manually tested roughly 25 functions, and the experiments had mixed results. In particular, we observe for shorter functions (i.e., less than 50 lines) involving only a few variables, the results looks promising, so that all the UPR variables are identified with only very few false positives; however, when analyzing longer functions, the result is not satisfactory. For example, when we ask GPT-4 to analyze the function `bftpd_login` in `login.c` for program `bftpd`, it gives unsatisfying reply, which is shown in Figure 2.

In this response, in total 4 UPR variables are missing and 2 identified are false positives, including:

- 1) `user` is not UPR variable, marked in red.
- 2) `config_getoption("DO_CHROOT")` is not even a variable, marked in red.
- 3) `state` is missing, which is used to indicate the user is authenticated in this function.

The critical variables in this code snippet are:

1. `password`: This is the password provided by the user for authentication. If tampered, it can lead to unauthorized access.
2. `file_auth`: This variable points to the file used to authenticate users. If tampered, it can lead to unauthorized access or bypassing of authentication.
3. `home_directory`: This variable holds the home directory retrieved from the auth file. If tampered, it can lead to unauthorized access to files/data.
4. `anonymous`: This variable holds the configuration option for anonymous user. If tampered, it can lead to unauthorized access or bypassing of authentication.
5. `change_uid_text`: This variable holds the configuration option for changing user id. If tampered, it can lead to privilege escalation.
6. `user`: This variable holds the username of the user trying to login. If tampered, it can lead to unauthorized access or bypassing of authentication.
7. `userinfo.pw_uid` and `userinfo.pw_gid`: These variables hold the user id and group id of the user. If tampered, it can lead to privilege escalation.
8. `str`: This variable is used to store various critical information like root directory, user name, etc. If tampered, it can lead to unauthorized access to files/data.
9. `config_getoption("DO_CHROOT")`: This configuration option decides whether to change the root directory or not. If tampered, it can lead to unauthorized access to files/data.

Fig. 2. GPT reply when prompted naively

- 4) `change_uid` and `change_uid_text` are missing, which is a configuration that can affect the `home_directory` and the execution of `set_uid` call.
- 5) `anon_ok` is missing, which is a configuration to allow anonymous users.

Consequently, we have derived a very important observation: *as the code snippet grows longer, the capability of the LLM in detecting UPR variables will decrease*. Unfortunately, GPT-4 is completely a blackbox, so that its parameters and meta-information are not available. Nevertheless, our observation does echo with many previous works. According to the experiments in previous works [7], [8], [9] focusing open-sourced LLMs, their results show that as the input length approach the context window size of an LLM, the perplexity of the output will increase. In cases of normal chat, the perplexity can be viewed as reversed-fluency, which may be less of an issue for the purpose of communication. However, identifying UPR variable is a reasoning task, so that such "surprise" is not acceptable. Another work [10] conclude that the LLMs are **not** robust to changes in the position of relevant information: specifically that LLM will fail to retrieve information if the relevant information is in the middle of the input context. As

the function length grows, the potential UPR variables and its relevant code statements could "lost in the middle". In the field of LLM assisted program analysis, similar observation could be found. In SWE-Bench [11], authors also observe that when using LLMs to resolve GitHub issues (i.e., code repair task), Claude 2 and other LLM's performance drops as the input context length increases.

#### D. Problem Statement

Despite the mixed results shown in Section III-C, it is clear that LLMs can help with circumventing the challenge of lack of explicit pattern for UPR variables. However, as shown in Section III-C, the LLM is neither panacea nor silver bullets, so that adopting it naively could lead to failure in real world settings.

We aim to develop a UPR variable detection method that can significantly reduce the amount of time needed to identify UPR variables. Specifically, to address the challenges and issues mentioned in Section III-B, our problem statement is as followed: *Given a program and a capable LLM, how to leverage the LLM to find the UPR variables in the program, so that substantially more UPR variables could be identified while maintaining a reasonable false positive rate.*

Based on our experiments and observations, the most imminent challenge when adopting LLM to identify UPR variables is: when finding UPR variables in large code bases, the long input context length will deteriorate the robustness of LLM, leading to false positives and false negatives.

A straightforward solution would be trimming the input function body by using, for example, information retrieval [21], [11]. But it is difficult to decide what and how much code to be removed so that enough information is left to determine whether a variable is UPR. Furthermore, even if there exists a perfect trimming strategy, the remaining context of a variable could be already too long for an LLM to perform robustly.

In this paper, we propose a novel LLM-assisted workflow, which can output the **UPR score** of a variable, to assist a human analyst in identifying UPR variables. The UPR score is a quantitative metric assessing the degree of relationship between a variable and user privileges. To alleviate LLMs' performance deterioration issue related to the input context length, we only leverage LLMs to generate the **UPR ratings for code statements** instead of whole chunks of code, so that the input context length of the prompts could be minimized. Other necessary aspects to determine the UPR score of a variable are addressed by traditional program analysis techniques, specifically by constructing the program dependence graph (PDG) [14], [22] at code statement level.

## IV. METHOD

### A. Issue of Binary Verdict

Before introducing the detail of our workflow, we would like first elaborate the reasons why we choose to output a UPR score instead of a UPR verdict. Essentially, the diversity among the UPR variables are substantial, which implies that they may have different degrees of relationship to the user

privileges. Therefore, some variables in a program may be at the borderline, meaning that they could be either UPR or not, depending on the context of the program/function as well as the stored values at runtime.

Code 3 shows an example of such cases. In an FTP server, a common utility function is to read the `stat` information of files, including `uid` and `gid`. Unlike passwords, `uids` are less sensitive and just leaking the values of the `uids` may not cause severe consequences. Typically, a variable saving `uids` is UPR only when it is used as the argument of function calls such as `chown` or `setuid`. Let us assume a function called `get_file_uid`, whose returned value is going to be stored in a local variable called `file_uid`. This function could be used by both the `ls` command to list the `dir` for the user, and the `chown` command to change the owner for a file. Needless to say, only in case of `chown` command will the variable `file_uid` be UPR. In this case, it is very difficult to make a UPR verdict for the local variable `file_uid` in function `get_file_uid`.

Although the borderline cases are relatively rare, they can distract the analysts if not being properly handled. By outputting a score rather than a verdict would alleviate this issue, and fortunately, the powerful LLMs make this possible.

### B. Workflow Overview

Figure 3 illustrates the complete workflow of our approach, which include 6 major steps. Starting from the source code of a target program as input, the workflow eventually output the UPR scores for all variables in the program.

① First, from the source code files, we generate the PDG for each function through static analysis. Each nodes in the PDG corresponds to a statement expression. ② The PDGs are then sliced into variable subgraphs, each of which corresponds to a single variable in the program. ③ Next, based on the variable subgraphs, we collect the code statements, creating a set of code statements that are either control and data dependent on or by the variables. ④ This set of code statements are subsequently rated by an LLM, using an initial prompt that is designed upon the purpose of the target program. The ratings shall reflect the significance of the statements in terms of their impacts toward potential non-control data attacks against the target program. ⑤ Then, the statement ratings will be used to update the variable subgraphs, serving as node score. ⑥ Finally, we compute the score of the variable by aggregating the node scores using a aggregation algorithm based on the graph topology. To confirm a UPR variable, subsequent manual review is necessary for those with highest ratings.

### C. PDG Generation

Since source code level PDGs are very well defined [14], we adopt the well-known definition of PDG and will not reiterate any formal definitions and specifications.

Essentially, we parse the source code files and generate a PDG for each function. As discussed in Section III, to control the length of the input length of the prompts as well as to balance the information carried by the graph topology

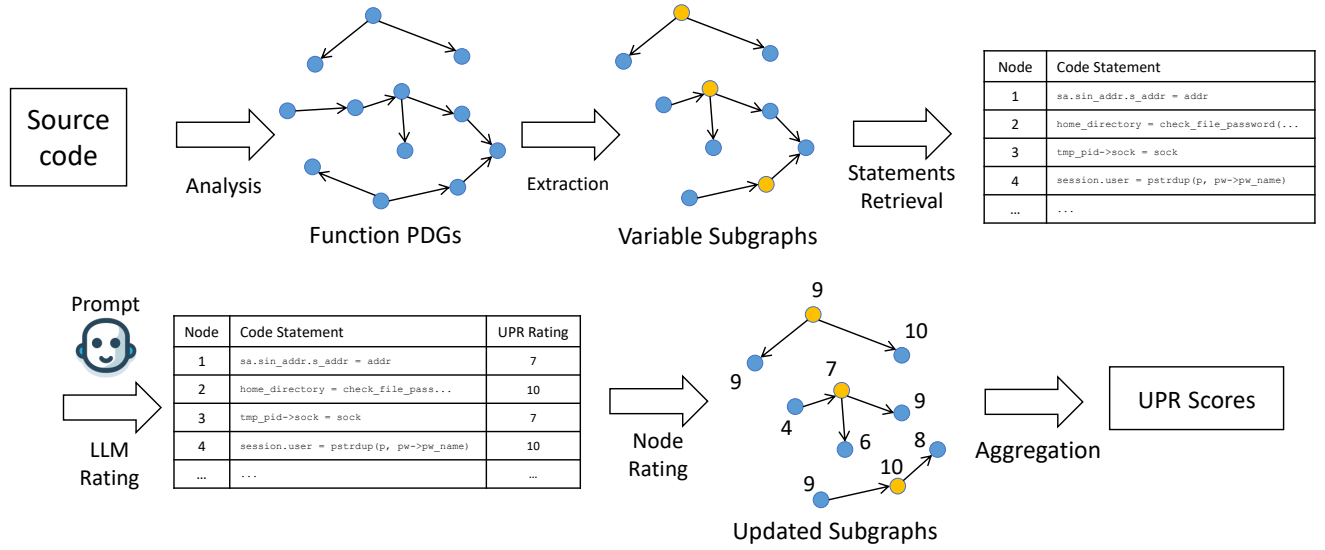


Fig. 3. Workflow

```

1  int get_file_uid(char* filename) {
2      int file_uid; // Is this UPR?
3      ...
4      return file_uid;
5  }
6  int do_ls(char* cmd) {
7      ...
8      int uid = get_file_uid(fname);
9      ...
10 }
11 int do_chown(char* cmd) {
12     ...
13     int uid = get_file_uid(fname);
14     if (uid == current_uid) {
15         ...
16         chown(pathname, new_uid, new_gid);
17         ...
18     }
19     ...
20 }

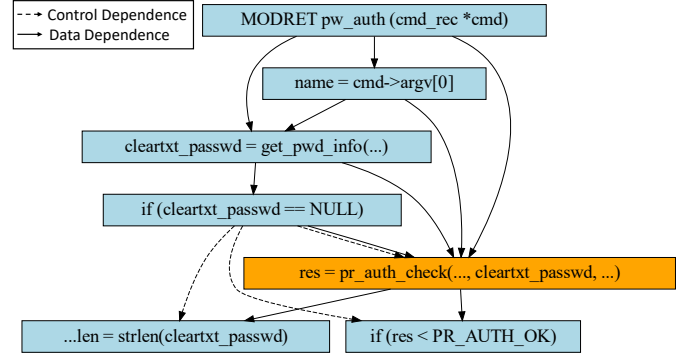
```

Code 3. An example of a borderline UPR variable.

and nodes, our PDG is designed to only have **simple statements** as nodes, including assignment, (function) call, assertion, goto, and return. In cases of a compound statement (i.e., block of simple statements), including for, while, do loops, switch, if, else branches, etc., the condition is treated as a single node (i.e., simple statement), and each simple statement in the block will also be a single node, with an control dependence edge connected to the condition node.

#### D. Variable Subgraph Extraction

From the function PDGs, we then extract variable subgraphs for each referred variable in the function, including global variables and local variables. Operating on the PDG, the extraction process essentially can be viewed as a customized program slicing [12] process, so that we slice the PDG in both backward and forward directions from the criterion nodes.

Fig. 4. A pruned variable subgraph for a single criterion node of variable `cleartxt_passwd` in program `proftpd`.

First, we identify all the variables referred in the function through AST analysis. For each variable, we then mark all the nodes corresponding to the statements referring it as criterion node (for program slicing). Since a variable could be referred multiple times in the function, intuitively each variable could have multiple criterion nodes, so that between a variable and its criterion nodes is there a one-to-many relation, as shown in Figure 3. Next, we extract the variable subgraph for each variable by exclusively including only immediate (i.e., 1-hop) neighbors of the corresponding criterion nodes, regardless the direction. One one hand, both directions are imperative, as whether a variable is UPR depends not only on its depender statements, but also on the dependee statements; on the other hand, statements that are not immediate depender/dependee of the criterion nodes are trivial, if not detrimental, to determine whether a variable is UPR, as further dependency relations are more noisy and less relevant. Figure 4 shows an pruned example of the extracted variable subgraphs, which only contains a single criterion node. Note that nodes are represented in different colors, where orange ones are criteria nodes. The variable subgraphs are extracted from the PDG, and therefore

they have two types of edges: control and data dependence edges.

Since nodes in the PDG are statements, it is completely possible that two variables are referred by a same set of statements, rendering identical variable subgraphs for different variables. In case one of them were UPR but the other were not, either false positive or false negative will present. This issue is observed in our experiments, and we have it discussed in Section VI-B

### E. Code Statement UPR Rating Using LLM

After variable subgraphs are constructed, we collect code statements from all the subgraphs, creating a set of statements to be rated by the LLM. To clarify, since the statements are collected from the variable subgraphs, there is no sequential order between the statements in the set. The LLM will evaluate and rate the statements one after another, so that the ratings of the statements are independent to each other. Each statement in the set is unique, so that the cost of using the LLM can be minimized.

To instruct the LLM with necessary information and context, our prompts consist two parts: an initial prompt consist of instructions, and the code statement to be rated. The initial prompt serves for two goals: 1) prompt the LLM to generate ratings for each statements; and 2) provide the specifications of the UPR property. The detail of the prompt design will be discussed in Section V.

Indeed, LLM is built upon neural networks after all, and thus it is a black-box. Therefore, the ultimate UPR ratings for code statement given by the LLM may be based on not only the criteria in the initial prompt, but also LLM's own "interpretations". With that being said, the code statement UPR ratings given by the LLM could be disagreed by a human analyst, so that our workflow need to be designed to tolerate this inconsistency.

After the code statements are rated, effectively their corresponding nodes in the variable subgraphs will receive UPR ratings, which can be further used to determine the UPR scores of the variables.

### F. Generating Variable UPR Score

The output of our workflow is a UPR score for each variable, which is obtained by aggregating the node scores in the updated variable subgraph. Recall that our variable subgraph is essentially a pruned PDG, so that the topology of the graph implies the dependence relationships between the code statements. As discussed in Section III-B, whether a variable is UPR pivots on not only the how the variable is used and defined, but also how the variable will affect other part of the program. Thus, the UPR score is computed by considering all the nodes in the updated variable subgraph.

We formalize the procedure of generating UPR scores for variables: Given a updated variable subgraph  $G = (N, E)$  of variable  $v$ , where  $N$  is the set of nodes and  $E$  is the set of edges in the graph, and a parameter  $\lambda$  weighing between the node score and its neighbors' score, we can compute the UPR score of a variable, as shown in Algorithm 1.

---

### Algorithm 1 Compute UPR score

---

**Require:**  $G = (N, E)$ ,  $\lambda \in [0, 1]$

```

1:  $S = \{\}$   $\triangleright$  Array to store aggregated score of all criteria nodes
2: for  $n \in N$  do
3:   if  $n$  is criterion node then
4:      $R = \{\}$   $\triangleright$  Array to store score of neighbors
5:      $s_c = \text{get\_node\_score}(n)$ 
6:      $V = \text{find\_neighbor\_nodes}(n)$ 
7:     for  $v \in V$  do
8:        $R = \{R, \text{get\_node\_score}(v)\}$   $\triangleright$  Append
9:     end for
10:     $s_v = \max(R)$ 
11:     $s_n = \text{Normalize}(s_c + \lambda s_n)$   $\triangleright$  Aggregated score
12:     $S = \{S, s_n\}$   $\triangleright$  Append
13:   end if
14: end for
15: return  $\max(S)$ 

```

---

Algorithm 1 is essentially a aggregation process, whose most important operation is at line 10, 11 and 15. There are two aggregating operations. The first one is at line 11 of Algorithm 1, where we compute the aggregated score of each criterion node. Do note that this score of a criterion node is affected by both the criterion node itself and its neighbors. Here, the value of  $\lambda$  is between 0 and 1, and the larger the more important neighbor score is. The normalize function ensures final UPR score is in a fix range, that is, between 0 to 10. The second aggregating operation is at line 15, so that we obtain a single UPR score for a variable from scores of multiple criteria nodes. It is worthwhile to note that, as mentioned in Section IV-E, our workflow needs to tolerate potential inconsistent ratings with human. Based on the principle of reduce false negatives, we use `max` as our aggregation function.

After collecting UPR scores for all variables, the variables will be sorted by their UPR scores from high to low. An analyst can then set up a threshold so that any variables with a UPR score above the threshold will considered as UPR variable candidates, subject to further manual analysis. As such, analysts do not have to manually check every single variables in a program.

## V. PROMPTING LLM

A fundamental component of our workflow is an LLM, which helps rate a code statement in terms of the degree of the relationship between the code statement and user privileges. Obviously, the quality of the rating result is one of the most important factors toward generating appropriate UPR scores and therefore helping analysts to identify UPR variables. However, there is no principled method to design prompts, due to the black-box nature of LLM. Therefore, we designed our prompt based on empirical case studies: we have selected 30 code statements from different programs of different lengths, and then gradually improve the prompt based on the human feedback of how well the UPR ratings are for the selected code statements.

As briefly introduced in Section IV-E, the prompt consists of two parts: initial prompt and the code statement to be rated. Since the code statement has to be presented as it is, the prompt engineering effort is mainly for the initial prompt. Eventually, we decide to adopt the following template for initial prompt:

I want you to act as a security critical code statement identifier. I will type a code statement and you will reply with a criticalness rating from 0 to 10 of the code statement. 0 is the least level of criticalness. Only reply with the criticalness rating and do not write explanation. Regarding what is a critical code statement, when it is tampered, it may result in:

1. **\*\*Malicious Goal 1\*\***
2. **\*\*Malicious Goal 2\*\***
3. ...

In this design, after describing our request to the LLM, we include a list of potential malicious goals that an attacker may be interested related to the user privileges. The list of malicious goals can be compiled based on the target programs, as different programs may have different kinds of user privilege related issues. In this paper, we compiled a list of common user privilege security incidents, generating a list of malicious goals used in the prompt:

1. One can bypass authentication and/or authorization
2. One can achieve behaviors that are being banned
3. One can gain elevated privileges (uid changes)
4. One can access files/data that are not accessible (chown related)
5. One can steal secrets (keys, passwd, etc.)

As will discussed shortly in Section VI, the UPR security incidents in selected programs in our experiments can be mostly covered by this list. In addition, the effect of the prompt design will be also studied and discussed in Section VI-E

## VI. EXPERIMENTS AND EVALUATION

We evaluate our approach in terms of both practicality and reliability, proposing 4 research questions: 1) How practical our approach is in terms of the false positive rate? 2) Can our approach find more UPR variables than heuristic-based method? 3) How to interpret distribution of the rating spectrum? 4) What is the effect of the prompt design? 5) How does it perform on proprietary programs? These research questions will be answered in later sub-sections.

Our evaluation uses a code base constructed using in total 7 programs: `bftpd`, `proftpd`, `vsftpd`, `ghttpd`, `nginx`, `sshd`, `sudo`. The programs are selected to ensure diversity in both program functionality (i.e., business logic) and size (i.e., lines of code).

To reduce both time cost and economical cost of our experiments, we did not analyze all the variables in the selected program. Specifically, for programs with a small code base, we run all variables though our method; whereas for

larger programs, we only include variables in the core modules (so that a minimum functional binary can be compiled). For example, `sshd` is one of the binary in the `OpenSSH` code base, so that in our experiments, all the variable in modules for other binaries (e.g., `ssh`, `ssh-keygen`, `scp`, etc.) are excluded.

### A. Experiment Setup and Implementation

The prototype of our tool leverages several existing tools. From the source code, the raw PDG is extracted using `Joern`. The extracted raw PDG is further processed using Python scripts to generate variable subgraphs. Subsequently, we use GPT-4 as the LLM code statement rater, as it is one of the best LLMs available. We send each statement along with the initial prompt discussed in Section V to the GPT-4, and collect the UPR ratings evaluated by the LLM. Finally, the UPR score of the variable is computed based on Algorithm 1 implemented in Python.

To reduce the number of variables may be missed, we conduct static analysis to identify variables in the program using both LLVM (at LLVM IR level) and `Joern` (at AST level). In total, the prototype of our tool used in the experiments are implemented in more than 1600 lines of code in Python and 1300 lines of code in C/C++.

### B. How practical is our approach regarding FPR?

It would be meaningless if the proposed method could not reduce manual efforts of the analyst. In particular, we want the analyst to find as many UPR variables as possible by examining as less variables as possible. Thus, without an adequate false positive rate (FPR) will never our approach be practical.

We evaluate the practicality of our approach in terms of the FPR by experiments. Simulating how our proposed tool may be used in real world, we first take the target program and evaluate the variables using our workflow discussed in Section IV-B. After obtaining the UPR scores of each variable, we set up a rating threshold of 9.0 (recall that the output UPR scores are between 0 and 10), above which are considered as UPR variable candidates. Then, the UPR variables are going through a manual analysis: only those variables which contain sensitive information, or/and can be exploited to cause security incidents (e.g., bypass authentication, gain additional privileges, etc.) are marked as UPR.

Program	Analyzed	Detected	FPR
<code>bftpd</code>	551	60	0.1500
<code>proftpd</code>	1153	235	0.1069
<code>vsftpd</code>	829	73	0.1311
<code>ghttpd</code>	71	3	0.0000
<code>nginx</code>	708	15	0.1333
<code>sshd</code>	679	195	0.1485
<code>sudo</code>	464	64	0.1905
All	4455	645	0.1349

TABLE I  
FALSE POSITIVE RATE OF OUR METHOD

The experimental results are shown in Table I. Since the selected programs are diverse regarding both their functionalities and their size, it is not surprising that the number of UPR variables varies between programs. We only conduct false positive analysis because manually identifying UPR variable is extremely time-consuming. By only focus on false positives can we only inspect 645 *Detected* variables. Were we to compute the false negative rate, instead 4455 *Analyzed* variables need to be inspected.

The results suggest that the FPRs are 13.49% on average. In other words, for every 100 of variables reported, less than 14 of them will be false positive. Nevertheless, when we take a glance at the number of variable *analyzed* in Table I, without our tool, in total 4455 of variables will need to be manually checked comparing only 645 are now need to be checked. Thus, we believe our tool can significantly reduce the manual efforts required from the human analysts.

After careful analysis of the false positives, we found that the most common cause for false positives is due to the granularity of the PDG. As mentioned in Section IV-D, it is totally possible that two variables have almost or exactly identical variable subgraphs. Using an example shown in Figure 4, `cleartxt_passwd` and `cleartxt_passwd_len` (abbreviated in Figure 4 to save space). In code statements where `cleartxt_passwd_len` is used, many of them also will use `cleartxt_passwd`. Eventually, after the variable subgraph extraction, the variable subgraph for `cleartxt_passwd` and `cleartxt_passwd_len` is similar. Unfortunately, though the password itself is clearly a UPR variable, the length of it is not, and therefore the similar variable subgraphs will cause the false alarm.

### C. Can our approach find more UPR variables than heuristic-based method?

One key motivation of our method is to identify more UPR variables than heuristic-based methods. To answer this research question, we design an experiment to compare our method with heuristic-based methods commonly adopted in the industry. Specifically, heuristic-based methods usually adopt two kinds of heuristics: 1) code pattern matching and 2) taint analysis.

Accordingly, we design a set of heuristic rules that can be applied using `Semgrep` framework, which is an open-source static analysis tool to find bugs and security vulnerabilities in the source code. It allows developers to write custom rules using a simple syntax to search for patterns in codebases. Since it supports syntax matching and taint analysis, we develop two rules containing several heuristics to discover UPR in the source code, as shown in Code 4 and Code 5.

```

1  patterns:
2    - pattern-either:
3      - pattern-inside: |
4        if (...)
5      - pattern-inside: |
6        switch (...)
7    - pattern: $VARIABLE
8    - metavariable-regex:
9      metavariable: $VARIABLE

```




Fig. 5. UPR positive variable reported by our method and the heuristic-based method.

```

10  regex: .* (?![A-Za-z]) (auth|authenticate|
    login|admin|authorize|banned|allowed|
    uid|gid|username|permission|chown|key
    |passwd|password) (?![A-RT-Za-rt-z]).*

```

Code 4. Syntax based detection rules for UPR.

```

1  mode: taint
2  pattern-source:
3    - patterns:
4      - pattern-either:
5        - pattern-inside: |
6          $IOFUNC(...)
7        - pattern-inside: |
8          $VAR = $IOFUNC(...)
9      - metavariable-regex:
10     metavariable: $IOFUNC
11     @\Hilighct@regex: scanf|gets|fscanf|
12     fgetc|fgets|getchar|fread
12 pattern-sinks:
13   - patterns:
14     - pattern-inside: |
15       seteuid (...)
16     - pattern-inside: |
17       setuid (...)
18   ...

```

Code 5. Taint analysis based detection rules for UPR.

Both rules are straightforward, so we only provide brief explanations. Specifically, the rule shown in Code 4 detects the variables that contain privilege-related keywords. The most important item in this rule is the regular expression at line 10 matching privilege-related keywords in the variable names:

```

.* (?![A-Za-z]) (auth|authenticate|login|admin|
  authorize|banned|allowed|uid|gid|username|
  permission|chown|key|passwd|password) (?![A-
  RT-Za-rt-z]).*

```

The rule in Code 5 is essentially taint analysis, which specifies the taint source and taint sinks. It can detect variables whose taint comes from inputs that are used in user privilege related function calls.

We then perform the UPR variable detection using the heuristic-based method on exactly the same code base described at the beginning of this section. To make this experiment practical, we only focus on positive variables that is either reported by our method or the heuristic-based method. Those variables reported by neither methods will be excluded from our analysis. In order to make a fair comparison, we divide the detected positive variables into the following three sets: (**Set A**) variables only reported by heuristic-based method, (**Set B**) variables only reported by our method, and (**Set C**) variables reported by both methods. The result of analysis is shown in Figure 5. According to the overlapped circles in Figure 5, it is clear that our method can find substantially more variables than the heuristic-based method can,

especially with a significantly lower false positive rate (13.49% vs. 54.53%). Do note that our method detected 645 variables as positive, using a fairly conservative threshold (i.e., 9.0/10.0), out of which 13.49% of them are false positives. In contrast, when we use the heuristic-based method introduced above, out of 413 (126 + 287), more than half of them (53.44%) are false positives. Indeed the heuristic-based method could be tweaked to be more conservative and our method can be more aggressive, but the differences in performance is apparent.

This experiment also exposes some false negative cases for our approach. A common false negative reason is the quality of the UPR ratings for code statements from GPT-4, specifically the mediocre UPR ratings. On one hand, the mediocre ratings may be caused by limited capability of the GPT-4; on the other hand, they may be also due to the insufficient amount of the context. For example, considering following code statement:

```
if(!access_check_file(&s_filter_str))
```

Similar to the situation described in Section IV-A, this is a case of "borderline UPR statement": whether it is UPR cannot be asserted without knowing the value may be stored in `s_filter_str`. In our experiments, we only consider a variable as a UPR candidate if the UPR score is at least 9, which is a relatively high threshold. Therefore, a few variables with mediocre-rating code statement in their variable subgraphs are missed. If such mediocre ratings are very common, it would indeed made our method less practical. We will have more discussions about the UPR ratings from the GPT-4 in Section VI-D.

#### D. How well can an LLM rate the code statement?

Toward the core of our approach is the UPR rating of the code statements generated by LLMs, which lead to a research question: how well can an LLM rate a code statement in terms of their relationship with user privileges? Due to the blackbox nature of the LLM, it is futile to inspect the LLM internally. Since there is no benchmark regarding the task of rating a code statement for LLMs, instead of studying the question directly, we decide to focus on the distribution of the ratings given by the GPT-4 in our experiments.

We firstly set up a series of hypotheses to validate, based on the observations of the UPR variables:

- O1:** UPR variables, especially those ones critically related to user privileges, are rare.
- O2:** The majority of variables are **not** UPR.
- O3:** A few variables may or may not be UPR, depending on the context of the program. (See Section IV-A)

Based on the above-mentioned observations of UPR variables, regarding the UPR scores of code statements, we have following hypotheses:

- H1:** Code statements with higher UPR scores are the minority.
- H2:** The majority of code statements should receive low scores.
- H3:** A few code statements should have mediocre UPR scores, because a single code statement contains no context about the containing function, as well as the caller(s) of the function.




Fig. 6. Distribution of the statements UPR ratings from GPT-4

Obviously, **H1** is derived from **O1**, and it is natural to believe that the number of UPR variables and the number of high-UPR-score code statements are closely correlated. After all, whether a variable is UPR is depending on whether it is directly or indirectly related to a code statement related to user privileges (i.e. with a high UPR score). For similar reasons, **H2** is derived from **O2**.

**H3** is derived from **O3**, which reflect those variables whose UPR scores are depending how they may be used in the program, as well as the business logic of the program. As discussed in Section IV-A, though they are uncommon, when analysts encounter these variables, it would be difficult to make a verdict regarding whether the variables are UPR variables. and the `chown` command to change the owner for a file. In our workflow, only the code statement itself is given to the LLM when generating the UPR ratings, so that the context of how the variable is defined and used may be missing. We therefore expect the LLM may give mediocre ratings in these cases, as a human analyst would also prefer to do so.

Section VI-D shows the distribution of the UPR ratings given by GPT-4. The area marked blue are code statements with low UPR ratings, whereas the area marked in red are those ones with high UPR ratings. The grey areas are considered as mediocre ratings. According to Section VI-D, it is evident that all three hypotheses are validated:

- 1) Only few statements have high UPR ratings (small red area).
- 2) Large number of statements have low UPR ratings (large blue area).
- 3) Only a few statements receive mediocre ratings (tiny grey area).

Analysis based on the distribution of the UPR ratings may neglect those corner cases. Nevertheless, Section VI-D shows that the selected LLM, GPT-4, is capable of generating appropriate UPR ratings for code statements.

#### E. What is the effect of the prompt design?

As discussed in Section V, our prompt template include a list of malicious goals to provide the LLM with more specifications about UPR variables. While the presence of this list can improve the UPR ratings for the code statements, it is intriguing to study whether the presence of this list would affect the **final UPR scores for variables**.

Accordingly, we evaluate our method using a modified prompt template for code statement UPR ratings acquisition. Specifically, the prompt template are modified in two ways: one kind only contains partial list of the malicious goal; whereas the other kind contains no list of malicious goals. In the modified prompt template containing partial list, two items in the original list are kept: 1) One can access files/data that are not accessible (chown related); 2) One can steal secrets (keys, passwd, etc.) We use the two modified prompts in our workflow to re-detect the confirmed UPR variables found in previous experiment.

The experiment is evaluated by measuring the percentage of the confirmed UPR variables are re-detected using different modified prompts, which is shown below:

- Prompt with partial list: 43.89%
- Prompt with no list: 54.30%

It may seem to be unexpected that the result when using completely no list is better. As mentioned in Section IV-E, LLM may have their own "interpretations" for some words or phrases in the prompts. Therefore, if there is no given specification for UPR variable, the UPR ratings may become very unpredictable. For example, in the initial prompt, we used the term "security critical", which is a rather broad phrase, so that GPT-4 can have various interpretation of such a term. Nevertheless, it is still evident that providing specification such as a list of malicious goal definitely enables our method to uncover more UPR variables.

## VII. DISCUSSION AND LIMITATION

### A. Proprietary Programs

All the LLMs are trained on massive amount of data, which probably include almost all the open source source code. Consequently, from machine learning perspective, essentially the UPR rating for code statements in open source programs are in-distribution data (as oppose to out-of-distribution data). Though it is believed that LLMs perform well in term of the generalizability, nevertheless out-of-distribution data still may affect the performance of the model. In our case, the out-of-distribution data is essentially the code statement in proprietary programs.

Since it is not realistic to comprehensively evaluate our approach on proprietary programs, we have conducted a simple experiment about using GPT-4 to re-rate the code statements, but with the names in the statements changed. For example, a function in `vsftpd` named `vsf_sysutil_read` will be renamed as `xyz_sysutil_read`, pretending it is a function in different program. We have randomly selected 30 code statements to have the GPT-4 to re-rate, after renaming all the names. Our result shows that the average difference between the ratings is -0.45 with a standard deviation of 0.94, which implies that the UPR rating does change, but only moderately. Nevertheless, thought it is difficult to draw rigorous conclusion regarding the performance on proprietary program, as LLMs are considered as black boxes, this may becomes a limitation when adopting our method to some proprietary programs.

### B. Cost of LLM Service

Our method will prompt LLM for each code statements, so that the number of token used in the input is dominated by the length initial prompt. While the disadvantage in terms of the cost is that every single statement will need to pay the price of the initial prompt, the advantage is that the output token used, which is typically more expensive than that of input, could be just 1 or 2, because we only want the LLM to output the UPR ratings.

## VIII. CONCLUSION

In this paper, we introduce a large language model (LLM) workflow that can assist analysts in identifying such UPR variables, which is a very time-consuming task. Our tool will analyze all the variables in a program to output a UPR score for each variable, which is the degree of relationship between the variable and user privileges. When our approach is used in real world, those variables with high UPR scores will then be manually investigated by the analysts. Our experiments show that in a typical setting (i.e., UPR score >0.8), the false positive rate (FPR) is only 13.49%, while the number of UPR variables found is significantly more than that of the heuristic based method.

## REFERENCES

- [1] H. Shacham, "The Geometry of Innocent Flesh on the Bone: Return-into-libc Without Function Calls (on the x86)," in *CCS*, 2007.
- [2] S. Checkoway, L. Davi, A. Dmitrienko, A.-R. Sadeghi, H. Shacham, and M. Winandy, "Return-Oriented Programming Without Returns," in *Proceedings of the 17th ACM Conference on Computer and Communications Security*, 2010.
- [3] H. Hu, S. Shinde, S. Adrian, Z. L. Chua, P. Saxena, and Z. Liang, "Data-Oriented Programming: On the Expressiveness of Non-control Data Attacks," in *Proceedings of the 37th IEEE Symposium on Security and Privacy*, 2016.
- [4] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang *et al.*, "Codebert: A pre-trained model for programming and natural languages," *arXiv preprint arXiv:2002.08155*, 2020.
- [5] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, A. Svyatkovskiy, S. Fu *et al.*, "Graphcodebert: Pre-training code representations with data flow," *arXiv preprint arXiv:2009.08366*, 2020.
- [6] X. Li, Q. Yu, and H. Yin, "Palmtree: Learning an assembly language model for instruction embedding," *arXiv preprint arXiv:2103.03809*, 2021.
- [7] O. Press, N. Smith, and M. Lewis, "Train short, test long: Attention with linear biases enables input length extrapolation," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=R8sQPpGCv0>
- [8] S. Chen, S. Wong, L. Chen, and Y. Tian, "Extending context window of large language models via positional interpolation," *arXiv preprint arXiv:2306.15595*, 2023.
- [9] G. Xiao, Y. Tian, B. Chen, S. Han, and M. Lewis, "Efficient streaming language models with attention sinks," *arXiv preprint arXiv:2309.17453*, 2023.
- [10] N. F. Liu, K. Lin, J. Hewitt, A. Paranjape, M. Bevilacqua, F. Petroni, and P. Liang, "Lost in the middle: How language models use long contexts," *arXiv preprint arXiv:2307.03172*, 2023.
- [11] C. E. Jimenez, J. Yang, A. Wettig, S. Yao, K. Pei, O. Press, and K. Narasimhan, "Swe-bench: Can language models resolve real-world github issues?" *arXiv preprint arXiv:2310.06770*, 2023.
- [12] M. Weiser, "Program slicing," *IEEE Transactions on software engineering*, no. 4, pp. 352–357, 1984.
- [13] Z. Yu and V. Rajlich, "Hidden dependencies in program comprehension and change propagation," in *Proceedings 9th International Workshop on Program Comprehension. IWPC 2001.* IEEE, 2001, pp. 293–299.
- [14] J. Ferrante, K. J. Ottenstein, and J. D. Warren, "The program dependence graph and its use in optimization," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 9, no. 3, pp. 319–349, 1987.

- [15] S. Kang, J. Yoon, and S. Yoo, "Large language models are few-shot testers: Exploring llm-based general bug reproduction," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2312–2323.
- [16] C. Lemieux, J. P. Inala, S. K. Lahiri, and S. Sen, "Codamosa: Escaping coverage plateaus in test generation with pre-trained large language models," in *International conference on software engineering (ICSE)*, 2023.
- [17] C. S. Xia, M. Paltenghi, J. L. Tian, M. Pradel, and L. Zhang, "Universal fuzzing via large language models," *arXiv preprint arXiv:2308.04748*, 2023.
- [18] Nergal, "The Advanced Return-into-lib(c) Exploits," <http://phrack.com/issues.html?issue=67&id=8>.
- [19] T. Bletsch, X. Jiang, V. W. Freeh, and Z. Liang, "Jump-Oriented Programming: A New Class of Code-reuse Attack," in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, 2011.
- [20] S. Chen, J. Xu, E. C. Sezer, P. Gauriar, and R. K. Iyer, "Non-Control-Data Attacks Are Realistic Threats," in *Proceedings of the 14th USENIX Security Symposium*, 2005.
- [21] S. Robertson, H. Zaragoza, and M. Taylor, "Simple bm25 extension to multiple weighted fields," in *Proceedings of the thirteenth ACM international conference on Information and knowledge management*, 2004, pp. 42–49.
- [22] T. M. Austin and G. S. Sohi, "Dynamic dependency analysis of ordinary programs," in *Proceedings of the 19th annual international symposium on Computer architecture*, 1992, pp. 342–351.