

On Detecting and Measuring Exploitable JavaScript Functions in Real-World Applications

MARYNA KLUBAN, Concordia University, Canada

MOHAMMAD MANNAN, Concordia University, Canada

AMR YOUSSEF, Concordia University, Canada

JavaScript is often rated as the most popular programming language for the development of both client-side and server-side applications. Because of its popularity, JavaScript has become a frequent target for attackers who exploit vulnerabilities in the source code to take control over the application. To address these JavaScript security issues, such vulnerabilities must be identified first. Existing studies in vulnerable code detection in JavaScript mostly consider package-level vulnerability tracking and measurements. However, such package-level analysis is largely imprecise as real-world services that include a vulnerable package may not use the vulnerable functions in the package. Moreover, even the inclusion of a vulnerable function may not lead to a security problem, if the function cannot be triggered with exploitable inputs. In this paper, we develop a vulnerability detection framework that uses vulnerable pattern recognition and textual similarity methods to detect vulnerable functions in real-world JavaScript projects, combined with a static multi-file taint analysis mechanism to further assess the impact of the vulnerabilities on the whole project (i.e., whether the vulnerability can be exploited in a given project). We compose a comprehensive dataset of 1,360 verified vulnerable JavaScript functions using the Snyk vulnerability database and the VulnCode-DB project. From this ground-truth dataset, we build our vulnerable patterns for two common vulnerability types: prototype pollution and Regular Expression Denial of Service (ReDoS). With our framework, we analyze 9,205,654 functions (from 3,000 NPM packages, 1892 websites and 557 Chrome Web extensions), and detect 117,601 prototype pollution and 7,333 ReDoS vulnerabilities. By further processing all 5,839 findings from NPM packages with our taint analyzer, we verify the exploitability of 290 zero-day cases across 134 NPM packages. In addition, we conduct an in-depth contextual analysis of the findings in 17 popular/critical projects and study the practical security exposure of 20 functions. With our semi-automated vulnerability reporting functionality, we disclosed all verified findings to project owners. We also obtained 25 published CVEs for our findings, 19 of them rated as “Critical” severity, and six rated as “High” severity. Additionally, we obtained 169 CVEs that are currently “Reserved” (as of Apr. 2023). As evident from the results, our approach can shift JavaScript vulnerability detection from the coarse package/library level to the function level, and thus improve the accuracy of detection and aid timely patching.

CCS Concepts: • **Security and privacy** → **Software and application security**.

Additional Key Words and Phrases: JavaScript security; vulnerability detection; vulnerable functions; prototype pollution; Regular Expression Denial of Service; static analysis; semantic pattern detection; taint analysis

ACM Reference Format:

Maryna Kluban, Mohammad Mannan, and Amr Youssef. 2023. On Detecting and Measuring Exploitable JavaScript Functions in Real-World Applications. *ACM Trans. Priv. Sec.* 1, 1, Article 1 (January 2023), 36 pages. <https://doi.org/10.1145/3630253>

Authors' addresses: Maryna Kluban, mkluban@gmail.com, Concordia University, , Montreal, Canada; Mohammad Mannan, m.mannan@concordia.ca, Concordia University, Concordia University, Montreal, Canada; Amr Youssef, youssef@ciise.concordia.ca, Concordia University, Concordia University, Montreal, Canada.

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.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2471-2566/2023/1-ART1 \$15.00
<https://doi.org/10.1145/3630253>

1 INTRODUCTION

JavaScript is a programming language used in 97.5% of all web applications [120]. Such applications provide users with direct access to the source code, and are therefore a very common target for attackers. According to GitHub [9] and StackOverflow [111], JavaScript has been a dominating language for software development for at least the last 7 years. Moreover, JavaScript language can be used in both client-side and server-side applications, and has a large amount of frequently used development frameworks, such as React (reactjs.org), Angular (angular.io), Vue.js (vuejs.org) (client-side) and Express.js (expressjs.com), Koa (koajs.com), Nest.js (nestjs.com) (server-side).

Due to its wide usage, JavaScript is a very common target for attackers. Client-side attacks occur on a user's system when the user initiates malicious actions knowingly (as an attacker), or unknowingly (as a victim). These attacks exploit vulnerabilities in the client-side JavaScript code, which is fully accessible to the user. Client-side attacks include Cross-site Scripting (XSS), data infiltration, content injection, etc. In contrast, server-side attacks target the vulnerabilities in server-based JavaScript applications that manage databases, web/mail/file servers, etc. These attacks include SQL injection, improper access control, authentication bypass, Denial of Service (DoS), etc. The consequences of exploiting vulnerabilities in JavaScript source code include information theft or forgery [4], malicious code injection [109], redirection to attacker-controlled sources [88], disruption of an application functionality [89] and much more.

Previous work [3, 17, 18, 23, 27, 126, 127, 131] on JavaScript vulnerabilities mostly cover the propagation of vulnerable NPM packages among real-world projects, primarily based on the project dependency information. While such work is very useful in identifying projects with vulnerable dependencies, they do not provide fine-grained information on using actual vulnerable functions from the vulnerable packages. This leads to the flagging of projects as vulnerable due to their use of vulnerable dependencies, when in reality, many such projects (73.3% according to Zapata et al. [126]) are not vulnerable as they do not actually use the vulnerable functions.

On the other hand, only a few studies rely on code-based approaches for JavaScript vulnerability detection. Ferenc et al. [30] use static code metrics, generated for each function by static analysis tools (OpenStaticAnalyzer [85] and escomplex [28]), as features for machine learning (ML) algorithms that predict the probability of the function being vulnerable. Mosolygó et al. [79] use generalized representations of code lines, and calculate vulnerability likelihood based on cosine distance between vulnerable and analyzed code lines. However, results from both approaches are not very encouraging (F1-measure of 0.7 for [30], and 97.3% false positives for [79]).

The issue of insufficient vulnerable code datasets also hinders research in this area. Ferenc et al. [30] compiled the first publicly accessible dataset of JavaScript functions which includes 1,496 functions marked as vulnerable. Later, Mosolygó et al. [79] reduced the dataset from [30] to 443 vulnerable functions by manually filtering out false positives. We still found several non-vulnerable functions in this filtered dataset after briefly examining them manually.

The objective of our study is to assess active JavaScript projects in the real-world to find vulnerable functions in their source code. We gather JavaScript code from three different sources: NPM packages, Chrome web extensions, and top popular websites. Then we develop a vulnerability detection framework based on the following approaches: (1) vulnerable pattern recognition, which uses our manually developed patterns to perform a semantics-based search in code; (2) search based on textual similarity, which uses content-sensitive and cryptographic hash comparison; and (3) static taint analysis to verify the exploitability of the vulnerability.

For the first approach, we utilize a static analysis tool Semgrep [97], which allows us to develop advanced patterns based on JavaScript semantics. For the second one, we tokenize each function from vulnerable and real-world datasets and generate a content-sensitive hash value using the

SimHash [1] algorithm along with a SHA-1 cryptographic hash value. For each hash value from the real-world dataset, we then search for matches in the vulnerable function dataset. In the third approach, we develop a static taint analysis tool to verify that a flagged vulnerability is reachable by an attacker in the related project. We combine an Abstract Syntax Tree (AST) representation with our proposed File Dependency Graph (FDG), which enables multi-file taint-tracking. By combining these three approaches into a vulnerability detection and verification framework, we achieve good scalability from the first two approaches, and excellent precision from the third one.

Contributions. Our contributions can be summarized as follows:

- (1) We automatically crawl for vulnerable JavaScript functions from Snyk [103] and VulnCode-DB [34], allowing us to add the newly reported vulnerable functions and keep our dataset up-to-date. Then, we create vulnerability patterns and a web-based tool for efficient manual verification of vulnerable functions. In the end, we compose an informative, semi-automatically verified dataset of 1,360 JavaScript vulnerable functions.
- (2) We develop an experimental vulnerability detection and verification framework that consists of the combination of pattern and textual-similarity-based approaches, and of a static taint analyzer with a novel representation of file dependency graphs. The framework includes our manually developed rule sets for vulnerable pattern detection of two common JavaScript vulnerability types: prototype pollution and Regular expression Denial of Service (ReDoS).
- (3) We gather a large dataset of 9,205,654 JavaScript functions from active real-world projects from three different application types (NPM packages, Chrome extensions, and top websites). This collection process is also fully automated. We use this dataset in our evaluation framework.
- (4) We detect 124,934 vulnerable functions with an estimated average precision of 94.5% (based on manual verification of a small subset). We perform a case study on 17 popular/critical projects that contain 20 vulnerable functions; all these functions flagged by our framework are found to be, indeed, exploitable.
- (5) With our taint analysis, we identify 301 cases from 134 NPM packages (5.7% of all findings in NPM packages), which are exploitable in the project context. Manual verification of 100 cases detected no false positives produced by the taint analysis mechanism.
- (6) To deal with a large number of disclosure notices, we develop a semi-automated technique to report our findings. We firstly search for duplicates of our findings in the CVE database (and identify 19 cases), and then automatically compose readable vulnerability reports for the remaining 290 findings and send them to 112 responsible project developers. For 11 of these findings, we received a response from the developers, who are expected to submit these findings for CVEs through the GitHub Security Advisory [44]; we submitted the remaining 279 cases directly to MITRE [78].
- (7) We obtain 25 published CVE IDs as of writing (Apr. 2023), with 19 severity scores “Critical” and six scores “High”. Additionally, we obtain 169 CVE IDs that are currently in the “Reserved” status. We are working with the package owners to responsibly publish the CVEs.
- (8) We open-source our framework code and datasets at: <https://github.com/Marynk/JavaScript-vulnerability-detection>.

Differences with the ACM ASIACCS version [63]. This article is built on our initial findings reported in our conference publication. We make the following substantial changes. (1) We extend our initial framework with a static taint analyzer to effectively verify the exploitability of the identified vulnerable functions in NPM packages. For this approach, we develop our own data representation called File Dependency Graph to achieve multi-file tainting (to the best of our knowledge, no such functionality is present in open source taint analyzers for JavaScript) and combine it with an AST representation. As a result of applying the extended approach to 5,389 findings from NPM packages,

we verify the exploitability of 309 findings (from 134 NPM packages). (2) We utilize our vulnerable function detection framework (based on Semgrep patterns) for analyzing a large-scale dataset of real-world functions, processing 8,409,742 more functions from NPM packages, web extensions, and websites, compared to the conference paper. (3) We perform an experimental comparison of three existing static pattern-based vulnerability detectors with our Semgrep-based approach and discuss in detail the results of the evaluation. (4) We develop a semi-automated vulnerability reporting technique that checks for duplicates in the CVE database and sends readable reports by email to the responsible developers. We also discuss the process of responsible disclosure of the vulnerabilities and highlight the challenges we faced when dealing with a large number of vulnerabilities across several projects. (5) We submitted 279 findings to MITRE [78] and received 194 CVE IDs; 25 of them are already (as of Apr. 2023) published and have severity ratings critical (19 vulnerabilities) and high (six vulnerabilities).

2 BACKGROUND

In this section, we provide a brief background on vulnerability detection based on static analysis, and the JavaScript vulnerability types that we consider.

Vulnerability detection in source code is an active research area, and most existing approaches can be classified into two main categories: static and dynamic code analysis [60]. Static source code vulnerability analysis examines the code without running it. This approach has more coverage than the dynamic one, as it can find vulnerable code regardless of the execution conditions. Moreover, static analysis is usually more efficient and therefore more suitable for analyzing a large codebase. On the other hand, static analysis tools are more likely to produce higher false positive rates, as they do not actually follow the program logic, therefore omitting code paths that may mitigate the vulnerability [24]. Dynamic code analysis examines an application during (or after) its execution. Tools that use dynamic analysis for vulnerability detection are given specific rules and inputs for the code to run with. This approach is usually more precise and accurate, however, it suffers from such disadvantages as low code coverage, poor scalability, slow speed, etc. Additionally, taint analysis techniques, which track the propagation of the tainted vulnerable values, can be applied either directly to the code lines (statically), or to the program during its execution (dynamically).

2.1 Vulnerability detection with static code analysis

We divide static vulnerability analysis methods into two: (1) textual similarity approaches which use syntactical structures of the code to reveal new vulnerable syntax based on records of previously discovered vulnerable code; and (2) semantic similarity approaches which use previously discovered vulnerable code to extract vulnerable behaviors and develop patterns of this behavior in an abstracted representation. We also briefly discuss the static taint analysis approach that we use for distinguishing exploitable functions from the flagged vulnerable functions.

2.1.1 Textual similarity methods. The main principle of textual similarity methods is finding matches between examined real-world code instances and a known vulnerable code dataset. Depending on the pursued similarity degree of the code, several approaches are used. To detect slightly modified code (e.g., different layout, renamed variables), the instances are usually represented in a generalized version, e.g., bag-of-tokens, Abstract Syntax Trees (AST), and then compared either directly, or using transformations like cryptographic hashing or vectorization (e.g., with Word2Vec [76]). To match code with more significant differences, other abstract representations can be used, such as Control Flow Graphs (CFG), and Program Dependency Graphs (PDG). Some studies [16, 49] also use context-sensitive hashing (i.e., hashes, which depend on the content) to allow limited alterations in the examined code. However, this approach may introduce false

positives in detecting vulnerable functions, because it cannot distinguish a vulnerable piece of code from a patched one if the code modification is minor. This limitation is also applicable to Machine Learning (ML) based approaches, which create signatures for code instances based on their syntactic features, and then compare the signatures between vulnerable and targeted functions.

2.1.2 Semantic similarity methods. These methods detect semantic (functional) similarities by searching for vulnerability patterns and can be divided into two categories based on how patterns are developed: manual and automated. Compared to textual similarity approaches, these methods are generally better suited to functions with significant implementation differences, as they mostly focus on abstracted code structures without relying on the code text itself.

Manually created vulnerable patterns are developed by researchers based on their expertise. Each vulnerable pattern targets a specific implementation of one vulnerability type (e.g., XSS, SQL injection, DoS). Abstract code representations such as AST, and Code Property Graphs (CPG) can also be used. However, instead of performing a direct comparison with previously reported vulnerabilities, templates are generated based on useful properties from these representations. As part of our work, we manually develop vulnerable patterns as well. More precisely, we utilize Semgrep [97] – a static analysis tool, which allows us to develop patterns with regard to the language semantics based on AST representations of the code.

Automatically extracted vulnerable patterns (or features) are usually the result of machine learning algorithms (e.g., approaches in [91, 102]). Certain features, which are supposed to make a function vulnerable, are extracted by analyzing a large set of the representations of known vulnerable functions. While these tools aim to reduce manual effort, we did not choose machine learning approaches as such algorithms rely on a large dataset of vulnerable and patched functions with clear ground truth which, to the extent of our knowledge, does not exist for JavaScript.

2.1.3 Taint analysis. Taint analysis is a type of dataflow analysis that examines the path of input data through a program. Specifically, inputs from defined untrusted sources (i.e., inputs that can be controlled by the user) are tainted and followed from the “source” (an initial place in code, where the input was introduced) to the “sink” (potentially vulnerable functionality that can be exploited if the input is malicious). The objective is to determine whether a function containing a potential exploit is reachable by an attacker with her chosen inputs. Therefore, this approach requires correct identifications of “vulnerable sinks” and possible “malicious data sources”. Taint analysis can be implemented both statically and dynamically, and we choose to use a static approach. In terms of vulnerability detection, static taint analysis commonly refers to a traversal through an abstracted code representation that contains information on the data nodes and their relationships [59, 64, 67, 72]. The traversal instrument uses a set of language-specific taint propagation rules for every operation in the code. As a result, we can determine the behavior of the dataflow and whether the tainted input can reach a vulnerable functionality and potentially exploit it.

2.2 JavaScript vulnerability types

For JavaScript projects, the most frequent vulnerability reports in the Snyk vulnerability database [103] belong to four vulnerability groups (as enumerated in July 2021, see Table 7 in the appendix for all vulnerability types): cross-site scripting (XSS), code injection (CI), prototype pollution (PP), and Regular Expression Denial of Service (ReDoS). We discuss prototype pollution and ReDoS, as we consider detecting these two vulnerability groups in our work for the following reasons. (1) They belong to the list of the most frequently reported vulnerabilities. (2) Unlike

XSS and CI, they have a very specific set of mitigating patterns, allowing the elimination of “protected” functions and therefore help us avoid false positives. (3) The implementations of both of the vulnerability types are mostly consistent, so we can ensure good coverage.

2.2.1 Prototype Pollution. This vulnerability occurs in JavaScript when the “reserved” object keys are reassigned. In JavaScript, all data types are essentially objects (including functions and primitives). All objects have common “root” properties, which are `__proto__`, `constructor` (for objects created using the “new” operator, e.g., `new Date()`), and `prototype` for function objects. The attacker manipulates these properties by tampering with their values. Once one of the root properties is changed for one object, it is changed for all JavaScript objects in a running application, including those created after property tampering. The prototype pollution attack occurs when the objects receive properties and/or values that they are not designed to have. For example, a common way to represent a user on the server side of the web application is in an object with the following structure: `{"name": "Jane Doe", "age": 30, "education": {"primary": true, "secondary": false}}`. To change their education information, a user sends a request with the following information: `{"education": {"secondary": true}}`. This data then gets recursively merged into the user’s record. If the functionality that performs the merging operation does not check for the validity of the received data, the attacker can send a forged request, for example: `{"__proto__": {"isAdmin": true}}`.

When merging object properties, the program performs the following assignment operation: `userID.__proto__.isAdmin = true`. As a result, all users in the system will inherit the `isAdmin` property by default, which grants them full access to the system. Depending on the implementation, prototype pollution vulnerability can cause several attack types, including XSS, remote code execution, DoS, and SQL injections [25].

To avoid prototype pollution attacks, modification of an object’s root keys should be prohibited. This can be done while creating the object: by “freezing” it so that it becomes immutable, e.g., by using `Object.create(null)` which replaces the object’s prototype with `null`. In addition, the developers can include explicit checks of the object properties and performed operations, where the property names equal to `__proto__`, `constructor` and `prototype`.

2.2.2 Regular Expression Denial of Service. The second vulnerability type that we develop patterns for is ReDoS. It is a specific case of a DoS attack that happens when a program runs a user’s input through an evil regular expression [123], which takes exponential time to process specially-crafted complex strings. This usually happens when operators in the regular expression are used in a particular combination. For example, due to the irresponsible use of the repetition operator ‘+’ a regular expression `^(([a-z])+.)+[A-Z]([a-z])+$` will result in a long execution loop, which may cause denial of service, if ran on the input ‘aaaaaaaaaaaaaaaaaaaaaaaaa!’ [123]. The protection measures for ReDoS are either sanitizing the input (e.g., limiting the length or discarding repetition patterns), or rewriting regular expressions without using dangerous combinations of operators (such as ‘+’ and ‘*’, ‘?’ and ‘>’, ‘?’ and ‘=’).

3 METHODOLOGY: DATASET OF VULNERABLE FUNCTIONS

To detect vulnerable JavaScript functions in the wild, we first need a dataset of such functions. Since we are unable to find a suitable dataset, we develop one, relying on the existing approaches for function collection, and developing our own methodology to verify the dataset.

3.1 Vulnerable function dataset preparation

In this section, we describe the process of creating our vulnerable functions dataset. First, we collect possibly vulnerable functions from major vulnerability sources. We then analyze the collected

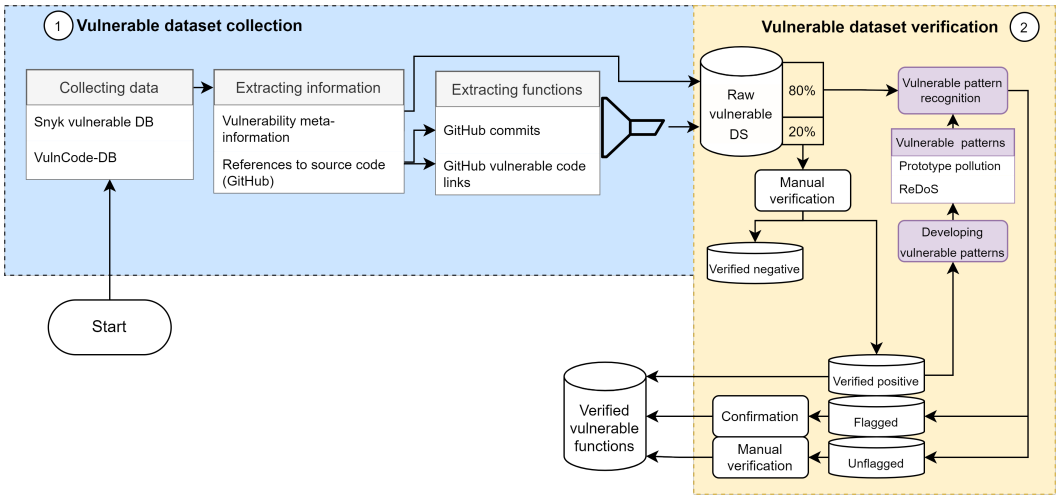


Fig. 1. Vulnerable dataset compilation approach.

functions and develop an approach to assist us in vulnerability verification. Finally, we perform a semi-automated filtering step to distinguish the truly vulnerable functions. Our approach is depicted in Figure 1.

3.1.1 Collecting possibly vulnerable functions. We use two sources for vulnerable functions collection: Snyk vulnerability database [103] and Google VulnCode-DB project [34].

Snyk database gathers information from other vulnerability databases (e.g., NVD, CVE), from threat intelligence systems, research and developer communities, and from scanning multiple platforms by the Snyk security team. Then each entry is manually verified and enriched by Snyk, and published on their website [103]. Snyk collects data for various programming languages. For JavaScript, they store vulnerabilities from NPM packages only, even though their sources for vulnerability collecting include other types of projects [82]. While having a large codebase (350,000 packages), the NPM repository is only one part of a massive JavaScript ecosystem, hence we reach out to an additional JavaScript vulnerability source that collects vulnerabilities from other environments. Like Snyk, VulnCode-DB also gathers data from CVE and NVD but additionally includes projects beyond NPM packages. VulnCode-DB also relies on the community to add relevant information for each vulnerability [34], as well as actual vulnerable code examples.

An entry from either of the vulnerability sources contains the following: vulnerability scores in several categories (difficulty, impact, scope); CVE (common vulnerabilities and exposures [78]) identifier; CWE (common weakness enumeration [77]) identifier; affected project name, version, and a short description; remediation actions; references; and other relevant information, e.g., a detailed description of the vulnerability, and proof of concept attacks.

We develop a JavaScript program to automatically collect available JavaScript vulnerability entries from Snyk and VulnCode-DB websites. We only collect entries that have at least one link under the “References” section, because later we extract source code from the provided links.

We accessed Snyk and VulnCode-DB websites to collect all applicable entries for our dataset on April 2021. The Snyk database had 2,975 entries under the category “NPM”, 2,810 of which had at least one code reference link. In the VulnCode-DB project, there were 3,680 entries, 201 of which were for JavaScript projects with at least one reference link.

3.1.2 Function extraction and preliminary filtering. To extract functions for our dataset, we first collect all the links that are provided for each entry of our vulnerability sources and isolate the links that lead to the source code. We also collect detailed metadata on each vulnerability. Out of the 3,011 collected entries with links, only four entries overlapped between Snyk and VulnCode-DB.

We then sort the reference links for each vulnerable entry into multiple categories. If an entry has multiple references, we rank them by priority, based on our assumption that some categories are more useful and convenient for our purposes than others (e.g., links directly leading to the vulnerable source code), and save the link with the highest rank. Other links in the same entry are discarded. The distribution of the remaining 2906 links by category is as follows: GitHub commits (1291, 44.43%), Advisories (474, 16.31%), GitHub issues (244, 8.40%), Pull requests (201, 6.92%), Vulnerable code (170, 5.85%), Bug reports (123, 4.23%), and other links (403, 13.87%).

For function extraction, we take two categories of links into consideration: GitHub commits and GitHub vulnerable code (50.3% of all links). Links of the first category point to the GitHub “diff” page, where all code changes in the given commit are displayed with the vulnerable files on the left and patched files on the right. Links of the second category point to a specific vulnerable file and include a range of lines that contain vulnerable code. The function extraction process from GitHub commits and GitHub vulnerable code is straightforward, while other categories of the links either do not contain the source code or have too many unrelated code parts that cannot be automatically filtered (e.g., pull request links). We extract functions from GitHub commits as follows. First, for each commit link, we use GitHub REST API to get vulnerable and patched versions of the committed files, along with positions of code lines that were modified (added, deleted, or changed). Secondly, we parse both versions of files with `espre` [31] and receive the range of each function (first and last symbol positions in the string). Lastly, we extract all functions with modified lines within their range. We also collect nested functions that include modified lines.

Links of GitHub’s “vulnerable code” type contain a modified line range at the end, in the format “#Li-j”, where *i* is the first affected line, and *j* – the last. We save the modified line range and retrieve code from the same link using a GitHub REST API. We then repeat the procedure of function extraction same as for commit links.

For each vulnerability entry, we place the affected files in the “files” property (see Figure 2). If the link category is “GitHub commit”, we include references to both vulnerable and fixed files (“link” and “fixedLink” properties) and a list of extracted functions in “vulnerable-fixed” pairs (“affectedFunctions” property). For “vulnerable code” links, we just add a vulnerable link to the “files” property, and vulnerable functions in the “affectedFunctions” property.

After collecting these possible vulnerability entries, we perform preliminary filtering to remove the following: (1) test files (links to such files contained “spec.js” or “test” keywords); (2) files that do not contain JavaScript functions; (3) empty functions (with no body); and (4) cases where the code snippets for both the vulnerable and fixed functions were identical.¹

After filtering, the number of functions from Snyk and VulnCode-DB “GitHub commit” links reduced to 4,288 (from 9,552) and 184 (from 538), respectively; “vulnerable code” functions remained unaffected (169). At this point, our dataset contains 4,870 functions (from 895 entries). We also group entries by vulnerability type. Note that we clustered 116 vulnerability types into 25 generalized groups—e.g., code injection group includes SQL, template, hash, and content injection; and procedure bypass group includes sandbox, signature, and authentication bypass. The most frequently reported vulnerabilities are XSS (228 entries), command injection (167 entries), ReDoS (121 entries), and prototype pollution (101 entries); see Table 7 in the appendix.

¹This happened when, even though the range of the function overlapped with the range of affected lines, modified parts belonged to a different function – such as where one function ends and another one begins on the same line.


```

{
  "link": "GitHub link with source code",
  "name": "category of the link (Commit, Pull Request, etc.)",
  "page": "vulnerability entry in Snyk / VulnCode-DB",
  "CVE": "CVE identifier",
  "CWE": "CWE identifier",
  "packageName": "affected package / project",
  "versions": "affected versions of packages / projects",
  "files": [
    {
      "link": "link to a file with vulnerable code",
      "fixedLink": "link to a file with fixed code",
      "affectedFunctions": [
        "a list of functions in pairs vulnerable-fixed"
      ]
    }, "..."
  ],
  "errors": "files that could not be processed",
  "details": "paragraph of detailed description of vulnerability",
  "vulnType": "category of a vulnerability"
}

```

Fig. 2. Structure of an entry in our vulnerable function dataset.

Following a similar approach, Ferenc et al. [30] produced a dataset of 1,456 vulnerable functions; however, at least 1,013 of these functions are in fact not vulnerable [79]. We identify two main reasons for this: (i) not considering cases with identical vulnerable and fixed functions—i.e., item (4) above; and (ii) not excluding the non-vulnerable functions from commits, where a given commit includes patched functions along with several unrelated/new functions. The second case is difficult to resolve automatically as it requires a clear distinction of vulnerable functions from the rest, which we address using a semi-automated verification step (see Sec. 3.3).

3.2 Manual verification of vulnerable functions

To perform further verification, we develop a web application that makes manual verification faster and easier. We upload our collected data to the web interface, which allows us to easily navigate between entries, files, and functions. For each function, the objective is to make a decision, on whether the function is vulnerable or not. For that, we examine the vulnerability description and the differences between vulnerable and fixed functions. If necessary, we also check Snyk and GitHub for additional information (sometimes other technical sources: e.g., blog articles, relevant forums) to make a concrete decision for each entry.

With help of our tool, we manually analyzed 150 vulnerability entries (~17% of the collected data) and found that, while some vulnerability types do not have any consistent patterns (each case is very specific to each project), others are often implemented in a similar way. Based on this observation, we implement a pattern-based detection approach for the functions that follow specific patterns. For these functions, the manual verification process involves less scrutiny than for other functions.

3.3 Semi-automated function verification

To improve the efficiency of the verification of vulnerabilities in our collected functions we develop an approach that uses vulnerable pattern search to detect functions of certain vulnerability types. For that, we utilize the Semgrep [97] static analysis tool, which performs an advanced semantic search in code based on provided patterns. Semgrep can understand variables and structures unlike a simple grep search. Thus Semgrep can also detect patterns in minified (and, in most cases, uglified)

| Semgrep rule | Function examples |
|--|---|
| <pre> rules: - id: prototype_pollution patterns: - pattern: \$SOME_OBJ[\$KEY] = ... - pattern-inside: function ...(...,\$KEYS,...) { ... } - pattern-inside: for (\$KEY in \$KEYS) ... - pattern-not-regex: ((__proto__[\\s\\S]*prototype[\\s\\S]* constructor))[\\s\\S]* - pattern-not-regex: Object\\.freeze\\([\\s\\S]* - pattern-not-regex: Object\\.create\\(null\\)[\\s\\S]* message: loop through keys in object severity: WARNING languages: [javascript] </pre> | <pre> function writeConfig(output, key, value, rec urse) { if (isObject(value) && !isArray(value)) { o = isObject(output[key]) ? output[key] : (output[key] = {}); for (k in value) { if (recurse && (recurse === true recurse[k])) { writeConfig(o, k, value[k]); } else { o[k] = value[k]; } } } else { output[key] = value; } } function recursiveMerge(base, extend) { if (!isPlainObject(base)) return extend; for (var key in extend) base[key] = (isPlainObject(base[key]) && isPlainObject(extend[key])) ? _recursiveMerge(base[key], extend[key]) : extend[key]; return base; } </pre> |

Fig. 3. Example of a Semgrep rule for prototype pollution.

code. Note that this does not apply to the obfuscated code, as the syntax and semantics may be distorted by obfuscation.

3.3.1 Semgrep rule development for selected vulnerability types. Semgrep search is performed with rule sets, provided in .yaml format.² Each rule defines which patterns the tool should look for, which to dismiss, and whether the target is inside of a specific structure or not. There are multiple open-source rule sets for many programming languages developed by Semgrep authors, as well as by the community. Most of them are written to find common bugs and inconsistencies in the code, but some rule sets also target vulnerabilities.

We performed a Semgrep search with the available community rules for JavaScript on our dataset, but unfortunately, it produced no matches. To understand the reason, we examined several patterns from those rule sets and reached several conclusions. Some rules, while covering a certain vulnerability type, come with patterns that are too specific. For example, a pattern requiring two code lines to be placed together would not be triggered if other code separates them. In other situations, our functions are simply not covered by the rules.

Therefore, we decided to develop our own rules for Semgrep pattern search. To write a Semgrep rule, we need to create a pattern for a line (or lines) of code that we want to match. Semgrep also allows to add conditions like: pattern-not, pattern-inside, pattern-not-inside etc. to make rules more targeted and avoid false positives. In the end, we add meta-information to each rule to describe the pattern. To develop rules for pattern recognition, we need to clearly understand each targeted vulnerability type and the protection measures against the vulnerability. The inclusion of patterns for preventive measures helps us avoid flagging fixed functions as vulnerable.

We create rule sets for several common vulnerability types, covering as many functions from each vulnerability type as possible. However, to rely on a pattern as an indicator of a vulnerability,

²<https://yaml.org/>

we need to consider the context where it appears. This proved challenging for certain vulnerability types, such as cross-site scripting and command injection, where these vulnerabilities can be mitigated by the surrounding context in numerous ways. As such, creating patterns for these vulnerability types is bound to generate many false positives, which we want to avoid. Therefore, we choose to create pattern rules only if the vulnerability has a very specific set of mitigating patterns. Prototype pollution (11.3% of all types) and ReDoS (13.5% of all types) vulnerability types primarily meet our selection patterns.

By understanding the vulnerability and its preventive measures described in Section 2.2.1, we can develop Semgrep rules for prototype pollution. The pattern for this vulnerability is the object key assignment statement, e.g., `object[key] = value`. The key, and possibly the value and the object have to come to the function from the outside (through arguments, global variables, or in another way). Thus, we add more rule properties to account for the context of the pattern, including mitigating factors. Finally, we created seven rules with different prototype pollution scenarios. The scenarios differ by the way the function argument gets into the property reassignment. Examples of such scenarios are iteration through arguments with `map`, `flatMap`, `forEach` methods, and `Object.keys` applied to one of the arguments; see Figure 3 for a rule example, along with two vulnerable functions, targeted by the rule. Additionally, all rules and the information on them are available on our GitHub page.³

To detect ReDoS vulnerability we utilize a tool that analyzes regular expressions and can identify them as “evil” or “benign”. For this purpose, we choose to use the NPM module `safe-regex` [46]. To extract regular expressions from the functions, we utilize the AST function representation, and collect nodes representing regular expressions, such as “`new RegExp(“...”)`” and “`/regex/`”. Then we perform `safe-regex` check on each regular expression, and save the functions that are flagged.

The presence of the evil regular expression in the code is already potentially dangerous, but we also perform a pattern search to check whether the evil regular expression is applied to a user-supplied string. As a result, we created two Semgrep rules that account for four different ReDoS scenarios. Figure 4 shows an example of a ReDoS rule and two different targeted functions.

3.4 Final vulnerable functions dataset

With two rule sets developed for prototype pollution and ReDoS vulnerability types, we executed a Semgrep pattern search on the remaining unverified functions from our vulnerable function dataset. We repeated the search several times, each time modifying the rule sets to match as many true positives as possible while reducing the rate of false positives. We summarize the results in Table 1; for each vulnerability type, the table includes the number of vulnerability entries, the number of all functions in this type, how many of them were flagged by our pattern search, and how many from all functions were then manually verified. Iteratively we adjusted our rule sets to not match any false positives in our dataset.

Although we cannot rely on the pattern search completely and all flagged functions still need to be manually confirmed, it does speed up the process of verification. Instead of performing an in-depth analysis of all the available information about a given vulnerability, we simply ensure that the flagged code is not an exception to the patterns.

During the manual verification, we observed that in some cases, the “fixed” file version that is supposed to eliminate the vulnerability contains only partial fixes. For example, for the prototype pollution vulnerability, the function performs a check for the value `__proto__`, but not for constructor or prototype that can still be used in a malicious way. An example of the ReDoS

³github.com/Marynk/JavaScript-vulnerability-detection/tree/main/semgrep

| Semgrep rule | Function examples |
|---|---|
| <pre> rules: - id: redos1 patterns: - pattern-either: pattern: <...\$ARG.\$METHOD(/.../,...)> - pattern: <...\$ARG.\$METHOD(new RegExp('...'),...)> .> - metavariable-regex: metavariable: \$METHOD regex: (match replace) - pattern-inside: function ...(..., \$ARG, ...) { ... } - pattern-not-inside: function ...(..., \$ARG, ...) { ... if(<...\$ARG.length...>) ... } message: direct use of match or replace severity: WARNING languages: [javascript] </pre> | <pre> function (name, value) { if (name == 'user_host') { return value.replace(/(\.[*?\\])+ /g, ''); } return value; } function isExternal(url) { let match = url.match(/^[^:\/?#]+:)?(?:\\\/\([^\/?#]*\))?(?:[^\?#]+)?(?:\?[^#]*)?(#.*)?/); if (typeof match[1] === 'string' && match[1].length > 0 && match[1] !== location.protocol) { return true; } return false; } </pre> |

Fig. 4. Example of a Semgrep rule for ReDoS.

| | Proto. Pollution | ReDoS |
|---------------------------------|------------------|-------|
| No. vulnerability entries | 101 | 121 |
| No. functions | 356 | 582 |
| No. functions flagged | 141 | 68 |
| No. functions manually verified | 166 | 164 |

Table 1. Summary on the vulnerable pattern findings in our vulnerable function dataset.

partial fix scenario is when one evil regular expression is removed, but the others remain present, or new evil regular expressions are introduced.

We had another unexpected but positive outcome from our manual validation. Since we ran a pattern search on the whole vulnerable function dataset, and not only on the targeted vulnerability types, the other functions were checked too, and matches were found. This means that a function that was reported to Snyk or VulnCode-DB under one vulnerability type also has another potential vulnerability present. In general, Semgrep search with our rule sets matched 195 prototype pollution patterns and 121 evil regular expressions (106 of them have ReDoS patterns) in other entries. This is an important finding, because while developers might fix a specific vulnerability after it is reported, other vulnerable code snippets may remain unchanged (discussed more in Section 7).

As a result, from the semi-automated verification step, we confirm 1,360 unique vulnerable functions of 43 different vulnerability type groups. This makes our dataset three times bigger than the dataset created by Mosolygó et al. [79] and to our knowledge the biggest dataset of verified vulnerable JavaScript functions.

4 METHODOLOGY: JAVASCRIPT VULNERABILITY DETECTION

In this section, we discuss the collection and preparation of the target dataset of real-world functions, as well as the implementation of our JavaScript vulnerability detection and verification framework. For an overview of the approach, see Figure 5.

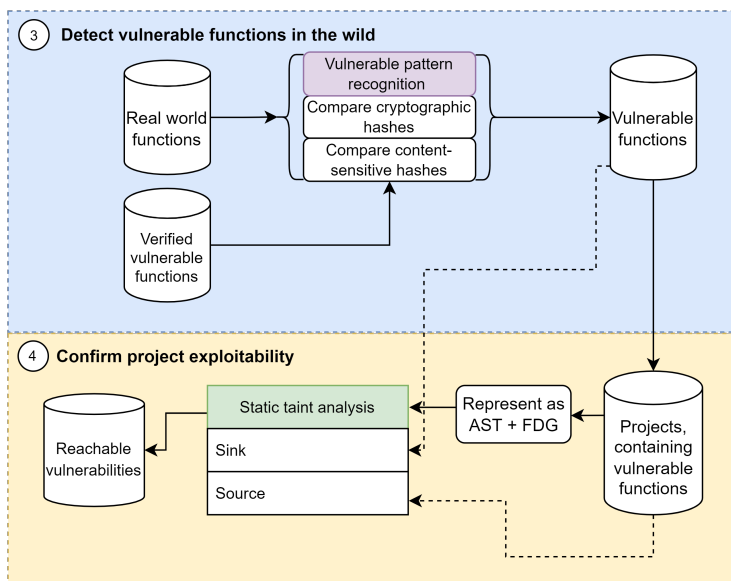


Fig. 5. Vulnerable function detection and exploitability verification approach.

4.1 Dataset of real-world functions

To collect functions from real-world projects we choose three sources: NPM packages, Chrome web extensions, and popular websites (as per the Cisco Umbrella Popularity List [20]). In the NPM open-source package registry, packages are mostly written in JavaScript. We extract JavaScript files from GitHub repositories of the 3,000 most popular NPM packages.⁴ These data sources were accessed between January and March of 2021.

The majority of scripts in Chrome web extensions are also written in JavaScript. To collect JavaScript files from extensions we download and unpack the source code for the 600 most popular extensions from Chrome Web Store. In 43 cases, our script was not able to retrieve the source code from Chrome Web Store API, hence only 557 extensions are processed further for our dataset.

From the Cisco Umbrella popularity list, we choose the 20,000 most popular websites. After sending an HTTP request to each website we receive a response with an HTML page. Then we either extract the JavaScript code from the `<script>` tag or save JavaScript files by following the links, defined in the same tag. In 18,108 cases the websites instead returned either a static HTML page with no JavaScript, a response in a different format, or an HTTP error. As a result, we collected JavaScript files for 1,892 websites.

From the crawled JavaScript files we extract all functions. The datasets contain a list of functions, mapped with the file link of the source URL. Finally, we filter all collected data to exclude test files (by searching for “spec.js”, “test” keywords in file links) and those functions that either had an empty body or only one statement such as printout to a command line or return. The resulting dataset may contain duplicate functions that belong to unique sources. A summary of the dataset is provided in Table 2 in the first 3 columns.

4.2 Vulnerable function detection

Our primary approach for detection is vulnerable pattern search. While this procedure can detect most of the vulnerable function variations, as long as the vulnerable pattern is present, it is limited

⁴We used a list of packages sorted by the frequency of their usage: <https://github.com/nice-registry/all-the-package-names>

to detecting only the patterns of the vulnerable types that we cover. To extend our detection range to the remaining vulnerabilities from our dataset, we also utilize textual similarity detection methods: content-sensitive fuzzy hashing and cryptographic hashing, which only target near-duplicate functions (used for completeness).

4.2.1 Vulnerable pattern search. For implementing vulnerable pattern search, we use the Semgrep static analysis tool with the rule sets we developed for prototype pollution and ReDoS vulnerability types; note that we also use these rule sets for a semi-automated verification of our vulnerable function dataset (see Section 3.3).

Our search logic is implemented using JavaScript, which iterates over the real-world functions, and for each of them executes the following steps: (1) Run Semgrep search with the rule set for prototype pollution, and flag the function if a match is found. (2) Find and extract regular expressions from the function. If successful, run the `safe-regex` module and flag the function if `safe-regex` finds any evil regular expressions. (3) Run Semgrep search with the rule sets for ReDoS patterns on functions with evil regular expressions; flag the function if the pattern is matched.

We save all flagged functions, including the projects/files where they are found. Note that for this approach, we do not need to refer to the functions of our vulnerability dataset until we want to find new patterns and develop additional rules.

Since our Semgrep rules are biased towards entries of our vulnerable functions dataset, it is reasonable to expect false positive matches among the real-world functions. To improve the precision of our approach, we make adjustments to the rules after analyzing around 100 flagged real-world functions (i.e., adding new conditions `pattern-not` and `pattern-not-inside`).

4.2.2 Textual-similarity based approaches. We use fuzzy hash and cryptographic hash comparison for vulnerable function detection.

Data abstraction. Before conducting the experiments with textual-similarity-based methods, we tokenize both vulnerable and real-world functions in order to bring functions to the same generalized format. We apply tokenization representation similar to the approach in [79]. After taking into consideration multiple syntactic parts of JavaScript code, and deciding on their importance and influence on the functionality, we decided to apply the following tokenization rules: (1) remove all space characters and comments; and (2) rename all variables and arguments to unified format (e.g., `varName1`, `varName2`). Note that we leave all punctuators, as assignments or arithmetical operators play a vital role in the meaning of the function. We also do not generalize the primitive variable values, such as strings, numbers, and regular expression patterns; note that we want to account for the variable values, since the difference between numbers 0 and 1 can be crucial for the vulnerability context (or, in the prototype pollution case, we want to check if the values of object keys are checked in the code).

To rename function variables, we parse each function to its AST and recursively process each node (i.e., tree leaf). In each round we look for “Variable declaration” and “Function declaration” nodes, save their position in the function (range), and then locate and replace characters between stored ranges. As for the arguments, we locate them by searching “Function declaration” nodes and looking for identifiers between the parentheses that follow. Then we perform the same renaming procedure as for variables. Figure 6 shows an example of a function and its tokenized version.

Content-sensitive hash comparison. Content-sensitive hashing (also known as fuzzy hashing), creates a fixed-size string that reflects the content of the input. It means that, unlike cryptographic hashing, small changes in the input result in small changes in the hash. It allows more flexibility in the function content, so if a few lines are added, removed, or modified, the content-sensitive hash will still be similar. It enables detecting slightly modified functions but may introduce false

| Function | Tokenized representation |
|---|--|
| <pre>function (n) { if (n === '.') return 1 return Buffer.byteLength(n)+2 }</pre> | <pre>function (varName1) { if (varName1 == = '.') return 1 return Buffer . byteLength (varName1) + 2 }</pre> |

Fig. 6. Example of a function and its tokenized version.

positives or false negatives. For example, this approach cannot distinguish the difference between vulnerable and patched functions, if the patch requires only small changes.

For fuzzy hash comparison, we use the SimHash [1] Python module. First, we create content-sensitive hashes for all functions of real-world and vulnerable datasets using the SimHash() method. Then we utilize the get_near_dups() method, which uses the hamming distance to compare the hashes. SimHash allows choosing the similarity threshold, by which it decides whether to flag hashes as near-duplicates. The value of the threshold can vary from 1 (strings are almost identical) to 64 (strings are completely different). To minimize the false positives rate we set the threshold to 1 so that it matches only highly similar functions. Finally, if SimHash finds a match, we save the function from the real-world dataset.

Cryptographic hash comparison. We also create SHA-1 hashes for all tokenized functions from both vulnerable and real-world functions datasets. We then perform a search for matches: for each hash of the tokenized function from the real-world project, we search for the same hash in the vulnerable function dataset. Finally, we save the real-world function and its source link, if a hash match is found.

4.3 Exploitability verification

In this section, we provide details of our design and implementation of a static taint analysis method to verify the potential exploitability of the flagged vulnerable functions (see Section 4.2).

Several existing tools offer taint analysis for JavaScript (see e.g., [32, 51, 97]); however as of writing, none of them support multi-file taint tracking—essential for real-world projects, which often consist of several JavaScript files. Our taint analyzer tracks an input datum, verifying that it has a path between a “source” and a “sink” across all the files in a given project. Our current prototype is tailored for vulnerable functions in NPM packages. Support for Chrome Web Extensions and websites can be added in future work, e.g., by leveraging the list of common sources for websites (such as URLs, Document.location, location.search, client-side storage) as defined by Kang et al. [56]. Note that capturing all input sources in extensions and websites comprehensively is non-trivial as there are numerous such options.

Note that we assume that the package is vulnerable if the vulnerable functionality of the package is reachable externally. The vulnerability can be exploited, only if a developer combines obtaining and processing a user’s input with a vulnerable functionality in the NPM package. Here the term “user” refers to a person who interacts with the developer’s project and enters some data in the context of this interaction. Despite the specific prerequisites for the exploitation, the vulnerabilities in the NPM packages are perceived seriously and often get the highest severity ratings (e.g., both CVE-2022-21189 and CVE-2021-23518 rated as critical with 9.8/10 scores).

NPM packages are meant to be included in projects as external dependencies. As such, they contain distinctly marked “exports”, which are the only objects accessible to the user. These objects are potential “sources” of user-controlled inputs. An NPM package consists of one or more JavaScript files that interact with each other by exporting and importing functionality, and a “package.json”

file that contains package metadata. Among other information, this file contains the top-level export location (or an “index file”) that is made accessible to the developer when this package is included in the project. However, there are multiple competing standards and formats to specify export information, and any NPM package owner can choose to implement any number of them. For an example of a typical implementation of “package.json” that contains three different export declaration standards, see the “Underscore” NPM package [5].

To address the absence of a single generalized data format in package.json, we create an automated approach that allows identifying which functionality may be accessible to the user without relying on unstructured metadata. Instead, we utilize the code structure to locate exported functionality that is not accessed within the project and therefore is likely intended for external use. Note that sometimes the package owners choose not to include some top-level exports in their “index file”, resulting in the unreachable module of functionality. This may be due to the module being shelved, under ongoing development, etc. We argue that if a vulnerability is found in such a module, it may become potentially exploitable in the future, if the package owners decide to change the module’s accessibility, or the developers use the source code of the package rather than downloading it from the NPM repository.

File Dependency Graph. To identify the package modules intended for external usage, and to enable a multi-file static taint tracking, we introduce our File Dependency Graph (FDG) structure. Firstly, a multi-file project is represented in AST, preserving the file structure. From each file’s AST we locate nodes that indicate an “import” or an “export” of certain functionality. For this, our algorithm uses a predefined set of rules, specific to each type of the import and export nodes (there are six import types, such as “ImportSpecifierNames” and “requireNames” and five export types, such as “ExportAllDeclaration” and “ModuleExportNames”). Each rule provides instructions to extract relevant information: the name of the imported/exported functionality, the origin file, and the node itself.

Secondly, after processing the nodes with the rules we iterate through the selected nodes to find and group the matching pairs: the functionality exported by one file, which is then imported by another file. Note that for each “import” node, there can be only one matching “export” node, whereas for the export nodes we include the array of all corresponding “imports”. If there is an export node, for which no matching import is found, the node is marked as top-level. Since the node is a function, we consider the arguments passed to it a taint “source”. This way we obtain a graph that indicates the relationship between project files, where each node has the information about its matching pair.

The taint approach implementation. The process to verify the exploitability of a vulnerable pattern in an NPM package consists of six steps (for the corresponding pseudo-code, Algorithm 1):

- (1) Download the project that contains the detected vulnerable function;
- (2) Create an AST representation of the project;
- (3) Create an FDG of the project based on its AST and find a “source”;
- (4) Get the location of the vulnerable “sink” from the Semgrep-based approach results, find the corresponding node in the AST;
- (5) Build a “taint graph” that goes over the AST with FDG and taints nodes, which are affected by input data; and
- (6) Look for a path in the “taint graph” that includes a “sink” variable. If there is such a path, then the user-controlled input can reach the vulnerable code.

Our algorithm starts by taking the resulting dataset of the functions marked as vulnerable by the previous step of our framework (see Section 4.2). Each vulnerable entry of the dataset has a URL link to a file within a project on GitHub. Using this link we get the information on the name

Algorithm 1 Algorithm of taint analysis tool.

```
1: downloadProject(projectUrl, projectPath + “/original”)
2: astgen.generate(projectPath + “/original”, projectPath + “/ast”)           ▶ create an AST
   representation of the project
3: asts ← []
4: for all file in (projectPath + “/ast”) do                               ▶ create an array of generated ASTs
5:   asts.push(parse(file))
6: end for                                                                 ▶ start: create FDG
7: allNodes ← []
8: importNodes ← []
9: exportNodes ← []
10: for all ast in asts do
11:   for all node in ast do
12:     allNodes.push(node)
13:     node.file = ast.file                                               ▶ add file information to each node
14:     if isImport(node) then
15:       node.isImport ← true                                             ▶ mark a node as an “import” node
16:       node.from ← resolveImportLocation(node)                       ▶ add origin to the import node
17:       importNodes.push(node)
18:     end if
19:     if isExport(node) then
20:       node.isExport ← true                                             ▶ mark node as an “export” node
21:       exportNodes.push(node)
22:     end if
23:   end for
24: end for
25: sources ← []
26: for all export in exportNodes do
27:   export.importedBy ← []
28:   for all import in importNodes do                                     ▶ using import origin connect export with import
29:     if import.from = export.file AND import.name = export.name then
30:       export.importedBy.push(import)
31:       import.importedFrom ← export
32:     end if
33:   end for
34:   if export.importedBy is empty then                                   ▶ if found export but no import, mark as “source”
35:     sources.push(export)
36:   end if
37: end for                                                                 ▶ end: create FDG
38: taints ← []
39: currentLeaves ← sources                                               ▶ define found sources as an initial “taintor”
40: while currentLeaves is not empty do                                   ▶ do while new values are tainted each iteration
41:   currentLeaves ← propagateTaints(currentLeaves, taints, allNodes)   ▶ this step is
   described in Algorithm 2
42:   taints ← taints + currentLeaves                                       ▶ add tainted nodes to an array
43: end while
```

and version of the project. We craft a new link with extracted information and clone the project repository in our temporary folder (using the `git clone` command).

Out of several JavaScript parsers that generate ASTs [2, 31, 40, 73], we found ASTgen [52] to be the most applicable. ASTgen is a multi-language parser that uses Babel [73]. The Babel parser efficiently deals with any JavaScript variations, bringing them to one general format when parsed. This parser is constantly adapting to the updates in ECMAScript standards, implementing changes even from *stage 0* proposals [6]. In comparison, Acorn [2] and esprima [40] only support basic JavaScript syntax that is already in the standard. ASTgen provides additional convenience by preserving the file structure of the project when parsing scripts to AST.

As a third step, we create an FDG of a project. Then we create one consolidated AST for the whole project and merge it with the nodes from our FDG (we replace the “import” and “export” nodes in the joined AST with the corresponding FDG nodes). For the unaffected nodes, we simply include a link to their origin file.

Next, we start an iterative process of data tainting. In the first iteration, we find all nodes in the AST that are tainted by defined “sources” and add them to the accumulator. The newly tainted nodes are then marked as new “sources” for the next iteration. This process is repeated until no new nodes are tainted. Additionally, each tainted node in the accumulator has information about its “taintor” (a node that caused tainting of the current node).

During the node tainting process, the algorithm first finds all instances of already tainted “source” identifiers. Since every node type in the AST has a different influence on the related nodes, a set of predefined taint rules is required, which regulates taint propagation based on the node type. For example, when a function is called with a tainted argument (e.g., `myFunc(tainted_arg)`), the corresponding argument is tainted in the function declaration (e.g., `function myFunc (tainted_arg, arg2) return tainted_arg;`). Similarly, if a tainted value is assigned to a variable, the variable becomes tainted as well. The pseudo-code of a single iteration is displayed in Algorithm 2. In total, we have 29 rules of tainted variable propagation for different syntactic constructs (e.g., variable declaration or reassignment, value assignment to the array and object entries, function calls, value returns, etc.). At the moment, the rules do not include handling for JavaScript classes and the `try/catch` construct, as well as the `eval()` construct. All rules are listed in our GitHub repository.⁵

Finally, we obtain the resulting accumulator that contains the list of the tainted nodes. We locate the “sink” node in the AST based on the information retrieved from Semgrep or a textual similarity algorithm. Then, to determine if the vulnerability is exploitable, we check if this node is present in the accumulator. If it is found, we can further obtain a full taint path by traversing the “taintor” properties. Later we use the taint path to provide the following information on the exploitability: where the vulnerability is introduced, which input is potentially controlled by the attacker, and where the functionality is exported.

5 RESULTS

In this section, we first present the results of running our vulnerable function detection approaches. Then we manually validate a randomly selected set of findings to provide an estimated precision measurement of our detection framework. Next, we provide the results of identifying the exploitable vulnerabilities in NPM packages (including our findings from manual validation). Furthermore, we describe our semi-automated vulnerability reporting process and discuss the overall response dynamics. Finally, we discuss several case studies for high-profile projects and ethical considerations for our responsible disclosure.

⁵github.com/Marynk/JavaScript-vulnerability-detection/blob/main/taintAnalysis/taintGraph.js

Algorithm 2 Algorithm of one iteration of taint propagation.

```
1: currentLeaves ← arguments[0] ▷ “taintors”
2: allTaints ← arguments[1] ▷ array with previously tainted nodes
3: allNodes ← arguments[2] ▷ all AST
4: nodesToPropagate ← currentLeaves
5: for all node in currentLeaves do
6:   sameIdNodes ← findSameId(allNodes, node) ▷ nodes with the same name
7:   for all sameIdNode in sameIdNodes do
8:     sameIdNode.taintedBy.push(node) ▷ add the information about the “taintor”
9:     if sameIdNode not in allTaints then
10:      allTaints.push(sameIdNode)
11:      nodesToPropagate.push(sameIdNode)
12:     end if
13:   end for
14: end for
15: newTaints ← []
16: for all node in nodesToPropagate do
17:   taints ← propagateTaint(node) ▷ with propagation rules for different nodes
18:   for all taint in taints do
19:     taint.taintedBy ← node ▷ add the information about “taintor”
20:     if taint not in allTaints and taint not in newTaints then
21:       newTaints.push(taint)
22:     end if
23:   end for
24: end for return newTaints
```

5.1 Vulnerable function detection results

We perform our textual-similarity tests and pattern recognition search on our real-world dataset of 9,205,654 functions. First, we run the Semgrep-based approach, which flagged 284,413 potential prototype pollution and 22,496 potential ReDoS vulnerabilities. The distribution of the detected vulnerabilities among all tested projects is presented in Table 2.

Note that sometimes a vulnerable pattern search matched multiple functions from the same file to one vulnerable finding. This happened because of nested functions that have the same vulnerable code. We create a script to find all nested functions and keep only the child function. As a result, we get 117,601 unique detected functions for prototype pollution and 7,333 functions for ReDoS.

The SimHash algorithm matched 1,320 functions from the whole real-world dataset (Table 3). The vulnerability types of the detected functions are distributed as follows: cross-site scripting (307), ReDoS (306), prototype pollution (138), command injection (133), directory traversal (72), SQL injection (68), denial of service (59), and others (such as arbitrary script injection (16), directory traversal (12)). Similar to the vulnerable pattern search algorithm, SimHash also detected several nested functions. We apply the same filtering script to all detected functions and as a result, we get 965 unique vulnerabilities.

Finally, the cryptographic hash matching algorithm produced 131 matching cases from all real-world functions (Table 3). All of these matches were already detected by SimHash. However, since the findings of cryptographic hash matching are guaranteed to be identical (except for variable names) copies of vulnerable functions, there is no need to manually verify them, and we can

| | # sources | # functions | # prototype pollution | | # ReDoS | |
|-------------------|-----------|-------------|-----------------------|---------|---------|--------|
| | | | total | unique | total | unique |
| NPM packages | 3,000 | 413,774 | 10,338 | 4,896 | 1,325 | 493 |
| Websites | 1,892 | 5,739,271 | 179,626 | 77,504 | 14,586 | 4,648 |
| Chrome extensions | 557 | 2,659,649 | 94,449 | 35,201 | 6,585 | 2,192 |
| Total | 5,449 | 9,205,654 | 284,413 | 117,601 | 22,496 | 7,333 |

Table 2. Results of vulnerable function detection in real-world projects by the Semgrep-based approach.

| | NPM packages | Extensions | Websites | Total | Unique |
|-------------|--------------|------------|----------|-------|--------|
| Fuzzy Hash | 56 | 201 | 1,063 | 1,320 | 965 |
| Crypto Hash | 30 | 85 | 16 | 131 | 131 |

Table 3. Detected vulnerable functions by textual similarity methods (all vulnerability types).

automatically count them as true positives. The findings belonged to the following vulnerability types: ReDoS (29), cross-site scripting (26), command injection (25), prototype pollution (14), timing attack (5), denial of service (1), and others (e.g., arbitrary script injection (1), directory traversal (1)).

Note that all ReDoS and prototype pollution vulnerabilities detected by textual similarity methods were also detected by our Semgrep rules. This is expected, as the rules are based on functions from our vulnerability dataset, and SimHash and cryptographic hash are detecting near-duplicate versions of those functions.

As a result, our experiment identifies 125,555 unique functions, detected by at least one method in our framework. 124,934 findings belong to either prototype pollution or ReDoS vulnerabilities (i.e., the total number of unique functions detected by Semgrep), and the remaining 621 to other types of vulnerabilities that appeared in our vulnerable functions dataset.

5.2 Manual validation of the results

To evaluate the performance of the vulnerable pattern search, we randomly picked 100 functions for prototype pollution and 100 for ReDoS vulnerability from all the detected functions. For each finding, we examine three main features. Firstly, we confirm that the flagged pattern is detected correctly. Secondly, we verify that the input to the pattern comes from the outside of the function. Lastly, we make sure that there are no sufficient protection measures, missed by our Semgrep rules. As a result, we identified 8 false positives for prototype pollution and 3 for ReDoS, resulting in the precision rate of 92% and 97%, respectively.

Among the prototype pollution functions, patterns got false matches for mainly two reasons. Firstly, in JavaScript the following assignment syntax: `obj[key] = value` is valid for an object and an array (ordered collection of elements). In the case of the assignment to the array element, the “key” is an index of the element. However, the prototype pollution applies only to the object’s properties modification. Since JavaScript does not have explicit data types, the patterns fail to distinguish array and object assignments. The second reason for false positives for prototype pollution was the use of a function obfuscation technique, which did not modify the vulnerable pattern but heavily obfuscated the protection measures that prevent the vulnerability.

In three false positive matches from ReDoS, the function contained multiple regular expressions. The input variable was matched with some of the regular expressions, but not with the ones, marked as dangerous in the previous steps of our experiment. Since the dangerous regular expression was not used with the user input, we do not see a possibility of a ReDoS attack in this scenario.

Since the fuzzy hashing comparison approach may introduce false positives (as described in the previous section), we need to manually verify the findings as well. Therefore, for 90 out of 965 matches we checked the differences between the real-world function and the vulnerable function

with similar content. As a result, we identified two false positive matches, which gives us a precision rate of 98%. There were two reasons for false positives: one function was patched, but the fix was only in one line, so the function remained similar enough; the other flagged function matched a ReDoS vulnerable function, but did not have a regular expression in its body (otherwise it was identical to the vulnerable function).

5.3 Exploitable project identification results

Next, we perform our static taint analysis on the findings from Section 5.1. We aim to identify the functions that make corresponding projects exploitable by an attacker. We run this experiment on 5,389 findings in 204 NPM packages flagged as vulnerable. Our tool identified taint paths for 259 prototype pollution findings and 52 ReDoS functions (from 134 NPM packages), reducing the number of flagged functions by 94%. After manual verification of 100 randomly-chosen findings in this experiment, we detected no false positives, meaning that all the identified paths correctly indicate a relation between the vulnerable “sink” and top-level export. We list top-20 NPM packages (as of April 2022) containing such potentially exploitable findings in Table 4. The second column in the table shows the number of NPM packages that depend on the vulnerable one. The third column refers to the weekly downloads number of vulnerable packages. In the fourth column, we present the number of findings by Semgrep, and in the last column, we show the number of vulnerable findings that are reachable by the top-level exported functions. Note, that for the “Lodash” package we do not include values in the second and third column, as our findings belong to the master branch of the package, which is different from a branch that is located in the NPM package registry (for detailed explanation see Section 6 about Lodash).

During a detailed analysis of vulnerabilities, we came across two scenarios, which are not considered in our approach. According to principles of best practices in coding, we assume that the protection measures for vulnerable functionality are implemented within one function to achieve better granularity and code reusability. However, in one scenario the protection measures were taken out in a separate function, which was then called directly before the vulnerable pattern, therefore negating the vulnerability. In the second scenario, the protection measures were implemented outside of the function, on an earlier step of the tainted variable propagation. These cases are current limitations of our approach, which can be solved in the future, particularly by searching for protection measures within the whole project and establishing their connection (or lack thereof) with the vulnerable pattern using taint paths.

5.4 Vulnerability reporting

From taint analysis, we obtain vulnerability entries with the following details: the taint path, the detected vulnerability type, the “sink” variable, the link to the targeted project, and the code itself.

Before reporting the findings, we need to make sure that they were not disclosed earlier in the MITRE CVE database [78]. We develop a script that automatically compares project vendors and vulnerability types from our findings with existing entries in the CVE database. If a match is found, the CVE ID is added to the finding, and afterward, we check manually if the existing vulnerability is the same as our finding. The script matched 11 findings as duplicates. We noticed that the naming of vendors and products is not regulated, and thus project vendor names in the CVE database may not match the actual names used in GitHub repositories. Therefore, we perform an additional search for vulnerability duplicates manually, and as a result, find additional 10 matches. At this point, our exploitable dataset contains 290 unique vulnerability reports.

Next, we use the information in our unique findings objects to automatically create reports in a readable format. Firstly, we extract three nodes from each taint path: the exposure of the vulnerability to the user (i.e., top-level function export), the point of entry for a user-controlled

| NPM package | # depended projects | # weekly downloads | Identified vulnerable functions | Identified paths |
|-----------------|---------------------|--------------------|---------------------------------|------------------|
| Lodash (master) | unknown | unknown | 86 | 56 |
| Request | 35,681 | 21,155,428 | 12 | 6 |
| Async | 22,704 | 46,511,068 | 72 | 4 |
| Minimist | 11,310 | 40,869,389 | 6 | 4 |
| Body-parser | 11,207 | 26,897,728 | 2 | 2 |
| Ramda | 5,016 | 8,958,095 | 48 | 20 |
| Morgan | 3,897 | 3,502,973 | 2 | 1 |
| Qs | 3,790 | 56,395,909 | 9 | 2 |
| Loader-utils | 3,777 | 41,352,810 | 2 | 1 |
| Validator.js | 2,445 | 5,671,275 | 2 | 2 |
| Grunt | 2,059 | 601,235 | 16 | 4 |
| Js.merge | 1,411 | 1,905,969 | 3 | 1 |
| Js-beautify | 1,313 | 2,246,662 | 6 | 1 |
| Nodemon | 1,277 | 4,895,728 | 1 | 1 |
| Vinyl | 1,234 | 4,077,423 | 1 | 1 |
| Html-minifier | 1,203 | 3,207,164 | 134 | 47 |
| Clean-css | 1,164 | 11,415,710 | 22 | 11 |
| Reselect | 1,089 | 3,960,959 | 1 | 1 |
| Url-parse | 858 | 9,114,628 | 1 | 1 |
| Restify | 754 | 188,600 | 9 | 4 |

Table 4. Top-20 NPM packages and vulnerable findings as of April 3, 2022.

variable, and the vulnerable sink. We then use the GitHub link to the corresponding file, and for each node add a specific line reference to the link, so that it points directly to the location of the taint node in question. We also use GitHub APIs to extract information on the project name and version. Finally, we write a readable template for the reports, and then we automatically insert all of the extracted data into placeholders to generated easy-to-understand reports.

For the responsible disclosure of vulnerabilities, we choose to send emails to the developers responsible for NPM projects. Using GitHub APIs and file links for each finding, we query the repository owner’s information and collect their email address. If the email of the owner is unavailable, we query the repository information and find the list of its contributors. Then we iterate through the list until one of the contributors has a publicly available email address. Then we map collected emails to the findings, and group readable reports by the same project, and then by the same email (sometimes one person is responsible for multiple repositories). Lastly, we use the “nodemailer” NPM package [90] to create and send emails to the responsible developers.

Using this vulnerability reporting algorithm we notified 112 responsible developers about 290 exploitable vulnerabilities in their NPM packages.

5.5 Ethical considerations during vulnerability disclosure

We performed our experiments in two batches: between May 2021 and October 2021, and between February 2022 and March 2022. After detecting and manually verifying vulnerabilities (during the first period of our experiments), we responsibly disclosed the findings to the owners by contacting them privately via email. In the statement, we included a detailed description of the finding, the corresponding proof-of-concept attack (in selected cases), and our suggested fix. We only received a response from 8/20 contacted owners, all of them are the developers of NPM packages.

Async's owner replied after 10 days, acknowledging the vulnerability, and notifying us about the fix applied in the next version. **Minimist**'s owner replied on the second day after the disclosure, promising to look into it soon. However, after more than 4 months the issue was not addressed. After the CVE ID was obtained, we contacted the owner again, notifying them about publicly disclosing the vulnerability as a GitHub issue. Since it was over 120 days since our first contact, we followed the best practices of ethical disclosure. After the issue was created, the developers of the package deployed a fix in two days.

QueryString cooperates with TideLift, a company that acts as an intermediary in the disclosure process. After several weeks of discussion between TideLift representatives and the package owner, they claimed that the developer who uses the affected functionality of the QS package must provide protection measures themselves. We provided further explanations, however, received no response. After 4 months of the last communication, the CVE ID on this issue was received with a 7.5/10 severity rating, and we contacted them again with a notification about publicly disclosing the vulnerability. We received no response. After some time we discovered that the owners sent a request to MITRE to withdraw the vulnerability. Their request was successful, and the CVE ID was revoked on April 11, 2022. **Ramda**'s owners asked us to disclose security vulnerabilities in GitHub issues and, after our statement, started a discussion with fellow developers as well as with an external security company, which advise the developers on related issues. After the CVE ID with severity rating 9.1/10 was issued, Ramda owners tried to withdraw it from MITRE, but did not succeed. The vulnerability is not addressed as of July 27, 2022.

SailsJS owners argued that the vulnerable functionality is not intended to be accessible to the attacker ("This code runs when Sails starts up, not on a per request basis"). We provided them with a proof-of-concept attack video but received no response. Then we obtained a CVE ID for this issue and made it public. Instead of fixing the vulnerability, the owner then provided a way for the developers to work around this problem. The **AngularJS** team decided not to fix the findings ("the reported problems don't pass the threshold of a dangerous vulnerability that would require changes to AngularJS"). The requested CVE IDs are being processed by MITRE as of writing.

Finally, **SheetJS** did not acknowledge the disclosed vulnerability as significant, as it does not belong to their main functionality, and **Highland**'s owner mentioned that this package was no longer maintained, despite being downloaded by over 44 thousand users weekly.

After the second period of our experiments, we used our semi-automated vulnerability reporting process. As of writing, we received 22 feedback emails. 13 project owners claimed that either there is no vulnerability, or the vulnerability is insignificant and will not cause any damage. Three owners requested additional information on the vulnerabilities. Another three acknowledged the vulnerabilities and have started to address them. One project owner notified us, that the project contains the preventive measures for the vulnerability. The applied protection measures are specific to the project functionality and hence were not detected by our approach. Finally, two of the contributors provided the wrong contact information, so the emails bounced back.

6 CASE STUDIES

After analyzing manually the results of our vulnerability detection and verification framework we conduct several case studies on some of them to understand how these vulnerabilities affect the projects containing them, and what threats they pose to users. We choose the projects from the top, middle, and bottom of the popularity list to understand the vulnerabilities in different project types. Additionally, we select two JavaScript frameworks (both client-side and server-side) to explore possible implications of the found vulnerabilities. All rankings, as well as the usage numbers in this section, are dated by April 2022.

We isolate 20 findings from 17 different projects to describe their impact in detail in this section. 11 projects are from the NPM registry, while the remaining six are from popular websites. From the NPM projects, we select nine packages that were flagged by the taint analysis algorithm; additionally, we select two NPM projects that are JavaScript frameworks (AngularJS, SailsJS), rather than packages (libraries). Note that the vulnerable functions found in the frameworks cannot be detected by taint analysis due to the differences in the usage of such projects by the developers. More precisely, while the functionality from NPM libraries is just imported by the developer in their project, frameworks work as engines, which the developer uses to execute their project. Therefore, the top-level export concept that we utilize to define “sources” of user-controlled data is not applicable in the framework scenario.

By vulnerability types, we targeted 12 prototype pollution findings, six ReDoS findings, and one finding with both vulnerability types (in the SailsJS project). Note that after searching automatically through the MITRE CVE database [78] and manually checking the information available on the web on specific projects and the functions in question, we were unable to locate any reports on these vulnerabilities in publicly accessible sources.

Minimist [36]. This NPM package (40.8 million weekly downloads) is used to parse command-line arguments in Node.js applications. Minimist was previously reported for prototype pollution⁶ and the authors issued an updated version with a fix. However, our framework flagged this code piece (with the fix) as vulnerable in the latest version of the package. After the manual analysis, we discovered that the fix does not cover all malicious scenarios. Therefore, we were able to implement PoC attacks that violate both integrity and availability of a JavaScript application.⁷ This vulnerability received a critical (rated 9.8/10) CVE entry: nvd.nist.gov/vuln/detail/CVE-2021-44906. It was fixed by the authors 4 months after the initial disclosure.

SailsJS [7]. This framework is a model-view-controller web application framework written in JavaScript. Currently, there are 8,415 active websites built on this framework [12]. At least 24 of these websites appear on the top 1-million Tranco list [87]. Our findings indicate that the `loadActionModules()` method is vulnerable to both ReDoS and prototype pollution, due to the absence of sanitization of the strings extracted from filenames. There is a conceivable scenario, where filenames are controlled by the end-user (e.g., dynamic creation of API endpoints). In this case, the method can be exploited in a form of a prototype pollution attack that leads to denial of service.⁸ This vulnerability received a critical (rated 9.8/10) CVE entry: nvd.nist.gov/vuln/detail/CVE-2021-44908. As of writing, it has not been addressed. Additionally, certain filenames may also cause availability issues due to the usage of an “evil” regular expression in the method.

Ramda [96]. This NPM package (8.9 million weekly downloads) provides utility functions with a focus on functional programming style. Our findings indicate that the method `mapObjIndexed()` is vulnerable to object property injection vulnerability. Due to insufficient protection measures, it is possible to pollute `Function.prototype` by supplying a crafted object, causing threats to the integrity and/or availability of the JavaScript application.⁹ This vulnerability received a critical (rated 9.1/10) CVE entry: nvd.nist.gov/vuln/detail/CVE-2021-42581. However, the owner did not acknowledge the threat, and thus no fix is available.

Async. Async [13] is an NPM package (46.5 million weekly downloads) that provides functionality for working with asynchronous JavaScript. Similar to Ramda, the method `mapValues()` in `async`

⁶Prototype pollution in Minimist, <https://security.snyk.io/vuln/SNYK-JS-MINIMIST-559764>

⁷The PoC Minimist attack is available at github.com/Marynk/JavaScript-vulnerability-detection/blob/main/minimist%20PoC.zip

⁸The PoC Sails.js attack is available at github.com/Marynk/JavaScript-vulnerability-detection/blob/main/sailsJS%20PoC.zip

⁹The PoC Ramda attack is available at jsfiddle.net/3pomzw5g/2/

package is vulnerable to object property injection.¹⁰ This vulnerability received a high (rated 7.8/10) CVE entry: nvd.nist.gov/vuln/detail/CVE-2021-43138. It was fixed by the owner 10 days after disclosure in package version 3.2.2.

Lodash [54]. This is a library for JavaScript with various utility functions, mainly for array and object manipulations. It is used in approximately 3.6% of websites as of July 2021 and has 38.9 million weekly downloads [53]. Our findings indicate that the master branch of the lodash GitHub repository contains code exposed to a prototype pollution attack. Vulnerabilities of this type have been reported and subsequently fixed in multiple lodash functions [105–108]. However, lodash does not apply security fixes to the source code in their master branch. While in the distributables that are supplied to the NPM registry, the found vulnerability is fixed, if developers decide to clone source code from lodash and use it locally, they can clone the master branch. Hence, the discovered vulnerability still bears a significant impact on some applications. More specifically, the internal function `baseAssignValue()` offers no protection measures against modifying prototype or constructor properties, which potentially exposes any project that uses specific methods from cloned lodash source code, such as `_.set()`, `_.copyObject`, `_.keyBy`, `_.countBy`, `_.groupBy`, to a prototype pollution attack. We implemented a proof-of-concept attack,¹¹ which targets `_.set()` method and successfully injects a custom property to `Object.prototype`, thus adding this property for all objects of the running application.

Additionally, among detected vulnerable functions we discovered that two of them are from projects that contain local copies of the lodash library: the Highland NPM package [14] and `accompany.com` web domain; these projects use lodash functionalities, but do not automatically receive security updates for it. As a result, both projects are exposed to attacks that can exploit both newly detected and previously found vulnerabilities that were reported and fixed in the original lodash. In particular, the Highland package and `accompany.com` website contain an older version of the `baseAssignValue()` function from our findings, pre-dating even the partial fix.

AngularJS [33]. This is a client-side JavaScript framework for developing web applications. This version of Angular has been discontinued since the end of 2021, however, there are still more than 1 million live websites written with this framework [11]. We have identified three prototype pollution vulnerabilities in the source code of AngularJS:

- (1) The AngularJS routing system, implemented by the `$routeProvider` service, contains vulnerable code that exposes certain AngularJS-based applications to a prototype pollution attack. It is possible in the scenario when the paths in such applications are created dynamically based on user inputs. As a consequence, the attacker can disrupt the routing of the application (e.g., navigating between UI components). We implemented a proof-of-concept attack for this.¹² If the AngularJS-based application allows the user to supply a custom payload to the `$routeProvider.when()` method, an attacker can manipulate the prototype of the routes object by providing `__proto__` as the route path.
- (2) A misuse of AngularJS's `Select` directive can lead to prototype pollution. If the developer of an AngularJS-based application allows the user to dynamically add/modify options associated with `Select`, a special select value can be crafted to perform a prototype pollution attack.
- (3) Programmatic navigation within an AngularJS-based application can be performed via the `$location` service. Particularly, query parameters can be added to the current URL with the `$location.search()` function. This functionality can be abused to perform a prototype pollution attack if the query parameters in question are dependent on user inputs.

¹⁰The PoC async attack is available at jsfiddle.net/oz5twjd9/

¹¹The PoC lodash attack is available at jsfiddle.net/evmjxaq1/

¹²The PoC AngularJS attack is available at jsfiddle.net/mspc3f8n/

QueryString [38]. This is an NPM package (56.4 million weekly downloads) that provides parsing and *stringifying* (i.e., transforming any data type into a string literal) functionality of query strings. We discover a prototype poisoning vulnerability in the internal `merge()` method. Due to insufficient input sanitation, attacks on the integrity and availability can be implemented by manipulating the input to `Qs.parse()`.¹³

Grunt-usemin [125]. This is an NPM package that creates a minified version of web files (HTML, CSS, JavaScript) in a project, and it has 29,673 weekly downloads. It can also automatically replace links to scripts in code with the minified versions of the same scripts. We found a potential ReDoS vulnerability in the function that searches for the internal file references in the project and replaces them. In terms of this functionality, the `getBlocks()` function aims to parse an HTML file line by line, and extracts specific information. During this process, a dangerous regular expression is applied to match patterns in every line. A maliciously crafted HTML file can cause ReDoS when processed by Grunt-usemin.

JSON.parse polyfill [80]. Four web domains were flagged due to the use of polyfill for a built-in JavaScript method `JSON.parse`. A polyfill is a piece of code (usually JavaScript) that provides modern functionality in older browsers that do not natively support it. The flagged domains and their rank in Cisco Umbrella [20] are: `acdc-direct.office.com` (#259), `ad.crowdctrl.net` (#5806), `activedirectory.windowsazure.com` (#6050) and `360.cn` (#6291). The implementation of `JSON.parse` polyfill by these domains uses an unsafe regex to sanitize the input, which can lead to a ReDoS attack. Furthermore, it can be exploited to run arbitrary code due to the use of `eval()` on the parsed text as part of the functionality implementation (command injection). This attack can be exploited using older versions of Firefox (v.2-3), Opera (v.10.1), Safari (v.3.1-3.2), and Internet Explorer (v.6-7).

Gravatar.com [35]. This is a service for creating *Gravatars* (globally recognized avatars), which are integrated into over a million websites as of July 2021. On `gravatar.com`, a WordPress module called `cookie-banner` processes cookies with a dangerous regular expression, and if the user modifies their cookies, such cookie processing implementation will lead to a ReDoS attack.

SheetJS [100]. This is an NPM package (1.4 million weekly downloads) that provides functionality for working with spreadsheets. It has both commercial and open-source versions, and the open-source version is used in more than 72,000 projects. The GitHub repository of the open-source SheetJS project includes source code for an online demo, where we found a function `deepset()` that is vulnerable to prototype pollution. This function is applied to a JSON representation of an `xlsx` file supplied by users. An attacker can supply a specially crafted file to overwrite properties of `Object.prototype`, thus exploiting the vulnerability. While this function is not in a library functionality itself, it is reasonable to expect it to be used by other developers as a foundation for actual projects, potentially retaining the vulnerability.

Highland [14]. This is an NPM package that provides functionality to work with data streams in JavaScript, and it has 47,759 weekly downloads. Highland exports a function `inspect()`, which we found to be vulnerable to a prototype pollution attack. This function assigns custom options to the context object. By supplying a specially crafted options object, a malicious actor can reassign the prototype of the inspection context, potentially affecting the execution of this function (e.g., causing denial of service, unexpected behavior, and procedure bypass).

7 EXPERIMENTAL COMPARISON

Static code analysis tools are popular among software developers, as they can detect vulnerabilities in the early stages of the project's life cycle. These tools also reduce the complexity and time

¹³The PoC QS attack is available at jsfiddle.net/pb6an1dy/

for the developers to fix the vulnerabilities. There are several commercial [71, 92, 103] and open-source [32, 58, 86, 97] code scanning tools, and most of them were initially created to detect mistakes, bad practices, bugs, etc. However, since software security became an essential component of software development, static code scanning tools have also been adapted to detect code vulnerabilities.

For direct comparison with state-of-the-art tools, we explore static code analysis tools that use semantic patterns to detect vulnerabilities. Specifically, we examine tools that include security patterns for JavaScript. We empirically compare our semantic vulnerable function detection approach (see Section 3.3) and publicly available JavaScript vulnerability patterns from three open-source static code scanning tools: Semgrep [97], CodeQL [32], and ESLint security plugin [86].

As of March 2022, Semgrep has 181 security patterns for JavaScript,¹⁴ three of which address prototype pollution, and only one for ReDoS. CodeQL offers 90 patterns, where four of them are for prototype pollution and four for ReDoS.¹⁵ Finally, ESLint has 13 patterns in total, including two for prototype pollution and two for ReDoS.¹⁶ For this experiment, we isolate entries for prototype pollution and ReDoS vulnerabilities from our manually verified dataset of vulnerable functions. As the ground truth dataset, we used all 112 functions with prototype pollution vulnerability, and 164 with ReDoS. The results are summarized in Table 5. The second and fourth columns in the table show the number of functions where the vulnerability was correctly identified by at least one pattern from each tool. Compared to Semgrep and CodeQL, our approach’s detection rate is significantly higher. ESLint detected more functions than us, with a lot of additional matches (discussed below). The additional matches flagged by each tool are of three categories: (i) false positives, (ii) new true positives, not included in our ground truth, and (iii) vulnerability duplicates.

New true positives are the ones that were not previously found and reported to vulnerability registries, and therefore are not listed in the Snyk database. We first identify these cases while working on our custom patterns and testing them on our whole vulnerable JavaScript functions dataset (see Section 3). Vulnerability duplicates occur, when a single function has multiple vulnerable code instances (i.e., “sinks”) related to the same vulnerability. Note that during manual verification, we flag a function as vulnerable if at least one vulnerable “sink” exists.

To further comprehend the precision of each tool, we perform manual analysis of the additional matches for every analyzed approach. The results are summarized in Table 6. Note that for default Semgrep and ESLint patterns, we randomly choose 50 matches for each of the vulnerability types from different functions. As evident from the results, Semgrep and ESLint patterns produced between 90–100% false positives, with zero new true positives for the randomly selected flagged functions. The reason for prototype pollution false positives in ESLint is that these patterns match almost every syntax that has square braces in it. Semgrep also has rather generic prototype pollution patterns and matched simple property access syntax cases, which were irrelevant to other conditions that cause a true vulnerability (e.g., input data from outside of the function, object property assignment). In ReDoS false positives, patterns from both tools did not check if a regular expression is used on data from the outside, flagging instead any use of regex on non-literal data. ESLint additionally flagged regex that it considers “unsafe”, however none of these matches were flagged by our approach.

CodeQL produced no prototype pollution false positives and flagged one new true positive case (also found by our approach). However, all five extra ReDoS findings are falsely positive, as the flagged evil regex is unreachable by any potentially user-controlled input.

As for our approach, most additional matches (78%) are repetitions of vulnerable sinks. We want to highlight 12 extra true positive findings that we flag during this experiment. Five of the

¹⁴semgrep.dev/r?lang=JavaScript&cat=security

¹⁵codeql.github.com/codeql-query-help/javascript-cwe/

¹⁶github.com/nodesecurity/eslint-plugin-security/tree/main/rules

| | Prototype pollution | | ReDoS | |
|-------------------------|---------------------|--------------|------------------|--------------|
| | # true positives | # additional | # true positives | # additional |
| All (manually verified) | 112 | - | 164 | - |
| Our approach | 70 | 24 | 127 | 40 |
| Semgrep | 45 | 73 | 36 | 241 |
| CodeQL | 8 | 9 | 31 | 5 |
| ESLint | 104 | 12377 | 107 | 642 |

Table 5. Results of evaluation of static code scanning tools.

| | Additional prototype pollution | | | Additional ReDoS | | |
|----------------------|--------------------------------|-----------------|------------|------------------|------------------|------------|
| | TP | FP | Dup. | TP | FP | Dup. |
| Our approach (64/64) | 9 (37.5%) | 2 (8.3%) | 13 (54.2%) | 3 (7.5%) | 0 | 37 (92.5%) |
| Semgrep (100/314) | 0 | 49 (98%) | 1 (2%) | 0 | 50 (100%) | 0 |
| CodeQL (14/14) | 1 (11%) | 0 | 8 (89%) | 0 | 5 (100%) | 0 |
| ESLint (100/ 13,019) | 0 | 47 (94%) | 3 (6%) | 0 | 45 (90%) | 5 (10%) |

Table 6. Categorization of additional matches from static code scanning tools evaluation.

matches belong to the list of the most popular NPM packages (i.e., Lodash [54], yargs-parser [116], marked [48], validator [84], SheetJS [100]), and as such, they were included in our dataset of real-world functions. Therefore, they were already found by our framework (see Section 4) and reported to the developers. The rest of the cases concerned less popular packages (e.g., Schema-inspector [45], Ducktype [22], Subtext [37]). In two cases, ReDoS vulnerabilities were found in functions reported for prototype pollution (yargs-parser, Subtext). In three cases, there was an additional prototype pollution vulnerability to the one reported to Snyk (Lodash, Validator, Ducktype). The remaining six cases include prototype pollution vulnerabilities found in ReDoS-vulnerable functions. All new findings were manually reported to the developers and submitted to the MITRE CVE registry.

Note that we compiled our Semgrep patterns from Snyk’s list of previously reported vulnerabilities. As the selected functions for this experiment also came from the same list, it may appear that the dataset is biased in favor of our approach. On the other hand, the dataset contains the most common existing/reported implementations of prototype pollution and ReDoS vulnerabilities, which are expected to be identified by any tool detecting such vulnerabilities.

8 RELATED WORK

In this section, we summarize the most relevant past work on code vulnerability detection. This includes: approaches based on static taint analysis, and textual and semantic similarity detection; the studies that target specifically prototype pollution and ReDoS vulnerabilities; and the ones creating JavaScript vulnerability datasets. We also briefly describe other approaches for vulnerability detection in source code in general.

8.1 Vulnerable code detection

JavaScript static taint analysis. There are several studies that apply taint analysis and vulnerable patterns to detect vulnerable code, utilizing various code representations. Li et al. [67] introduce Object Property Graph (OPG) that tracks tainted object values to detect specifically prototype pollution vulnerabilities in JavaScript. In the following work, Li et al. [68] present Object Dependency Graph representation and the extended framework to model and detect various JavaScript vulnerability types. The WALA framework developed by IBM uses Abstract Syntax Trees (AST) together with call graphs [43]. Initially the framework was developed only for Java, but later the developers introduced JavaScript support. Unfortunately, when we tried to utilize the framework,

we found that their JavaScript normalizer does not convert some of the language’s features properly (e.g., *this* keyword, and new features in ECMAScript 6 such as arrow functions, and block-scoped variables), which results in the wrong representation of the code flow. Joern [51] combines ASTs with Code Property Graphs (CPG), a representation introduced by Yamaguchi et al. [124]. This tool also supports JavaScript. However, we found that it also fails to process several JavaScript language features correctly (e.g., the *module.exports* syntax, “first-class function” concept). Due to this limitation, Joern currently does not support multi-file taint tracking for JavaScript. The CodeQL team presents Dataflow Graphs (DFG) [19] for taint analysis. These graphs are based on the language-agnostic abstract code representation (called CodeQL database), developed by the authors, and therefore can be applied to JavaScript code. As a comparison, our framework uses the traditional AST representations combined with an FDG - a graph structure that allows multi-file taint tracking due to the correct processing of file relationships in JavaScript.

Textual similarity detection. We base our textual similarity detection approach on comparing tokenized code representations similar to previous studies [55, 65, 93]. Apart from tokenization, several other representations are also used for textual similarity, such as Abstract Syntax Trees (AST) [8, 16, 49], Control Flow Graphs (CFG) [129], Program Dependency Graphs (PDG) [66] or a combination of these representations [10].

To optimize the textual similarity-based vulnerability detection, we use both cryptographic hash and locality-sensitive hash comparison on our tokenized functions, similar to [47, 61, 110] and [16, 49] studies respectively. Also Several approaches [29, 91, 118] use machine/deep learning algorithms in order to detect near-duplicate code instances.

Semantic similarity detection. Li et al. [67] uses their OPG to create some of the common prototype pollution vulnerable patterns. Note that since their representation is based on the JavaScript object features, the patterns can only be created to object-related vulnerabilities. Li et al. [68] also provide an ODG representation that allows creating patterns for several non-object-related vulnerabilities (e.g., XSS, code injection, path traversal). GitHub’s CodeQL engine [32] also allows to create semantic patterns. The vulnerable patterns for multiple languages (including JavaScript) are written on a specially designed object-oriented query language called QL. CodeQL has a registry of vulnerable patterns available to the developers. To the best of our knowledge, so far there are no studies that incorporate CodeQL-based patterns (or “queries”) in JavaScript vulnerability studies.

8.2 Vulnerability detection by type

Prototype pollution detection. Prototype pollution was first identified by Snyk researchers in 2018 in a popular NPM package “Lodash” [25]. Kim et al. [59] proposed an approach to detect prototype pollution with pattern recognition and utilize ASTs and CFGs for code representation. However, their approach produces a large number of false positives (50.6%) and false negatives (84.6%). Later, Li et al. [67] presented a novel approach called *ObjLupAnslys* for detecting prototype pollution in NPM packages with a static taint analysis algorithm based on their own representation called Object Property Graph (OPG). However, *ObjLupAnslys*’ false positive rate is relatively high, as it cannot confirm if the found paths between “sources” and “sinks” are valid. It also suffers from poor scalability, as large projects may lead to the path explosion problem. In the following study, Kang et al. [56] introduce the *ProbeTheProto* tool that targets websites. The tool applies dynamic taint analysis and introduces “joint taint flows”. Here, the algorithm first traverses from a “source”, which has to match a predefined set of code samples (e.g., `document.location`, `location.search`), to a “sink”, which is a concrete prototype pollution syntax (meaning that out of all possible syntactical and semantic implementations of prototype pollution, the tool focuses on one specific implementation). Then it traverses from a found “sink” to a “secondary sink”, which has to match a predefined set

of potential exploit code samples (e.g., eval, innerHTML). This approach is very precise but has a low recall, in addition to being time- and resource-intensive. A similar approach *SilentSpring* is developed by Shcherbakov et al. [99] to detect prototype pollution vulnerabilities that lead to remote code execution. The authors use a combination of static and dynamic analysis techniques, as opposed to *ProbeTheProto*, and identify 11 “secondary sinks” that they refer to as “code gadgets”. Additionally, unlike *ProbeTheProto*, Shcherbakov et al. [99] target NPM packages. *SilentSpring* focuses on achieving high recall (75%-96%), but unfortunately results in poor precision (23%-27%).

ReDoS detection. There are several tools [21, 46, 115] that can analyze regular expressions and flag the dangerous ones in a given codebase. Recall that a ReDoS vulnerability is exploitable under two conditions: (1) the regular expression is constructed in a way that certain examined strings take exponential processing time; and (2) the regular expression is applied to the data potentially controlled by an attacker. While the second condition is shared by multiple vulnerability types, the first one is targeted separately by studies that are focused on creating exploits/attacks on regular expressions using static [62, 70], dynamic [74, 112] and hybrid [69] approaches. Static regex detectors suffer from low coverage of different regular expression features, and it is difficult to adapt them to detect continuously added features [74]. Additionally, Li et al. [69] report that both static and dynamic approaches tend to suffer from low recall (e.g., static tools: RXXR2 [62] - 2.27%, NFAA [122] - 8.73%; dynamic tools: SDL [114] - 1.06%, ReScue [101] - 1.86%). ReDoSHunter achieves high precision and recall rates (100% in both), although it detects only most common regex extensions and characters and does not consider less frequently used ones.

8.3 Vulnerability datasets

To create vulnerability datasets, several past studies collect meta-information, packages, and functions/code-snippets from known/reported vulnerabilities. For example, VulData7 [50] uses the National Vulnerability Database (NVD) to extract information on vulnerabilities. However, it does not extract the code from files, and the data is not classified by the language. Ferenc et al. [30] use the GitHub repositories of the Snyk vulnerability database [104] and the Node Security Platform (NSP [83]) as vulnerability sources, and create the first publicly available JavaScript dataset consisting of 12,125 functions, 1,496 of which are flagged as vulnerable. However, the Snyk database on GitHub that they used has not been maintained since 2018, and NSP has been made private by NPM and thus is no longer publicly available. Mosolygó et al. [79] analyzed vulnerable functions from the dataset in [30]. They manually filtered all 1,496 vulnerable functions and extracted only 443 functions, which they deemed to be actually vulnerable. However, when examining the resulting filtered set of functions, we noticed that it still contained some false positives. This lack of a reliable vulnerability dataset is the primary motivation for creating our own dataset, which is vital for deriving our approach (e.g., Semgrep rules). We use up-to-date and well-maintained vulnerability sources (Snyk [103] and VulnCode-DB project [34]), and design a comprehensive methodology to avoid flagging unrelated or benign functions as vulnerable.

Another popular approach for dataset compilation is to perform a search through GitHub projects, find commits that include specific keywords in the commit message, such as “bug”, “fix”, “CVE” etc. and extract functions from commits [10, 61, 65]. This method relies on developers’ use of proper conventions when they give a name to their commits;¹⁷ as such, only vulnerabilities with correctly formulated commit messages can be detected.

¹⁷conventionalcommits.org/en/v1.0.0/

8.4 Alternative vulnerable code detection approaches

Several textual similarity tools aim to detect vulnerable code regardless of the programming language used. To be able to do that, they require any code to be brought into a generic abstracted representation, such as a high-level tree-based representation [119], PDG and CFG [10, 61, 118]. However, cross-language tools are mostly implemented for textual similarity search due to the specificity of each language's design, so they can only detect nearly identical copies of the code.

Yamaguchi et al. [124] target code vulnerabilities with a semantic-similarity approach. They model patterns for several vulnerability types in C/C++ (e.g., buffer overflow, memory disclosure) by combining multiple function representations into a Code Property Graph (CPG) and extracting properties that form a vulnerable pattern. Then they use the same representations to transform the target dataset, and perform a pattern search against the obtained data representations.

Russell et al. [91] attempt to automatically extract vulnerable patterns/features by using convolutional and recurrent neural networks on vectorized C/C++ code, reaching a 0.56 value for the F1 measure. Sheneamer and Kalita [102] extract features from vectorized AST and PDG representations of Java code using 15 previously published different ML algorithms. Both studies, as well as multiple others [29, 75, 130] focus on the textual similarity, and thus their target datasets mostly consist of nearly-duplicate code instances.

Dynamic analysis. The dynamic analysis methods commonly used to detect vulnerabilities in code are symbolic execution, concolic testing, and fuzzing. However, to the best of our knowledge, there are presently no studies that use concolic testing for JavaScript vulnerability detection.

The symbolic execution method simulates program execution by replacing concrete values with symbolic variables for inputs. The program is executed several times, each time modifying the symbolic inputs to reach the execution of new paths. It has been used for finding concrete vulnerabilities in code, such as file parsing vulnerability [41], DoS [81, 117], code injection [15]. Saxena et al. [95] use symbolic execution for JavaScript to find client-side code injection vulnerabilities.

Fuzz testing executes the program with self-generated inputs, which may contain invalid, unexpected, or random values. The analysis tool then monitors the behavior of the program and reports undesired results, e.g., data leaks, unexpected modifications, and program crashes. The objective is to generate an input that exploits vulnerabilities in the program. Fuzzers are most commonly used for DoS detection [94, 113], but also for XSS [26], and buffer overflow [121]. Fuzzing algorithms have also been developed for finding vulnerabilities in JavaScript [39, 42]. The main challenge in fuzzing is the difficulty of automatically generating useful inputs. For example, the CodeAlchemist fuzzer [42] can generate more than 40% semantically invalid JavaScript inputs [39].

Several studies conduct dynamic taint analysis on JavaScript [98, 128], although only a few target vulnerability detection [57]. Dynamic taint analysis shares the disadvantages with other dynamic approaches (i.e., low coverage, high complexity, time- and resource-consuming).

9 CONCLUSIONS

In summary, we propose a framework for function-level detection of JavaScript vulnerabilities in the wild with further exploitability verification, shifting the focus from package-level vulnerability tracking/measurements considered in the past work. We also design a semi-automated vulnerable function collection mechanism to build a reliable dataset of known vulnerable JavaScript functions. Our dataset contains 1,360 verified vulnerable functions. By testing 9,204,654 real-world JavaScript functions from popular NPM packages, Chrome web extensions, and websites, we detected 124,934 potentially vulnerable functions with a precision of 94.5% (calculated based on a small randomly chosen dataset) by vulnerable pattern search. We then checked the dataset against fuzzy and cryptographic hashes, and detected 131 and 965 vulnerable functions with the estimated precision

of 100% and 98%, respectively. Our static taint analysis on the 5,389 findings in NPM packages verified 301 cases with no false positives (19 of them already have CVE IDs). All cases were privately reported to the npm projects repository owners in terms of responsible disclosure. In addition, we conducted an in-depth analysis of 20 detected vulnerabilities from 17 projects and described the attack vectors that can be exploited for these vulnerabilities. Moreover, we performed successful proof-of-concept attacks on seven projects by exploiting the detected vulnerabilities. We obtained 194 CVE IDs from MITRE, from which 25 are already published and have “critical” and “high” severity ratings. Finally, for reproducibility and further research in JavaScript security, all the outcomes of our work are publicly available at: <https://github.com/Marynk/JavaScript-vulnerability-detection>.

REFERENCES

- [1] 1e0ng. 2020. SimHash. <https://github.com/1e0ng/SimHash>.
- [2] AcornJS. 2012. Acorn: A Tiny, Fast JavaScript Parser. <https://github.com/acornjs/acorn>.
- [3] Mahmoud Alfadel, Diego Elias Costa, Emad Shihab, and Mouafak Mkhallati. 2021. On the Use of Dependabot Security Pull Requests. In *MSR’21*. Madrid, Spain, 254–265.
- [4] Andrew Smith. 2021. Content Spoofing. https://owasp.org/www-community/attacks/Content_Spoofing.
- [5] Jeremy Ashkenas. 2009. Underscore.js. <https://github.com/jashkenas/underscore>.
- [6] Babel. 2020. Babel progress on ECMAScript proposals. <https://github.com/babel/proposals>.
- [7] Balderdash Design Co. 2012. Sails.js: The MVC Framework for Node.js. <https://sailsjs.com/>.
- [8] Ira Baxter, Andrew Yahin, Leonardo de Moura, Marcelo Sant’Anna, and Lorraine Bier. 1998. Clone Detection Using Abstract Syntax Trees. In *ICSM’98*. Bethesda, MD, USA, 368–377.
- [9] Fabian Beuke. 2021. GitHub Language Statistics. https://madnight.github.io/github/#/pull_requests/2021/1.
- [10] Benjamin Bowman and H. Howie Huang. 2020. VGRAPH: A Robust Vulnerable Code Clone Detection System Using Code Property Triplets. In *IEEE EuroS&P 2020*. Virtual, 53–69.
- [11] BuiltWith.com. 2022. AngularJS Usage Statistics. <https://trends.builtwith.com/javascript/Angular-JS>.
- [12] BuiltWith.com. 2022. SailsJS Usage Statistics. <https://trends.builtwith.com/framework/SailsJS>.
- [13] Caolan McMahon. 2011. Async. <https://caolan.github.io/async/v3/>.
- [14] Caolan McMahon. 2014. Highland: The High-Level Streams Library for Node.js and the Browser. <https://caolan.github.io/highland/>.
- [15] Daming D Chen, Maverick Woo, David Brumley, and Manuel Egele. 2016. Towards Automated Dynamic Analysis for Linux-based Embedded Firmware. In *NDSS Symposium*. San Diego, CA, USA.
- [16] Wai Cheung, Sukyoung Ryu, and Sunggun Kim. 2015. Development Nature Matters: An Empirical Study of Code Clones in JavaScript Applications. *Empirical Software Engineering* 21 (March 2015), 517–564.
- [17] Bodin Chinthanet, Raula Gaikovina Kula, Shane McIntosh, Takashi Ishio, Akinori Ihara, and Kenichi Matsumoto. 2021. Lags in the Release, Adoption, and Propagation of NPM Vulnerability Fixes. *Empirical Software Engineering* 26 (May 2021).
- [18] Bodin Chinthanet, Serena Elisa Ponta, Henrik Plate, Antonino Sabetta, Raula Gaikovina Kula, Takashi Ishio, and Kenichi Matsumoto. 2020. Code-Based Vulnerability Detection in Node.js Applications: How far are we?. In *35th IEEE/ACM International Conference on Automated Software Engineering*. Los Alamitos, CA, USA, 1199–1203.
- [19] CodeQL. 2019. CodeQL: About Data Flow Analysis. <https://codeql.github.com/docs/writing-codeql-queries/about-data-flow-analysis/>.
- [20] Dan Hubbard. 2016. Cisco Umbrella 1 Million. <https://umbrella.cisco.com/blog/cisco-umbrella-1-million>.
- [21] Jamie Davis. 2018. Detect Vulnerable Regexes in your Project. <https://github.com/davisjam/vuln-regex-detector>.
- [22] Jos de Jong. 2013. Ducktype. <https://github.com/josdejong/ducktype>.
- [23] Alexandre Decan, Tom Mens, and Eleni Constantinou. 2018. On the Impact of Security Vulnerabilities in the NPM Package Dependency Network. In *Conference on Mining Software Repositories, MSR 2018*. ACM, 181–191.
- [24] Cypress Data Defense. 2020. Differences Between Static Code Analysis and Dynamic Testing. <https://www.cypressdatadefense.com/blog/static-and-dynamic-code-analysis/>.
- [25] Ben Dickson. 2020. Prototype Pollution: The Dangerous and Underrated Vulnerability Impacting JavaScript Applications. <https://portswigger.net/daily-swig/prototype-pollution-the-dangerous-and-underrated-vulnerability-impacting-javascript-applications>.
- [26] Adam Doupé, Ludovico Cavedon, Christopher Kruegel, and Giovanni Vigna. 2012. Enemy of the State: A State-Aware Black-Box Web Vulnerability Scanner. In *21st USENIX Security Symposium (USENIX Security 12)*. 523–538.
- [27] Ruian Duan, Omar Alrawi, Ranjita Pai Kasturi, Ryan Elder, Brendan Saltaformaggio, and Wenke Lee. 2021. Towards Measuring Supply Chain Attacks on Package Managers for Interpreted Languages. In *NDSS’21*. Virtual.

- [28] Escomplex. 2015. Escomplex: Software Complexity Analysis of JavaScript Abstract Syntax Trees. <https://github.com/escomplex/escomplex>.
- [29] Chunrong Fang, Zixi Liu, Yangyang Shi, Jeff Huang, and Qingkai Shi. 2020. Functional Code Clone Detection with Syntax and Semantics Fusion Learning. In *ACM Symposium on Software Testing and Analysis*. Online, 516–527.
- [30] Rudolf Ferenc, Péter Hegedüs, Péter Gyimesi, Gábor Antal, Dénes Bán, and Tibor Gyimóthy. 2019. Challenging Machine Learning Algorithms in Predicting Vulnerable JavaScript Functions. In *RAISE@ICSE'19*. 8–14.
- [31] OpenJS Foundation. 2014. Espree. <https://www.npmjs.com/package/espree>.
- [32] GitHub. 2019. CodeQL, Semantic Code Analysis Engine. <https://codeql.github.com/>.
- [33] Google. 2010. AngularJS. <https://angularjs.org/>.
- [34] Google. 2019. The Vulnerable Code Database (Vulncode-DB). <https://www.vulncode-db.com>.
- [35] Gravatar.com. 2007. Gravatar. <https://en.gravatar.com/>.
- [36] James Halliday. 2013. Minimist. <https://www.npmjs.com/package/minimist>.
- [37] hapi.js team. 2014. @hapi/subtext. <https://github.com/hapijs/subtext>.
- [38] Jordan Harband. 2014. QS, a Query String Parsing and Stringifying Library. <https://github.com/ljharb/qs>.
- [39] Xiaoyu He, Xiaofei Xie, Yuekang Li, Jianwen Sun, Feng Li, Wei Zou, Yang Liu, Lei Yu, Jianhua Zhou, Wenchang Shi, and Wei Huo. 2021. SoFi: Reflection-Augmented Fuzzing for JavaScript Engines. In *ACM CCS'21*. 2229–2242.
- [40] Ariya Hidayat. 2012. ECMAScript Parsing Infrastructure for Multipurpose Analysis. <https://esprima.org/>.
- [41] Chaojian Hu, Zhoujun Li, Jinxin Ma, Tao Guo, and Zhiwei Shi. 2012. File Parsing Vulnerability Detection with Symbolic Execution. In *Theoretical Aspects of Software Engineering (TASE'12)*. Washington DC, USA, 135–142.
- [42] Han HyungSeok, Oh DongHyeon, and Kil Cha Sang. 2019. CodeAlchemist: Semantics-Aware Code Generation to Find Vulnerabilities in JavaScript Engines. In *NDSS'19*. San Diego, CA, USA.
- [43] IBM. 2015. The T. J. Watson Libraries for Analysis (WALA). <http://wala.sourceforge.net>.
- [44] GitHub Inc. 2017. GitHub Advisory Database. <https://github.com/advisories>.
- [45] Schema inspector team. 2014. Schema-inspector. <https://github.com/schema-inspector/schema-inspector>.
- [46] James Halliday. 2013. saferegex. <https://github.com/substack/safe-regex>.
- [47] Jiyong Jang, Maverick Woo, and David Brumley. 2012. ReDeBug: Finding Unpatched Code Clones in Entire OS Distributions. *IEEE Symposium on Security and Privacy* 37, 6 (May 2012), 48–62.
- [48] Christopher Jeffrey. 2011. Marked. <https://github.com/markedjs/marked>.
- [49] Lingxiao Jiang, Ghassan Mishserghi, Zhendong Su, and Stéphane Gloudu. 2007. DECKARD: Scalable and Accurate Tree-Based Detection of Code Clones. In *ICSE'07*. Minneapolis, MN, USA, 96–105.
- [50] Matthieu Jimenez, Yves Le Traon, and Mike Papadakis. 2018. Enabling the Continuous Analysis of Security Vulnerabilities with VulData7. In *IEEE Source Code Analysis and Manipulation*. Madrid, Spain, 56–61.
- [51] Joern. 2019. Joern - The Bug Hunter's Workbench. joern.io.
- [52] Joernio. 2021. AST Generator. <https://github.com/joernio/astgen>.
- [53] John-David Dalton. 2009. Lodash. <https://www.npmjs.com/package/lodash>.
- [54] John-David Dalton. 2009. Lodash: A Modern JavaScript Utility Library. <https://lodash.com/>.
- [55] Toshihiro Kamiya, Shinji Kusumoto, and Katsuro Inoue. 2002. CCFinder: A Multilinguistic Token-based Code Clone Detection System for Large Scale Source Code. *IEEE Transactions on Software Engineering* 28, 7 (2002), 654–670.
- [56] Zifeng Kang, Song Li, and Yinzhi Cao. 2022. Probe the Proto: Measuring Client-Side Prototype Pollution Vulnerabilities of One Million Real-World Websites. In *NDSS'22*. San Diego, CA, USA.
- [57] Rezwana Karim, Frank Tip, Alena Sochůrková, and Koushik Sen. 2020. Platform-Independent Dynamic Taint Analysis for JavaScript. *IEEE Transactions on Software Engineering* 46, 12 (October 2020), 1364–1379.
- [58] Kestrel Technology. 2020. CodeHawk Tool Suite. <https://github.com/static-analysis-engineering/codehawk>.
- [59] Hee Kim, Ji Kim, Ho Oh, Beom Lee, Si Mun, Jeong Shin, and Kyounggon Kim. 2022. DAPP: Automatic Detection and Analysis of Prototype Pollution Vulnerability in Node.js Modules. *International Journal of Information Security* 21 (February 2022), 1–23.
- [60] Seokmo Kim, R. Kim, and Young Park. 2016. Software Vulnerability Detection Methodology Combined with Static and Dynamic Analysis. *Wireless Personal Communications* 89 (August 2016), 1–17.
- [61] Seulbae Kim, Seunghoon Woo, Heejo Lee, and Hakjoo Oh. 2017. VUDDY: A Scalable Approach for Vulnerable Code Clone Discovery. In *IEEE Symposium on Security and Privacy*. 595–614.
- [62] James Kirrage, Asiri Rathnayake Mudiyansele, and Hayo Thielecke. 2013. Static Analysis for Regular Expression Denial-of-Service Attacks. In *Network and System Security (NSS'13)*. 135–148.
- [63] Maryna Kluban, Mohamman Mannan, and Amr Yousef. 2022. On Measuring Vulnerable JavaScript Functions in the Wild. In *Proceedings of the 2022 ACM Asia Conference on Computer and Communications Security (ASIA CCS '22)*, May 30–June 3, 2022, Nagasaki, Japan.
- [64] Aditya Kurniawan, Bahtiar Saleh Abbas, Agung Trisetyarso, and Sani Muhammad Isa. 2018. Static Taint Analysis Traversal with Object Oriented Component for Web File Injection Vulnerability Pattern Detection. *Procedia Computer*

Science 135 (2018), 596–605. <https://www.sciencedirect.com/science/article/pii/S1877050918315230>

- [65] Hongzhe Li, Hyuckmin Kwon, Jonghoon Kwon, and Heejo Lee. 2016. CLORIFI: Software Vulnerability Discovery Using Code Clone verification. *Concurrency and Computation: Practice and Experience* 28, 6 (May 2016), 1900–1917.
- [66] Jingyue Li and Michael D. Ernst. 2012. CBCD: Cloned Buggy Code Detector. In *ICSE'12*. Zurich, Switzerland, 310–320.
- [67] Song Li, Mingqing Kang, Jianwei Hou, and Yinzhi Cao. 2021. Detecting Node.js Prototype Pollution Vulnerabilities via Object Lookup Analysis. In *ESEC/FSE'21*. Athens, Greece, 268–279.
- [68] Song Li, Mingqing Kang, Jianwei Hou, and Yinzhi Cao. 2022. Mining Node.js Vulnerabilities via Object Dependence Graph and Query. In *31st USENIX Security Symposium*. Boston, MA, USA.
- [69] Yeting Li, Zixuan Chen, Jialun Cao, Zhiwu Xu, Qiancheng Peng, Haiming Chen, Liyuan Chen, and Shing-Chi Cheung. 2021. ReDoSHunter: A Combined Static and Dynamic Approach for Regular Expression DoS Detection. In *30th USENIX Security Symposium*. Virtual, 3847–3864.
- [70] Yinxi Liu, Mingxue Zhang, and Wei Meng. 2021. Revealer: Detecting and Exploiting Regular Expression Denial-of-Service Vulnerabilities. In *IEEE Symposium on Security and Privacy*. Online, 1468–1484.
- [71] Veracode LLC. 2006. Veracode Static Analysis. <https://www.veracode.com/products/binary-static-analysis-sast>.
- [72] Abdalla Wasef Marashdih, Zarul Fitri Zaaba, and Khaled Suwais. 2023. An Enhanced Static Taint Analysis Approach to Detect Input Validation Vulnerability. *Journal of King Saud University - Computer and Information Sciences* 35, 2 (2023), 682–701.
- [73] Sebastian McKenzie. 2014. Babel, the JavaScript Compiler. <https://babeljs.io/>.
- [74] Robert McLaughlin, Fabio Pagani, Noah Spahn, Christopher Kruegel, and Giovanni Vigna. 2022. Regulator: Dynamic Analysis to Detect ReDoS. In *31st USENIX Security Symposium*. Boston, MA, USA.
- [75] Nikita Mehrotra, Navdha Agarwal, Piyush Gupta, Saket Anand, David Lo, and Rahul Purandare. 2021. Modeling Functional Similarity in Source Code with Graph-Based Siamese Networks. *IEEE Transactions on Software Engineering* (August 2021).
- [76] Tomáš Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *Conference on Learning Representations*. Scottsdale, Arizona, USA.
- [77] MITRE. 2006. Common Weakness Enumeration. <https://cwe.mitre.org/>.
- [78] MITRE. 2021. Common Vulnerabilities and Exposures. <https://cve.mitre.org/>.
- [79] Balázs Mosolygó, Norbert Vándor, Gábor Antal, Péter Hegedűs, and Rudolf Ferenc. 2021. Towards a Prototype Based Explainable JavaScript Vulnerability Prediction Model. In *2021 International Conference on Code Quality (ICQ)*. Moscow, Russia, 15–25.
- [80] Mozilla. 2021. Polyfill. <https://developer.mozilla.org/en-US/docs/Glossary/Polyfill>.
- [81] James Newsome, David Brumley, Jason Franklin, and Dawn Song. 2006. Replayer: Automatic Protocol Replay by Binary Analysis. In *ACM CCS'06*. Alexandria, VA, USA, 311–321.
- [82] npm. 2010. Node Package Registry. <https://www.npmjs.com/>.
- [83] npm. 2018. The Node Security Platform Service is Shutting Down. <https://blog.npmjs.org/post/175511531085/insert-title-here.html>.
- [84] Chris O'Hara. 2018. Validator.js. <https://github.com/validatorjs/validator.js>.
- [85] OSA 2018. OpenStaticAnalyzer. <https://openstaticanalyzer.github.io/>.
- [86] The Node Security Platform. 2017. ESLint Security Plugin. <https://www.npmjs.com/package/eslint-plugin-security>.
- [87] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco - A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *NDSS'19*. San Diego, CA, USA.
- [88] PortSwigger.net. 2021. DOM-based Open Redirection. <https://portswigger.net/web-security/dom-based/open-redirection>.
- [89] Niels Provos. 2015. A JavaScript-based DDoS Attack as Seen by Safe Browsing. <https://security.googleblog.com/2015/04/a-javascript-based-ddos-attack-as-seen.html>.
- [90] Andris Reinman. 2014. NODEMAILER, Send Emails from Node.js – Easy as Cake! <https://nodemailer.com/>.
- [91] Rebecca Russell, Louis Kim, Lei Hamilton, Tomo Lazovich, Jacob Harer, Onur Ozdemir, Paul Ellingwood, and Marc McConley. 2018. Automated Vulnerability Detection in Source Code Using Deep Representation Learning. In *17th IEEE International Conference on Machine Learning and Applications (ICMLA)*. Orlando, FL, USA, 757–762.
- [92] SonarSource S.A. 2008. SonarCube: Code Security, for Developers. <https://www.sonarqube.org/features/security/>.
- [93] Hitesh Sajani, Vaibhav Saini, Jeffrey Svajlenko, Chanchal K. Roy, and Cristina V. Lopes. 2016. SourcererCC: Scaling Code Clone Detection to Big-Code. In *ICSE'16*. Austin, TX, USA, 1157–1168.
- [94] Raimondas Sasnauskas and John Regehr. 2014. Intent Fuzzer: Crafting Intents of Death. In *WODA+PERTEA'14*. San Jose, CA, USA, 1–5.
- [95] Prateek Saxena, Devdatta Akhawe, Steve Hanna, Feng Mao, Stephen McCamant, and Dawn Song. 2010. A Symbolic Execution Framework for JavaScript. In *IEEE Symposium on Security and Privacy*. Oakland, CA, USA, 513–528.
- [96] Scott Sauyet, Buzz de Cafe. 2020. Ramda. <https://ramdajs.com/>.

- [97] Semgrep.dev. 2020. Semgrep–Find Bugs and Enforce Code Standards. <https://semgrep.dev/docs/>.
- [98] Koushik Sen, Swaroop Kalasapur, Tasneem G. Brutch, and Simon J. Gibbs. 2013. Jalangi: a Selective Record-replay and Dynamic Analysis Framework for JavaScript. In *ESEC/FSE’13*. Saint Petersburg, Russia, 488–498.
- [99] Mikhail Shcherbakov, Musard Balliu, and Cristian-Alexandru Staicu. 2023. Silent Spring: Prototype Pollution Leads to Remote Code Execution in Node.js. In *USENIX Security Symposium*. Anaheim, CA, USA.
- [100] Sheetjs LLC. 2012. SheetJS Community Edition – Spreadsheet Data Toolkit. <https://github.com/SheetJS/sheetjs/>.
- [101] Yuju Shen, Yanyan Jiang, Chang Xu, Ping Yu, Xiaoxing Ma, and Jian Lu. 2018. ReScue: Crafting Regular Expression DoS Attacks. In *ACM/IEEE ASE’18*. Montpellier, France, 225–235.
- [102] Abdullah Sheneamer and Jugal Kalita. 2016. Semantic Clone Detection Using Machine Learning. In *IEEE Conference on Machine Learning and Applications*. Pasadena, CA, USA, 1024–1028.
- [103] Snyk.io. 2015. Snyk Vulnerability Database. <https://snyk.io/product/vulnerability-database/>.
- [104] Snyk.io. 2018. GitHub Snyk Vulnerability Database. <https://github.com/snyk/vulnerabilitydb>.
- [105] Snyk.io. 2018. Prototype Pollution in merge, mergeWith, and defaultsDeep, Lodash. <https://snyk.io/vuln/SNYK-DOTNET-LODASH-540455>.
- [106] Snyk.io. 2019. Prototype Pollution in defaultsDeep, Lodash. <https://snyk.io/vuln/SNYK-DOTNET-LODASH-540457>.
- [107] Snyk.io. 2020. Prototype Pollution in set/setWith, Lodash. <https://snyk.io/vuln/SNYK-JS-LODASH-608086>.
- [108] Snyk.io. 2020. Prototype Pollution in zipObjectDeep, Lodash. <https://snyk.io/vuln/SNYK-JS-LODASH-590103>.
- [109] Softwaretestinghelp.com. 2021. JavaScript Injection Tutorial: Test and Prevent JS Injection Attacks On Website. <https://www.softwaretestinghelp.com/javascript-injection-tutorial/>.
- [110] Xiaonan Song, Aimin Yu, Haibo Yu, Shirun Liu, Xin Bai, Lijun Cai, and Dan Meng. 2020. Program Slice based Vulnerable Code Clone Detection. In *IEEE TrustCom’20*. Guangzhou, China, 293–300.
- [111] Stackoverflow.com. 2020. Developer Survey. <https://insights.stackoverflow.com/survey/2020#most-popular-technologies>.
- [112] Cristian-Alexandru Staicu and Michael Pradel. 2018. Freezing the Web: A Study of ReDoS Vulnerabilities in JavaScript-based Web Servers. In *USENIX Security Symposium*. Baltimore, MD, USA, 361–376.
- [113] Nick Stephens, John Grosen, Christopher Salls, Andrew Dutcher, Ruoyu Wang, Jacopo Corbetta, Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. 2016. Driller: Augmenting Fuzzing Through Selective Symbolic Execution. In *NDSS*. San Diego, CA, USA, 1–16.
- [114] Bryan Sullivan. 2010. SDL Regex Fuzzer. <https://www.microsoft.com/security/blog/2010/10/12/new-tool-sdl-regex-fuzzer/>.
- [115] Superhuman Labs. 2016. RXXR2 Regular Expression Static Analyzer. <https://github.com/superhuman/rxxr2>.
- [116] Yargs team. 2016. yargs-parser. <https://github.com/yargs/yargs-parser>.
- [117] Nikolai Tillmann and Jonathan de Halleux. 2008. Pex–White Box Test Generation for .NET. In *Tests and Proofs (TAP’08)*. Prato, Italy, 134–153.
- [118] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2018. Deep Learning Similarities from Different Representations of Source Code. In *MSR’18*. Gothenburg, Sweden, 542–553.
- [119] Tijana Vislavski, Gordana Rakic, Nicolás Cardozo, and Zoran Budimac. 2018. LICCA: A Tool for Cross-language Clone Detection. In *IEEE SANER’18*. Campobasso, Italy, 512–516.
- [120] W3Techs. 2021. Usage Statistics of JavaScript as Client-side Programming Language on Websites. <https://w3techs.com/technologies/details/cp-javascript>.
- [121] Tielei Wang, Tao Wei, Guofei Gu, and Wei Zou. 2010. TaintScope: A Checksum-aware Directed Fuzzing Tool for Automatic Software Vulnerability Detection. In *IEEE Symposium on Security and Privacy*. Berkeley, CA, USA, 497–512.
- [122] Nicolaas Weideman, Brink Van Der Merwe, Martin Berglund, and Bruce Watson. 2016. Analyzing Matching Time Behavior of Backtracking Regular Expression Matchers by Using Ambiguity of NFA. In *Implementation and Application of Automata*. 322–334.
- [123] Adar Weidman. 2019. Regular Expression Denial of Service - ReDoS. https://owasp.org/www-community/attacks/Regular_expression_Denial_of_Service_-_ReDoS.
- [124] Fabian Yamaguchi, Nico Golde, Daniel Arp, and Konrad Rieck. 2014. Modeling and Discovering Vulnerabilities with Code Property Graphs. In *IEEE Symposium on Security and Privacy*. San Jose, CA, USA, 590–604.
- [125] Yeoman team. 2012. grunt-usemin. <https://www.npmjs.com/package/grunt-usemin/>.
- [126] Rodrigo Elizalde Zapata, Raula Gaikovina Kula, Bodin Chinthanet, Takashi Ishio, Kenichi Matsumoto, and Akinori Ihara. 2018. Towards Smoother Library Migrations: A Look at Vulnerable Dependency Migrations at Function Level for NPM JavaScript Packages. In *IEEE ICSME’18*. Madrid, Spain, 559–563.
- [127] Ahmed Zerouali, Valerio Cosentino, Tom Mens, Gregorio Robles, and Jesús M. González-Barahona. 2019. On the Impact of Outdated and Vulnerable JavaScript Packages in Docker Images. In *IEEE Conference on Software Analysis, Evolution and Reengineering*. Hangzhou, China, 619–623.
- [128] Matt Zeunert. 2016. FromJS. <https://www.fromjs.com/>.

- [129] Minmin Zhou, Jinfu Chen, Yisong Liu, Hilary Ackah-Arthur, Shujie Chen, Qingchen Zhang, and Zhifeng Zeng. 2019. A Method for Software Vulnerability Detection Based on Improved Control Flow Graph. *Wuhan University Journal of Natural Sciences* 24 (April 2019), 149–160.
- [130] Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. 2019. Devign: Effective Vulnerability Identification by Learning Comprehensive Program Semantics via Graph Neural Networks. In *Conference on Neural Information Processing Systems*. Vancouver, Canada, 10197–10207.
- [131] Markus Zimmermann, Cristian-Alexandru Staicu, Cam Tenny, and Michael Pradel. 2019. Small World with High Risks: A Study of Security Threats in the NPM Ecosystem. In *USENIX Security Symposium*. Santa Clara, CA, 995–1010.

APPENDIX

| Vulnerability type | # of entries | # of functions |
|----------------------------|--------------|----------------|
| Cross-Site Scripting | 228 | 1183 |
| Code Injection | 167 | 983 |
| ReDoS | 121 | 582 |
| Prototype Pollution | 101 | 356 |
| Denial of Service | 45 | 245 |
| Directory Traversal | 42 | 300 |
| Information Exposure | 23 | 201 |
| Insecure Download Protocol | 19 | 35 |
| Improper Input Validation | 18 | 50 |
| Request Forgery | 17 | 100 |
| Memory Exposure | 13 | 54 |
| Insecure File Access | 13 | 121 |
| Procedure Bypass | 12 | 316 |
| Improper Auth | 8 | 79 |
| Insecure Defaults | 7 | 26 |
| Improper Cred. Protection | 7 | 11 |
| Timing Attack | 6 | 19 |
| Open Redirect | 6 | 10 |
| Insecure Randomness | 6 | 15 |
| Improper Access Control | 6 | 18 |
| Other | 30 | 166 |
| Total | 895 | 4870 |

Table 7. Distribution of vulnerability types among all collected entries.