

# Gelato: Feedback-driven and Guided Security Analysis of Client-side Applications

Behnaz Hassanshahi, Hyunjun Lee, Paddy Krishnan, Jörn Güy Suß  
Oracle Labs, Brisbane, Australia  
{firstname.lastname}@oracle.com

**Abstract**—Even though a lot of effort has been invested in analyzing client-side web applications during the past decade, the existing tools often fail to deal with the complexity of modern JavaScript applications. However, from an attacker point of view, the client side of such web applications can reveal invaluable information about the server side. In this paper, first we study the existing tools and enumerate the most crucial features a security-aware client-side analysis should be supporting. Next, we propose GELATO to detect vulnerabilities in modern client-side JavaScript applications that are built upon complex libraries and frameworks. In particular, we take the first step in closing the gap between state-aware crawling and client-side security analysis by proposing a feedback-driven security-aware guided crawler that is able to analyze complex frameworks automatically, and increase the coverage of security-sensitive parts of the program efficiently. Moreover, we propose a new lightweight client-side taint analysis that outperforms the start-of-the-art tools, requires no modification to browsers, and reports non-trivial taint flows on modern JavaScript applications.

## I. INTRODUCTION

The client-server design of web applications allows developers to perform the logic and computation of their applications on both the client and server sides. The powerful and rich features of modern browsers have enticed developers to use languages, such as JavaScript, HTML and CSS, to implement complex and highly interactive user interfaces on the client side. All the client-side source-code that runs in the browsers is available to everyone, including attackers. While attackers have been actively exploiting client-side vulnerabilities, such as DOM-based XSS [23] for almost a decade now, the increase in complexity of JavaScript applications and frameworks has rendered the existing tools ineffective to prevent them. Moreover, from an attacker point of view, the client side of web applications can reveal invaluable information about the server side, such as REST end points, validation routines, and database queries.

In this paper, we propose GELATO to detect vulnerabilities in modern client-side JavaScript applications, which are often built upon complex libraries and frameworks. More specifically, we address the practical challenges necessary to deal with such web applications. For this purpose, first we enumerate the most crucial features a security-aware client-side analysis should be supporting, and report on the status of state-of-the-art dynamic analysis tools accordingly. We limit our study to dynamic analysis tools because they suit the dynamic nature of JavaScript applications.

Even though a lot of effort has been invested in analyzing client-side web applications during the past decade, the existing tools often fail to deal with the complexity of modern

JavaScript applications. To understand why, we provide an overview of the most essential features necessary to analyze a client-side application to detect security issues. Table I summarizes the features supported by the existing tools and compares them against GELATO.

From this table, we can make several interesting observations: (1) all of the tools except for GELATO miss at least six features necessary to analyze modern real-world JavaScript applications; (2) existing tools either focus on improving the crawling technology or security analysis, and there is no tool that focuses on both aspects, while one is not independent of the other; (3) most of the crawlers that focus on improving the coverage (lines of JavaScript code or number of discovered hyperlinks) are not guided towards specific locations, which can be essential for security analysis in practice; and (4) GELATO is the only tool that directly addresses complex libraries and frameworks, reducing the need to have manually crafted models.

Having support for all of the features listed in Table I to perform a client-side security analysis is challenging. In this paper, we take the first step in bringing together a state-aware crawler and a client-side security analysis, and in closing the gap between them. One of the challenges that state-aware crawlers face is that the search space they need to explore (number of paths in the state graph) can grow exponentially. Therefore, traversing the whole search space can result in poor performance. In practice, we have found the efficiency of crawlers to be one of the most essential factors for them to be used for testing. However, it is possible to devise an algorithm to cover specific paths of interest efficiently, without having to traverse the whole search space. By guiding GELATO towards specific targets, we improve the performance while achieving acceptable coverage. Furthermore, GELATO tries to incorporate most of the features in Table I that are crucial for a security analysis to detect vulnerabilities effectively in practice.

To analyze modern JavaScript applications for security vulnerabilities, crawlers play a substantial role by providing inputs to the program. However, the state-of-the-art security analysis tools [42], [49] either don't have a crawler or provide an insufficient support. Hence, while they can report vulnerabilities in the first page of applications for Alexa top [1] websites, they cannot report vulnerabilities in the other parts of the applications as shown in Sec. VI. During the past few years, many crawling techniques have been proposed to explore the search space on the client side.

Feature	GELATO	Crawljax [44]	jAk [50]	Feedex [46]	WATEG [55]	Artemis+SID [39]	Artform	DexterJS [49]	Kudzu [52]	CTT [42]
F1	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F2	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F3	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F4	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F6	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F7	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F8	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F9	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F10	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F11	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F12	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F13	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F14	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
F15	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

TABLE I: Comparing features supported by existing client-side analysis tools. Below, we provide a short description for each of these features:

- F1: Client-state-aware** refers to crawlers that interact with the client-side user interface and explore its different possible states at runtime.
- F2: Server-state-aware** refers to crawlers that aim to trigger different states at the server side.
- F3 & F4: Data input & event sequence generation** shows whether the analyzer generates data inputs and event sequences, respectively.
- F5: Dynamic event handlers registration** shows whether the analyzer supports event handlers that are registered dynamically.
- F6: Static link extraction** refers to the crawlers that can statically find hyperlinks embedded into HTML pages.
- F7: Prioritization** shows whether a crawler can prioritize the triggering of certain events to achieve a specific goal, such as improving coverage.
- F8: Filling forms** is required in some JavaScript applications, where users enter data via forms. This feature shows whether the crawler is able to fill such forms.
- F9: Bypassing guards in the code** is required to explore applications deeply and reach specific locations in the program. String constraint solving is a popular technique used to generate such inputs.
- F10: Handling nondeterminism** is required when the state-aware crawler cannot go to a previously visited state, or gets trapped in one state and cannot make progress due to usage of time stamps, randomization and being dependent on the existence of certain behaviors that change over time.
- F11: Security analysis** refers to having support for client-side vulnerability detection techniques, such as DOM-based XSS detection.
- F12: Guiding towards sinks** is a feature that allows a crawler to drive a security analysis more efficiently and effectively.
- F13: Triggering specific functionalities** through a sequence of events is required by some client-side applications that provide complex functionalities.
- F14: Improving coverage** is often the goal of all the crawlers. Coverage usually is measured in terms of the number of detected hyperlinks or the amount of executed JavaScript code.
- F15: Support for modern libraries and frameworks** is required for most of the modern JavaScript applications that are built on top of complex frameworks, such as React [?], Knockout.js [19].

Given the large search space, our key insight is that a crawler needs to be guided to accelerate a target dynamic analysis to solve a problem. We propose a hybrid analysis of JavaScript applications that is feedback-driven to direct crawlers towards reaching program locations of interest, such as DOM-based XSS sinks and REST calls.

Once the crawler identifies endpoints and drives the execution to reach them, a security analysis runs to report security vulnerabilities. In this work, we focus on DOM-based XSS and reflected XSS vulnerabilities. The DOM-based XSS analysis, in particular, requires a practical dynamic taint analysis to work on real-world applications. Being able to analyze modern real-world applications as our main objective, we propose a novel staged taint inference analysis to detect DOM-based XSS vulnerabilities. Compared to the state-of-the-art dynamic taint analysis tools [36], [49], [53], [34], [42], [43], our solution has better recall and is less intrusive, which makes our analysis less likely to break the semantics of the applications.

The rest of the paper is organized as follows: Sec. II summarizes the closely related works and how they compare against GELATO. Sec. III describes the design of our feedback-driven and guided crawler, and the input value generation to bypass guards in the program. Sec. IV explains the staged taint inference analysis used to detect DOM-based XSS vulnerabilities. Sec. V provides details on the implementation of GELATO and Sec. VI explains the experimental setup, evaluation results on various benchmarks and applications, and comparison against state-of-the-art tools. Finally, Sec. VII concludes the paper. In summary, we make the following contributions:

- a new crawler that can be guided towards program locations of interest using a call graph.
- a feedback-driven analysis that enables our guided crawler to support modern client-side JavaScript libraries and frameworks instead of using manually crafted models.
- a novel staged taint inference analysis that detects potential DOM-based XSS vulnerabilities with high accuracy.
- an input generator that supports both event and data value generation to increase the coverage of security analyses.

## II. RELATED WORK

Client-side web application analysis has a large body of literature, and over the past decade, many crawling and security analysis techniques have been developed. In this section, we compare GELATO against the related works in terms of the crawling technology and taint analysis.

### A. Client-side web application crawling

Crawljax [44] is a state-aware crawler that explores AJAX-based client-side applications using dynamic analysis. It computes the edit distance of the string representation of DOM trees to compare states, and performs both depth-first and breadth-first search strategies. Compared to Crawljax, our approach is targeted and aims to increase coverage for a set of

program locations of interest. We combine static and dynamic analysis and also generate data values as inputs to guide the runtime execution.

jÄk [50] is designed to analyze modern web applications to find server-side security vulnerabilities. The main goal of this work is to increase code coverage and trigger all interaction points in the client-side program to be able to find more vulnerabilities on the server side. For this purpose, it uses dynamic analysis of the client-side JavaScript program to detect dynamically generated URLs, the registration of events, etc. The dynamic analysis is combined with crawling to interact with the application and infer a navigational model.

Similar to jÄk, we perform dynamic analysis to collect runtime values, traces and to capture events that are difficult to detect statically. Additionally, we use static analysis to guide the crawler to specific locations in the program to avoid the state explosion problem. Moreover, we integrate input value generation to our crawling technique to increase the coverage of relevant parts of the program.

FEEDEX [46] uses a state-aware feedback-directed crawling technique to derive a test model for client-side web applications. The main focus of FEEDEX is to reduce the test model size and enhance coverage in three aspects of a test model: functionality, navigation and page structure. As FEEDEX crawls the web application, the coverage of these three aspects is fed back to the tool to prioritize next states and events. Compared to FEEDEX, our technique focuses on guiding the crawler towards specific locations in the program. We guide the analysis by combining static analysis, link extraction, state-aware crawling and input value generation. To analyze modern web applications that heavily use complex libraries for which static analysis is difficult, our feedback-directed analysis refines the static analysis results using the runtime execution. This novel design allows us to use the crawler to improve coverage of security analyses, such as DOM-based XSS detection and black-box REST fuzzing.

WATEG [55] takes a rule-directed test-case generation approach that focuses on increasing the coverage of business rules (functionalities) provided as specifications. As the program executes, WATEG checks whether the invariants derived from business rules are not violated. In this work, the crawler is directed towards pre-determined functionalities (business rules) using a two-phase approach: (1) uses coarse state comparison to create STD (abstract state transition diagram) and to determine the portions of the state space that are relevant; (2) uses a more fine-grained state comparison for the relevant parts. The abstract paths of STD are used in the second phase as starting points and refined to traversable paths that lead to triggering the pre-determined functionalities.

Even though our technique also guides the execution of the program, its goal is different. Unlike WATEG, our technique guides the execution towards specific locations in the JavaScript code or events and not to cover pre-defined business rules. Therefore, we analyze the JavaScript code and use program analysis techniques to prioritize the event and state prioritization.

In what follows, we provide an overview of other related works (not closely related) in the literature and compare against our technique.

1) *Scoped and guided crawlers*: Scoped crawling [48] of client-side web applications proposes strategies to limit the scope of the exploration based on textual content, such as topic, geography, language, etc. However, our approach uses program analysis techniques to guide the execution towards specific program locations.

Guided crawling aims to guide the exploration to achieve a particular goal, such as increasing code, functionality or navigation coverage [31], [46], [55]. Our crawling technique is also guided because it uses prioritization strategies to guide the execution. However, it aims to increase coverage of specific program locations to drive a target analysis.

[32] guides the exploration of the application, aiming to discover as many states as possible in a given amount of time. Compared to [31], [46], [55], which focus on increasing the diversity of crawled pages, this work mainly focuses on increasing efficiency for an anticipated model. Our crawling technique also improves efficiency for reaching program locations of interest without requiring a crawling model, and uses different analysis techniques to make it targeted.

Recently, guided fuzzing techniques have been proposed for C/C++ programs [33], [51]. Similar to our approach, AFLGo [33] uses call graph distance from target locations in the program as a metric to prioritize input generation. However, our approach goes one step further and also allows on-the-fly refinement of the statically constructed call graph to add missing and remove false positive edges in the call graph. Moreover, we analyze JavaScript applications that have a highly dynamic nature and require support for both events and input values.

2) *Traditional crawlers*: Many of the traditional web application crawlers available in the industry are not state-aware, and rely only on static link extraction over the HTML pages [3], [24], [28]. To be practically useful for vulnerability detection, security analysts need to manually interact with the application to take the browser to the desired states. However, our technique is fully automatic, and combines static and dynamic analysis of JavaScript code to guide the crawler towards target (security-sensitive) locations in the program.

3) *Data input generation*: In addition to event-based inputs (e.g., clicking), client-side web applications also accept input value in URLs and form elements such as input fields. Therefore, to explore a client-side application deeply enough, input value generation is required. For input fields, existing approaches mostly provide random data if no custom data is available [45]. More heavyweight analyses such as symbolic execution have also been proposed [52], [54], which are known to be prone to scalability issues. In contrast, we propose a lightweight data input generation technique based on taint inference that starts with random or custom data, but generates new inputs to bypass validation routines and reach target locations (sinks) in the program.

Autogram is a recent work that explores input generation by mining grammars [38]. It uses dynamic taint tracking to trace the data flow of each input character for a set of sample inputs. By grouping input fragments that are handled by the same functions, it produces a context-free grammar that can be combined with fuzzers to generate inputs. In contrast, our input generation assumes that the input grammar is known

(e.g., URL) and tries to bypass the validation routines using taint inference.

## B. Taint Analysis

In this section, we compare our approach against existing taint analysis techniques used for finding vulnerabilities in web applications.

Several static analysis techniques have also been proposed to analyze JavaScript applications [40], [37], [41]. However, the lack of static predictability and presence of dynamic typing in JavaScript as well as the asynchronous and event-based nature of web applications can be highly problematic for determining taint flows using static analysis techniques.

Dynamic analysis of JavaScript programs requires instrumentation. Two types of instrumentation techniques are used for dynamic analysis of web applications: engine-level instrumentation [42], [56], [43], and code-rewriting [53], [49], [17], [36]. Engine-level instrumentation involves adding hooks to the JavaScript engine. While this design can have performance benefits from being compiled into the engine itself, it is not portable across different engines and requires considerable effort to maintain [42].

On the other hand, source code-level instrumentation [14], [34] involves replacing and/or appending code to the existing program's source-code so that the runtime behavior can be analyzed without affecting the original behavior. This method has drawbacks in performance, but is often easier to write and test, and can be engine-agnostic.

Dynamic taint analysis involves tracking taint labels in a program during its execution. DexterJS [49] carries out character-level taint tracking using code-rewriting to discover potentially vulnerable taint flows. This approach involves tracking each character originating from a taint source individually. Even a single character reaching a sink can be detected, provided that it is propagated from the tainted value. DexterJS attaches taint labels to primitive values by wrapping (boxing) them. Because built-ins, browser APIs and DOM functions cannot be instrumented, DexterJS requires hard-coded models. However, coming up with models that capture all runtime behavior is very challenging. In fact, our experiments with DexterJS [49] show that incomplete models for built-ins result in missing valid taint flows.

Linvail [34] is a dynamic shadow execution framework based on source-code-level instrumentation that can be used to implement dynamic taint tracking. To make sure the analyzed program is not affected by wrapped values, Linvail permanently unwraps and wraps values around calls and uses JavaScript proxies to intercept object accesses from non-instrumented code. However, in order to handle side-effects and keep track of taint labels in non-instrumented code, it relies on an oracle of hard-coded models. Providing a complete oracle that preserves the semantics of the program is known to be very challenging.

Jalangi1 [53] (not maintained anymore) is another dynamic analysis framework based on source code-level instrumentation that provides shadow execution to run different types of analysis, e.g., dynamic taint tracking. To support primitive values in the shadow execution, it has hard-coded wrapping

and unwrapping operations for language-level operations, e.g., assignments. For external code that is not instrumented, its partial solution handles built-in calls that expect primitive values but receive wrapped values instead. However, it might break the semantics of the program when wrapped objects reach non-instrumented code. Also, during the offline mode in Jalangi1, external calls are replaced with concrete values, thereby ignoring side-effects and potentially deviating from the dynamic execution. Unlike Jalangi1, Jalangi2 [14] only provides syntactic traps to implement dynamic analyses to avoid the existing problems in Jalangi1.

Compared to the dynamic taint tracking approaches, our taint inference analysis is based on source code-level instrumentation. We use Jalangi2 to instrument JavaScript programs with our analysis code because it is easy to maintain and can work across different JavaScript engines. Our analysis is lightweight and can deal with the non-instrumented parts: built-ins, browser APIs and DOM functions.

Affogato [36] is an instrumentation-based dynamic taint inference analysis tool for Node.js applications. Similar to GELATO, it finds injection vulnerabilities by detecting flows of data from untrusted security-sensitive sources to sinks at runtime using a non-intrusive grey-box taint inference analysis. GELATO goes one step further and improves the precision by introducing a multi-staged approach.

### III. FEEDBACK-DRIVEN, GUIDED, AND SECURITY-AWARE CRAWLING OF MODERN WEB APPLICATIONS

Most of the existing crawlers are designed as standalone drivers to interact with an application with the goal of getting maximum coverage of the executed code or discovered hyperlinks. However, the search space in modern web applications is simply too large for such a coarse-grained approach to be useful in practice. As a result, based on a target analysis, the security analysts often need to manually identify which parts of the application should be prioritized and explored more deeply to find security vulnerabilities.

In this work, we take the first step to automate guiding a state-aware crawler based on the requirements of a target security analysis using a feedback-driven approach. Algorithm 1 shows the main feedback loop of our crawler. This algorithm takes as input the *URL* of a web application, a target security analysis, *TA*, and a set of program locations, *Loc*, that the crawler should be guided towards. The loop continues until the crawler reaches a fixpoint and there are no more new states to visit, and the results of the target analysis (*Results*) is reported as output.

At each iteration, we run a target security analysis, *TA*, over the JavaScript and HTML code in the given state, *S*, and collect results. At the same time, we run an approximate call graph analysis [35] on the newly discovered JavaScript code, and collect the execution trace using lightweight instrumentation. The execution trace helps us to determine the actual (true positive) function calls, thereby removing false positive call graph edges, and adding newly discovered ones to the *ACG*, which contains the call graph for the explored parts of the application. *PrioritizeEvent(S, ACG)* determines which event should be triggered next based on the metrics described in Sec. III-D that are computed using *ACG* and the current

state *S*. In Sec. IV, we describe a novel taint inference analysis to detect DOM-based XSS vulnerabilities as an example security analysis that fits well in our approach.

---

#### Algorithm 1 Feedback-driven and guided security-aware crawler

---

```

1: inputs: web application URL, target analysis TA, target
   program locations Loc
2: output: Results
3: ACG  $\leftarrow \emptyset$ 
4: browser.goto(URL)
5: while browser.newStateExists() do
6:   S  $\leftarrow$  browser.getNewState()
7:   Results  $\leftarrow$  Results  $\cup$  ANALYZE(S, TA) // See Algo-
   rithm 2
8:   cg  $\leftarrow$  computeACG(S)
9:   trace  $\leftarrow$  getExecutionTrace(S)
10:  ACG  $\leftarrow$  refineACG(ACG, cg, trace)
11:  e  $\leftarrow$  prioritizeEvent(S, ACG)
12:  browser.goto(e)
13: end while
14: report(Results)

```

---

#### A. State representation and comparison

The first key challenge in designing a state-aware crawler is to come up with a suitable state representation. Ideally, we would like to store the entire browser and server state to be able to switch back and forth between the states gracefully. However, keeping such a huge amount of information in each state is unrealistic. To address this challenge, we try to keep the size of the states minimal by storing only: (1) URL; and (2) DOM tree, which we have found to be the most crucial elements. We also record references to the parent and child states to be able to replay the sequence of events that have been triggered to reach the current state.

We determine whether a state has been visited before using the following heuristics: (1) the path segment of the URL is the same; (2) the size of the DOM tree has not changed dramatically (less than a threshold); (3) the hash computed for the DOM tree is exactly the same, or the difference in DOM tree structure is less than a threshold.

#### B. Search strategy

We perform Depth First Search (DFS) on the crawler state graph to explore the states. However, because we do not keep the entire browser state, it is not possible to directly backtrack and take the browser to a previously visited state. One way to go back to a visited state is to replay event sequences all the way from the root to reach a particular state. This approach requires triggering many unnecessary events that can significantly affect the efficiency of the crawler.

To deal with this challenge, we apply heuristics to take the browser to a previously visited state: (1) check whether the target state can be reached by triggering another event from the current state; and (2) compute the shortest path in the state graph to reach the target state.

```

1 <!DOCTYPE html>
2 <html lang='en'>
3 <script src='./knockout-min.js'></script>
4 <script>
5 function event_handler() {
6     fetch('');
7 }
8 </script>
9 <body>
10 <button id='knockout_data_bind_1' data-bind='click:
    ↳ button_click'>knockout_data_bind</button>
11 <script>
12 function KnockoutViewModel() {
13     this.button_click = function() {
14         event_handler();
15     };
16     this.button_text = 'knockout_data_bind';
17 }
18 ko.applyBindings(new KnockoutViewModel());
19 </script>
20 </body>
21 </html>

```

Listing 1: An example framework code that is hard for ACG to analyze soundly, and misses a critical call graph edge. Our call graph refinement approach, however, is able to detect and add it to the call graph.

### C. Call graph refinement

*ComputeACG* in Algorithm 1 generates an approximate call graph using ACG [35] for the JavaScript code executed in the state  $S$ . Note that the call graph construction is initially performed statically. As the crawler interacts with the user interface, we collect the function calls as part of the execution trace. This execution trace is next processed to examine whether an edge in the call graph is missing or is a false positive. The call graph is updated with this new information. An edge from node  $a$  to node  $b$  in the call graph is considered as a false positive if visiting  $a$  does not result in visiting  $b$ . And an edge from node  $a$  to node  $b$  is missing if the call graph does not include such an edge but the execution trace does. Listing 1 is a code-snippet from the knockout.js [19] framework. Due to a complex event delegation mechanism, ACG fails to find the edge from the `click` event in the button element to `event_handler` function. However, once the crawler clicks on this button, the execution trace records that `event_handler` gets triggered. This newly found edge is added to the call graph, *ACG*.

### D. Prioritization

We prioritize a state that is visited for the first time (not similar to any of the previously visited states) if it contains a target program location (*Loc* in Algorithm 1). For partially expanded states, we use a prioritization heuristic to choose the next event in the state that should be triggered by the browser. To guide the crawler towards target program locations, we prioritize an event if it has the minimum distance from the handler (registered to handle it) to a target location in the call graph.

### E. Input value generation

While the state-aware crawler interacts with the JavaScript application, we analyze (line 7 in Algorithm 1) the JavaScript code returned from the server side<sup>1</sup>, generating input values to bypass guards and increase coverage. There are several ways to

provide input values into a client-side JavaScript application: forms, URLs, cookies, local storage, etc. In this section, we show how we integrate input value generation for URLs to our state-aware crawler for simplicity. However, the same algorithm can be used for other sources of input values.

The input value generation is performed only on the states that are candidates to be analyzed by the target analysis (*TA* in Algorithm 1). A state is analyzed if its code contains a target program location. Algorithm 2 shows how we generate input values.

At high level, to bypass the guards on the execution path we collect runtime values of interest during the execution, construct path constraints, and solve them to generate inputs (*generateNewURLs* at line 19). The guards that we aim to bypass are validation routines that must be satisfied to let the analyzer reach the deeper parts of the program. Example runtime values of interest are operands in conditional statements (e.g., `if` statement) and arguments in string function calls (e.g., the `string.substring` built-in function) that are triggered on the execution path. We use such logged values in constraint generation if they are tainted (the taint analysis is explained in detail in Sec. IV). These constraints are used to replace tainted characters of input values that are compared in a conditional statement.

```

1 <!DOCTYPE html>
2 <html>
3 <body>
4 <button id="button1" onclick=safe()> Button 1 </button>
5 <button id="button2" onclick=unsafe()> Button 2 </button>
6 <textarea id='text1' style="display:none;">This is
    ↳ safe! </textarea>
7 <script>
8     function safe(){
9         document.getElementById("text1").style =
    ↳ "display:true;";
10     }
11     function unsafe(){
12         var loc = document.location.hash //url is
    ↳ "http://example.com#action"
13         if (loc.indexOf("show")!=-1) {
14             document.write("You are visiting site: " +
    ↳ document.location.href + ".");
15         }
16     }
17 </script>
18 </body>
19 </html>

```

Listing 2: An example of HTML/JavaScript code with constraints on input value.

We explain the input value generation in Algorithm 2 through an example. Listing 2 shows a simplified application that is vulnerable to DOM-based XSS attack. In a DOM-based XSS attack, the attacker-controllable input (e.g., URL) is loaded in the victim's browser and the injected payload flows to a DOM manipulation statement that results in running the malicious payload, such as stealing cookies, under the victim's session. In this example, a value obtained from the URL at line 12 is written to the `document` object at line 14, which modifies the DOM and allows to run the attackers malicious payload. However, the original URL used to load the page is "`http://example.com#action`", which does not contain "`show`". Therefore, line 14 is not executed when the original input value is used. Next, we show how we generate an input (URL) that bypasses the validation at line 13 and allows the execution to reach line 14.

<sup>1</sup>The JavaScript code is stored in the crawler state.

Initially, the execution path ( $\pi$ ) in Algorithm 2 is empty and the test input queue, *InputQ*, contains the original URL. Our input generator continues generating new test inputs until *InputQ* is empty. In each iteration, an input is removed from *InputQ* and passed to the *runTargetAnalysis* function, which runs the target analysis (*TA*) determined by the analyst. Before running the target analysis, we take the browser to state *S* by obtaining and triggering the corresponding event sequence (*eventSeq(S)*).

As the target analysis is performed at line 7, the conditional statements (e.g., *if* statements) are logged in  $\pi$ , which are used to generate path constraints and new test inputs at line 8. Going back to the example in Listing 2, when analysis executes line 13, we record the *if* statement together with the following runtime values in  $\pi$ : "show", -1 and "http://example.com#action" (value of *loc* variable).

The *ValueInputGen* function in Algorithm 2 generates new inputs using the values recorded in the execution path,  $\pi$ . If a value in a conditional statement or *string* function call is identified to be tainted by taint analysis (See Sec. IV for details), the tainted characters and the value that they are compared against are recorded in the *Constraints* map. For instance, the value of *loc* at line 13 in Listing 2 is inferred to be tainted by taint analysis, and the tainted characters are "action". Therefore, "action" is added to *Constraints*[3].*taintedVal* at line 15 in Algorithm 2. We also record "show", which is the value that the tainted value is compared against.

Finally, the *genConstraint* function at line 16 in the algorithm generates the *loc* == "show" constraint and stores it in the *Constraints* map. The *generateNewURLs* function at line 19 replaces "action" with "show" in the original URL<sup>2</sup> and generates a new input<sup>3</sup>. If the target analysis is DOM-based XSS detection, once the new URL is loaded and analyzed at line 6 in Algorithm 2, line 14 in Listing 2 is executed and a DOM-based XSS vulnerability is reported.

#### F. Support for other features

In this section, we elaborate on the remaining features mentioned in Table I that are supported by GELATO.

**Dynamic event handler registration.** We perform dynamic analysis to collect runtime values, traces and to capture events that are difficult to detect statically. To perform the dynamic analysis, we instrument the JavaScript code to collect traces and hook into the event handlers to capture the events that otherwise cannot be found at runtime.<sup>4</sup> There are also many cases where the event registration is done in a library through a complex mechanism that is not straightforward (e.g., the event delegation mechanism in *jQuery* [15]). In such cases, we use models for the common libraries and frameworks that allow us to extract the required information from their internal data storage.

**Static link extraction.** Similar to most of the existing crawlers, we statically extract links from the HTML pages. In addition, we integrate static link extraction to the state prioritization,

---

#### Algorithm 2 Input value generation

---

```

1: function ANALYZE(S, TA)
2:    $\pi \leftarrow \emptyset$  // JavaScript execution path
3:   InputQ  $\leftarrow$  URL //initial seed input value
4:   while notInputQ.isEmpty() do
5:     v  $\leftarrow$  InputQ.pop()
6:     browser.Goto(eventSeq(S))
7:      $\pi \leftarrow$  runTargetAnalysis(v, TA)
8:     InputQ.add(VALUEINPUTGEN( $\pi$ ))
9:   end while
10: end function
11: function VALUEINPUTGEN( $\pi$ )
12:   Constraints  $\leftarrow \emptyset$ 
13:   for n in  $\pi$  do
14:     if taintAnalysis(n.val) then
15:       Constraints[n.loc].taintedVal = taintedVal(n.val)
16:       Constraints[n.loc].cons = genConstraint(n)
17:     end if
18:   end for
19:   Return generateNewURLs(Constraints, URL)
20: end function

```

---

as follows. Before the state-aware crawling starts, GELATO statically extracts links from HTML pages starting from *URL*. This step allows us to retrieve a partial and coarse-grained structure of the web application. We also build call graphs statically for the JavaScript code of the extracted pages and prioritize them for the state-aware crawling if they contain user-specified target locations (*Loc* in algorithm 1).

**Filling forms.** GELATO fills forms with payloads provided as configuration. Furthermore, we use the input value generation technique described in Sec. III-E to fill form fields.

## IV. TARGET SECURITY ANALYSIS: DOM-BASED XSS DETECTION

In this section, we describe a novel DOM-based XSS detection technique as an example security analysis that can be integrated into our guided crawler to detect vulnerabilities in real-world JavaScript applications. For DOM-based XSS analysis, the crawler needs to be guided towards DOM manipulation locations, which are marked as sinks. Once the crawler reaches the states that contain such sinks, we perform taint analysis as described below to detect vulnerabilities.

Dynamic taint analysis is a common technique to detect injection vulnerabilities in applications written in dynamic languages. In practice, however, several technical challenges must be addressed to implement a dynamic taint analysis that will be able to analyze real-world applications. In this section, we first elaborate on the technical challenges that must be addressed to implement a dynamic taint analysis for JavaScript. Then we present our dynamic taint inference approach that overcomes or circumvents these challenges.

Dynamic analysis for JavaScript has a short but rich history [30] and we can already extract valuable lessons from existing work. This type of analysis requires instrumentation

---

<sup>2</sup>http://example.com#action

<sup>3</sup>http://example.com#show

<sup>4</sup>These are the events that are registered using *addEventListener*

either at JavaScript engine [42], [56] or source-code level [53], [49], [17] to be carried out. An engine-level instrumentation-based analysis would require substantial effort in order to support multiple engines and to be maintained in the long-term. Being engine-agnostic is important for security analysis because it is possible for an attack to work on one engine but not on another one [49]. The dynamic taint analysis should find taint flows regardless of the engine it is running on.

Previous works show that source-code level instrumentation-based dynamic taint analysis for JavaScript will face the following challenges:

- 1) Tracking taint through non-instrumented code.
- 2) Attaching taint labels to primitive values.

Because modern JavaScript engines are typically implemented in low-level languages (C, C++), an instrumentation-based dynamic analysis will not be able to instrument all built-in functions (e.g., array modification, and string operations). Furthermore, for efficiency reasons, it is often desirable not to instrument an entire application to leave some modules uninstrumented. Consequently, the analysis needs to model what happens in non-instrumented code and update taint labels accordingly. A common approach to deal with non-instrumented code is to use manually created models [34], [49], [17]. Unfortunately, our experience with these tools suggests that their models contain bugs and are incomplete.

Finally, because primitives cannot be extended with additional properties, instrumentation-based approaches need to *wrap* primitives in an object that will have a property representing the taint label of the primitive value – aka, boxing. Because wrapping primitives is an intrusive process that alters the execution of the original program, care must be taken not to alter its semantics. However, empirical evidence from previous work [34], [49], [53] suggests that wrapping primitives while preserving the semantics of the original program is extremely challenging. In general, wrapped primitives must be unwrapped before they exit the instrumented code, and wrapped before they enter the instrumented code. We suspect that because of these difficulties primitive wrapping has been dropped in the new version of Jalangi (Jalangi2) [14].

To highlight the challenges encountered by existing dynamic taint tracking solutions, consider the code-snippet in Listing 3. This example shows a JavaScript program that contains a vulnerable taint flow from the source, `location.hash` at line 1, to the security sensitive sink, `document.write` at line 5. In this example, the string value `#payload` is injected in the URL as the fragment identifier (i.e., the part of the URL following the `#` sign). DexterJS [49] fails to report the vulnerable taint flow due to an incorrect model used for the uninstrumented built-in function, `substring` at line 2. On the other hand, Linvail [34] and Chromium Taint Tracking [42] report two taint flows at both lines 4 and 5, even though `document.write` at line 4 is not tainted.

### A. Dynamic Taint Inference

To address the challenges listed above, we developed a non-intrusive, dynamic taint inference analysis based on source code-level instrumentation. Our analysis infers tainted flows by correlating values at sources and sinks, and observing the

```

1 var tmp = document.location.hash; // tmp = "#payload"
2 tmp = tmp.substring(3,7); // tmp = "yloa"
3 tmp = tmp + "123"; // tmp = "yloa123"
4 document.write(tmp.substring(5)) // "23" is written to
  ↳ DOM
5 document.write(tmp); // "yloa123" is written to
  ↳ DOM

```

Listing 3: Example JavaScript program that contains a vulnerable taint flow.

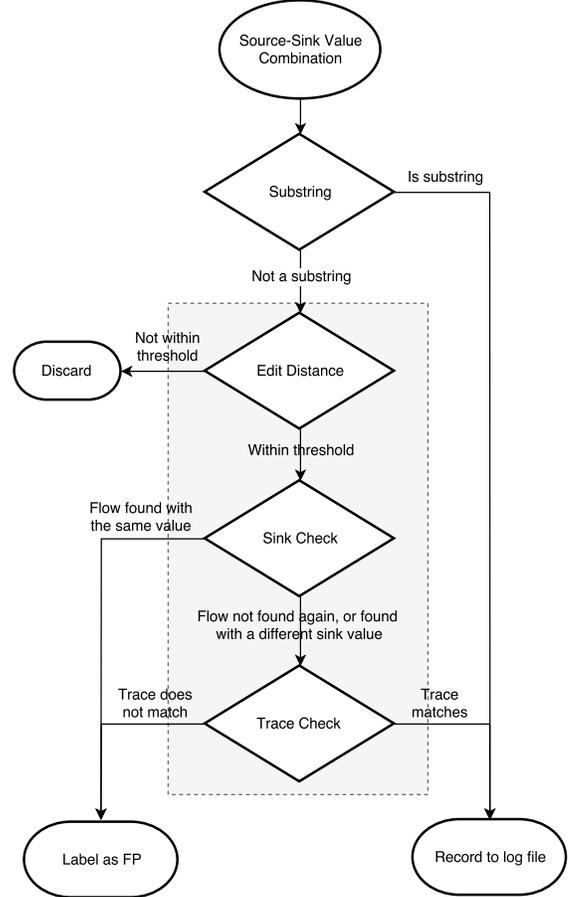


Fig. 1: Flow diagram of our staged dynamic taint flow inference technique.

behavior of the program instead of attaching and tracking taint labels.

Fig. 1 shows our staged approach to discover correlations between values at taint sources and sinks. The stages, represented as diamonds in Fig. 1, act as increasingly complex filtering steps that aim to maximise the precision of our approach.

1) *Stage 1: Substring*: The first stage looks for an exact substring match of length  $\geq \theta$  between string values observed at sources and sinks, i.e., whether either string is a substring of the other. If a match of length  $\geq \theta$  is found, a taint flow is immediately reported. Otherwise, the remaining three stages (shown in the grey box) are used to infer taint flows when the values at sources and sinks approximately match. The remaining three stages will be described in the following sections using the symbols outlined below.

- $A$  = a source (identified by location in source-code)

- $B$  = a sink (identified by location in source-code)
- $A_v$  = string value at source  $A$
- $B_v$  = string value at sink  $B$
- $F$  = a taint flow detected by our analysis

2) *Stage 2: Edit Distance:* If neither  $A_v$  or  $B_v$  is a substring of the other (i.e., the substring stage does not report a match), the edit distance filter performs approximate matching of  $A_v$  and  $B_v$ . Specifically, this stage computes the *longest common subsequence* (LCS) [47] between  $A_v$  and  $B_v$ , extracts  $D_i$  and  $D_d$ , the number of insertions and deletions required to compute the LCS, computes a similarity score, and compares it to a threshold  $\eta$ , as shown in Algorithm 3. The source-sink ( $A, B$ ) pairs that pass this test in this stage are recorded and processed in the next stages to filter out false positives (FPs).

---

#### Algorithm 3 Edit Distance Stage

---

```

1: function EDITDISTANCE( $A_v, B_v, \eta$ )
2:   Let  $D_i, D_d = LCS(A_v, B_v)$ 
3:   Let  $L = \max(\text{len}(A_v), \text{len}(B_v))$ 
4:   if  $\frac{L - (D_i + D_d)}{L} \geq \eta$  then
5:
6:     return "Yes"
7:   end if
8:
9:   return "No"
10: end function

```

---

3) *Stage 3: Sink Check:* According to the flow diagram in Fig. 1, when a source-sink pair ( $A, B$ ) reaches the sink check stage, we know that there is no exact substring match but that the two strings are *similar*. To weed out cases where the similarity happens by chance (i.e., there is no taint flow from  $A$  to  $B$ ), the sink check stage mutates  $A_v$  into  $A'_v$  by changing a few characters randomly, running the program again with the new source input, and observing  $B'_v$ . There are three possible outcomes, as shown in Algorithm 4:

- 1) Sink  $B$  is not reached ( $B'_v$  is NULL). The execution path triggered by  $A'_v$  has diverged from the execution path triggered by  $A_v$ . The pair ( $A, B$ ) proceeds to the next stage.
- 2)  $B'_v$  is different from  $B_v$ , indicating that the value at  $A$  has an impact on the value at  $B$ . The pair ( $A, B$ ) proceeds to the next stage.
- 3)  $B'_v$  is identical to  $B_v$ , indicating that the value at  $A$  probably has no impact on the value at  $B$ . The pair ( $A, B$ ) does not proceed to the next stage.

---

#### Algorithm 4 Sink Check Stage

---

```

1: function SINKCHECK( $(A, B)$ )
2:   Let  $A'_v = \text{mutate}(A_v)$ 
3:    $B'_v = \text{runApplication}(A'_v)$ 
4:   if  $B'_v$  is NULL then
5:
6:     return "Proceed to next filter"
7:   else if  $B'_v \neq B_v$  then
8:
9:     return "Proceed to next filter"
10:  else if  $B'_v == B_v$  then
11:
12:    return "No taint flow from  $A$  to  $B$ "
13:  end if
14: end function

```

---

4) *Stage 4: Trace Check:* Trace Check is the final and most expensive stage of our taint flow inference process. It aims at detecting real taint flows with high precision. This step involves recording the JavaScript execution trace and analyzing the string manipulation operations performed on  $A_v$  to determine whether  $B_v$  is derived from  $A_v$  (i.e., there is a taint flow from  $A$  to  $B$ ).

Algorithm 5 shows our trace check stage. Given a source value  $A_v$ , a number of insertions  $D_i$ , a number of deletions  $D_d$ , and the execution trace seeded with  $A_v$ , the TRACECHECK procedure determines whether the string operations in the trace can possibly transform  $A_v$  into  $B_v$ . The sub-procedure ISOPTAINTED in Algorithm 5 re-uses the Substring and Edit Distance stages, parameterised with  $\theta$  and  $\eta$ , to determine whether the base variable or any argument of a string operation matches  $A_v$ . If the base variable or any argument matches  $A_v$ , ISOPTAINTED returns true.

---

#### Algorithm 5 Trace Check Stage

---

```

1: function TRACECHECK( $A_v, D_i, D_d, \text{trace}$ )
2:   Let  $D_{ti} = 0$  and  $D_{td} = 0$ 
3:   for each  $\text{string}_{op}$  in  $\text{trace}$  do
4:     if ISOPTAINTED( $\text{string}_{op}, \theta, \eta, A_v$ ) then
5:       if ISOPINSERTION( $\text{string}_{op}$ ) then
6:          $D_{ti} += 1$ 
7:       else if ISOPDELETION( $\text{string}_{op}$ ) then
8:          $D_{td} += 1$ 
9:       end if
10:    end if
11:  end for
12:  if ( $D_i > 0 \ \&\& \ D_{ti} == 0$ ) || ( $D_d > 0 \ \&\& \ D_{td} == 0$ )
13:  then
14:    return "Trace does not match"
15:  else
16:
17:    return "Trace matches"
18:  end if
19: end function

```

---

The trace check stage counts the number of tainted string operations that are insertions ( $D_{ti}$ ), and deletions ( $D_{td}$ ) in the trace. Then, it weeds out traces where either no insertion happens while  $D_i > 0$  or no deletion happens while  $D_d > 0$ .

We now revisit the example in Listing 3 to show how our taint inference technique correctly reports an inferred taint flow at line 5 and does not report any flows at line 4. For the sink at line 4, the observed values are  $A_v = \text{"\#payload"}$  and  $B_v = \text{"23"}$ . Since  $A_v$  and  $B_v$  do not pass the check at Substring stage, they are passed to the Edit Distance stage. As these values also fail to pass the Edit Distance check, the analysis does not infer any taint flows.

For the sink at line 5 in this example,  $A_v = \text{"\#payload"}$ ,  $B_v = \text{"yloa123"}$ ,  $D_i = 3$ ,  $D_d = 4$ . Inspecting the trace between line 1 and 4 in this example, the Trace Checking filter is able to determine that one string concatenation and two substring operations occurred. Because both of these operations are performed on the `tmp` base variable with the string values `"\#payload"`, `"yloa"` and `"yloa123"`, they pass the ISOPTAINTED check at line 4 in Algorithm 5 that compares them against the source value `\#payload`. The trace check algorithm then computes  $D_{ti} = 1$ , and  $D_{td} = 2$  and concludes that the trace matches at line 15 in Algorithm 5.

## V. IMPLEMENTATION

We implemented our security-aware guided crawler in a tool called GELATO. To guide our input generator towards target locations, we use the pessimistic mode in the approximate call graph construction [35] to statically build call graphs. We have developed a new lightweight instrumentor to carry out the dynamic analysis of JavaScript code for dynamic event handler registration and call graph refinement. The dynamic analysis finds dynamically registered events, new pages and missing edges in the call graph at runtime. The lightweight instrumentor is also used to collect runtime values and generate constraints for creating new input values.

The source-sink identification, Substring and Edit Distance checkers in the taint flow inference system are implemented as an analysis written in JavaScript on top of the Jalangi2 [14] analysis framework. The Response Check and Trace Check components are written in Python, using WebSocket to communicate with the JavaScript part of the system. The taint flow reporting, crawler and event generator are implemented in Python using bindings for Puppeteer [25], a browser automation framework.

The taint flow inference framework requires a list of sources and sinks to begin with. The same sinks are used to guide the crawler for DOM-based XSS detection. These JavaScript methods and property read/write statements are intercepted using Jalangi to record and analyze their values.

## VI. EVALUATION

The experiments are performed in Google Chrome browser version 55.0.2883.87 and ChromeDriver version 2.27.440175 on Ubuntu 16.04 running on VirtualBox 6.0, Intel i7-7700 CPU @ 3.60GHz x 4 (4 cores assigned to VM) with 4096 MB memory. To evaluate our feedback-driven edge addition technique, we have manually created models for complex libraries, jQuery [15], Knockout.js [19], and React.js [26] that help find missing edges in the call graph to improve coverage of target locations. We compare the results of our edge addition technique with these manually created models. In our experiments we answer the following research questions:

- RQ1: How effective is GELATO’s feedback-driven and guided crawling technique in terms of coverage and performance?
- RQ2: How does GELATO compare against other tools for the number of discovered URLs?
- RQ3: How effective is GELATO’s DOM-based XSS detection technique compared to the state-of-the-art taint analysis techniques in terms of precision and recall?

### A. Choosing benchmarks

One of the challenges we have faced for evaluating GELATO is choosing benchmarks. Existing taint analysis works often evaluate on Alexa top [1] websites, however, we are constrained to fuzz and analyze websites and applications that are published as open-source projects for comparing analysis tools. Therefore, we have gathered deliberately vulnerable open-source applications, vulnerable libraries, our in-house applications, and micro-benchmarks. Our criteria for choosing benchmarks are to:

- be realistic, diverse, and use modern technologies
- use complex libraries and show the capability to detect known CVEs
- include both single-page and multi-page applications
- show the accuracy of the analysis by evaluating on relevant micro-benchmarks
- be tested by the related works if possible

We evaluate GELATO on two target analyses, for which we have collected different benchmarks. The first target analysis reports the AJAX calls and URLs found during exploring the client-side web application, which can be used for REST API testing of the server side. The second target analysis is DOM-based XSS detection using taint inference, as described in Sec. IV. Details of each set of benchmarks are described in the following sections.

### B. RQ1: effectiveness of our feedback-driven and guided crawling technique

First we evaluate the effectiveness of our feed-back driven and guided crawling technique by measuring the performance and coverage of AJAX calls on two internal Oracle applications, WebScanTest [29], which is the live instance of a program used to evaluate crawlers, and Juice Shop v8.3.0 [18], a modern and sophisticated deliberately vulnerable web application. We decided to use these applications for our experiments because they make use of various modern technologies and libraries, such as jQuery [15], Knockout.js [19], React.js [26], and AngularJS [2]. To compare the overall coverage of our approach with other crawlers, we ran well-known crawlers on open-source applications that are also used as benchmarks in [50].

Fig. 2 shows how our guided crawling strategy compares to a non-guided random crawling strategy. Note that the random strategy still benefits from our dynamic analysis techniques

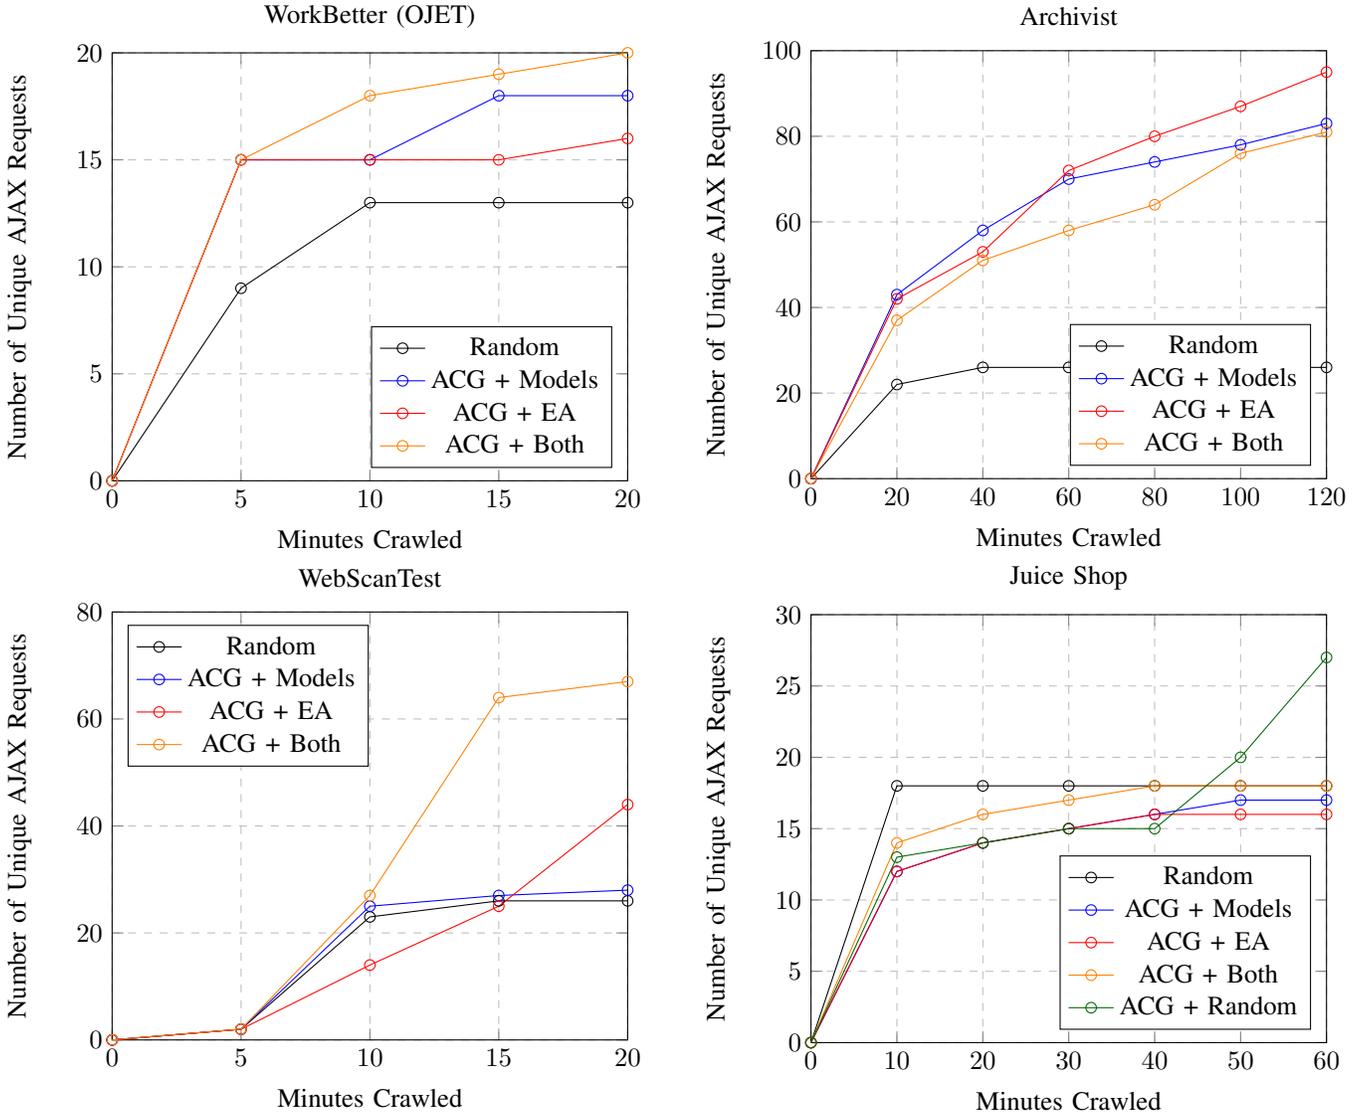


Fig. 2: Evaluating guided crawling against random crawling.

and only replaces the state and event prioritization functions in Algorithm 1 with random selection.

In this experiment, we count the number of distinct AJAX calls made by each strategy over time. The timeout for this experiment is 120 minutes, but we don't show the results once all the strategies start to plateau or they reach a fixpoint, i.e., finish running, before the timeout. We evaluate three versions of our guided crawling technique: (1) ACG + Models, which uses manually crafted models for libraries, such as jQuery when the approximate call graph fails to analyze them effectively; (2) ACG + EA, which is our feedback-driven edge addition technique to refine the statically generated call graph during runtime execution and add newly found edges; and (3) ACG + BOTH, which uses both Models and EA. By comparing these three versions, we evaluate the effectiveness of our novel ACG Edge Addition technique, which is fully automatic and can be used when manually crafted models for libraries are not available.

**Archivist:** Archivist is a highly interactive application for document management that uses the AngularJS [2] framework and jQuery [15] library. The random strategy performs significantly worse than our ACG guided strategies consistently. The ACG + EA is not as effective as the manually crafted models for this application but the gap between them is small. Therefore, replacing the manual models with the edge addition technique can be promising.

**WorkBetter:** WorkBetter is a tutorial application used to demonstrate the OJET [22] framework. Apart from OJET, it also uses the Knockout.js [19] framework and jQuery library. All of the ACG guided strategies outperform the random strategy on this application. Among the ACG strategies, ACG + BOTH outperforms the others and the EA is more effective than the manually crafted models. The results on this application show that not only EA can replace the manual models but also achieves better coverage within the same amount of time.

**WebScanTest:** The ACG Edge Addition strategy for this benchmark outperforms the two other strategies for the AJAX call detection experiment. The ACG Model finds almost the same number of AJAX calls as the random strategy. ACG + Both outperforms the other strategies and EA is more effective than the manually crafted models towards the end.

**Juice Shop v8.3.0:** Unlike the other applications in this experiment, the ACG strategies do not outperform Random, ACG + Both catching up with Random only after 40 minutes of crawling. Juice Shop is a highly AJAX-driven application and many of the events result in triggering an AJAX call. Therefore, even a Random strategy can be effective due to the nature of the application. In fact, our closer investigation revealed that ACG strategies overprioritize certain events in this case, leaving few chances for other events to be triggered. To further understand this behaviour, we experimented with a hybrid strategy (ACG + Random), where one in five events is chosen randomly and the rest of the events are prioritized based on the call graph distance metric. The result shows that this hybrid strategy significantly outperforms the pure ACG strategies after 45 minutes. In future, we plan to experiment with such hybrid strategies for applications that are similar in nature to Juice Shop.

In summary, we show that our guided crawling strategy finds AJAX calls more quickly than the random strategy in all three applications, which make use of modern technologies and are non-trivial to crawl. Moreover, our feedback-driven edge addition (ACG EA) technique is shown to be effective: the number of AJAX calls found by this technique are in the same ballpark as the carefully constructed manual models.

In addition to the guided crawling strategy, we integrate dynamic analysis of JavaScript, link extraction and state-aware crawling in a novel way. We compare this novel design against state-of-the-art crawlers on open-source applications that are also used as benchmarks in [50]. Table II shows the total number of URLs and AJAX calls recorded by each tool. The results show that our crawler records more URLs and AJAX calls than other tools except for Wivet benchmark (Arachni finds six more URLs).

TABLE II: Comparing our crawler against state-of-the-art crawlers for number of URLs and AJAX calls.

Crawlers	Tested Web Applications			
	Public firing range	DVWA	Wivet	WebGoat
Our prototype	268	76	77	85
Arachni [3]	257	27	83	3
Crawljax [44]	83	20	15	2
Htcp [13]	160	63	68	65
jÄk [50]	257	34	76	83

### C. Results for DOM-based XSS detection

For the second target analysis, DOM-based XSS detection, first we evaluate the effectiveness of our staged taint inference techniques on micro benchmarks and libraries with known DOM-based XSS vulnerabilities. We use two open-source micro benchmarks designed to evaluate DOM-based XSS detection tools: Firing Range [10]<sup>5</sup> from Google, and IBM

<sup>5</sup>For the Firing Range benchmark, we evaluate only against the DOM-related test cases: address tests [8], urldom tests [11] and dom tests (toxicdom) [9].

benchmarks [20]. We also compare our taint inference technique against dynamic taint tracking in DexterJS [49], and CTT (Chromium Taint Tracking) [42], the state-of-the-art DOM-based XSS detection tools. Next, we evaluate the effectiveness of our guided crawling strategy and input generation technique for DOM-based XSS detection analysis.

1) *Taint flow inference on microbenchmarks:* The inputs to a DOM-based XSS vulnerability are often provided in the URLs, which can easily be controlled by attackers. While some of the test cases in the Firing Range and IBM benchmarks contain a valid flow from a source to a sink, the value at the source cannot be directly tainted through URLs, e.g., `sessionStorage`. Table III considers only test cases that can be triggered through a user-controlled URL, effectively also omitting the test cases for which DexterJS and CTT report false positives. The criteria used to label test cases as being *controllable* are as follows:

- Contains at least one valid URL-controllable input (source) that enters the JavaScript program.
- Contains at least one valid sink.

TABLE III: Vulnerability detection on microbenchmarks – URL-controllable test cases

Benchmark	Firing Range			IBM
	address tests	urldom tests	dom tests (toxicdom)	
# Test Cases	28	26	3	66
# GELATO	23	22	3	60
# DexterJS	21	6	3	53
# CTT	18	10	3	54

Table III shows that our framework finds more taint flows than DexterJS and CTT. This is particularly apparent in urldom test cases, where our framework is able to detect a significantly higher portion of the URL-controllable test cases. In particular, we can report the taint flow in `Incorrect_Sanitizer/apollo_test_01.html` test case in the IBM benchmark, which is triggered if the URL has `topic=` query parameter. GELATO successfully finds this query parameter and reports the taint flow. This test case shows the effectiveness of our value input generator and constraint heuristics explained in Sec.III-E.

We investigated the results from the other tools further to understand why they fail to report many valid taint flows. We noticed that DexterJS does not handle all built-in functions, sanitizations and browser APIs. On the other hand, CTT does not handle property reads and writes, which result in false negatives. We also noticed problems in the event-generation component of DexterJS that leads to poor coverage and missing valid taint flows. Because CTT does not have support for event generation, it misses all the flows that require user interaction. Also, the sources and sinks in some of the test cases are not supported in these tools. However, our framework can handle most of these test cases and successfully reports valid taint flows.

Table IV shows how our framework outperforms the dynamic taint tracking used in the other tools for false positives. These results show that the false positive reduction strategies discussed in Sec. IV are very effective in lowering false positive rate and improving precision. However, whether all the 33 dom tests are actually false positives is debatable. The

TABLE IV: False positives detected on non URL-controllable test cases

Benchmark	Firing Range			IBM
	address tests	urldom tests	dom tests (toxicdom)	
# Test Cases	1	0	33	70
# GELATO	0	0	0	0
# DexterJS	0	0	26	3
# CTT	0	0	12	6

taint source in these test cases are not URL controllable and are resources such as cookies and `localStorage` that are set by the application if they are empty. Therefore, if they are not set, the application sets them to benign values. Listing 4 shows an example where the `badValue` item in the `localStorage` is set to a constant value that is benign if it is not set before, hence GELATO does not report any taint flows.<sup>6</sup> CCT still considers the benign values set to these resources as valid taint sources (instead of killing the taint). Because the taint analysis in [42] is followed up by an exploit generation step that confirms the taint flows, their DOM-based XSS tool might not report them as exploitable vulnerabilities.<sup>7</sup>

```

1 <html>
2 <head><title>Toxic DOM</title></head>
3 <body>
4 <script type="text/javascript">
5   if (!localStorage.getItem('badValue'))
6     localStorage.setItem('badValue', 'false
→   positive');
7   document.write(localStorage.getItem('badValue'));
8 </script>
9 </body>
10 </html>

```

Listing 4: An example of HTML/JavaScript code where the taint source is set by a benign constant value if it is empty.

2) *Taint flow inference on popular JavaScript libraries and open-source applications* : Table V reports the effectiveness of our taint inference mechanism on some of the JavaScript libraries that have known vulnerabilities, as reported in RetireJS [27]. RetireJS documents vulnerabilities in “retired” versions of JavaScript libraries. We have created test harnesses for these libraries that trigger the DOM-based XSS vulnerable paths. The payloads tested in these experiments are injected as payloads for all the tools. GELATO is able to report taint flows for all of these vulnerable libraries while DexterJS misses all of them and CTT misses two. Dojo is an interesting library in our benchmarks because its vulnerability can be found if the analysis can bypass the input validation. The code-snippet in Listing 5 shows a simplified version of the vulnerability in this library. GELATO successfully reports the taint flow from line 6 to 21 by appending `theme` as the query parameter in the URL and bypassing the input validations at lines 5 and 20 using the constraint heuristics explained in Sec.III-E.

Table VI shows that GELATO can successfully detect DOM-based XSS vulnerabilities in non-trivial modern web-applications, while the other tools miss reporting them. To detect the vulnerabilities in these applications a state-aware crawler is needed that triggers the vulnerable path. In this experiment, GELATO is the only tool that has an effective

<sup>6</sup>We do inject any payloads to any resources, such as cookies and `localStorage` in any of the experiments in this paper.

<sup>7</sup>We did not have access to the exploit generation component to confirm.

TABLE V: Vulnerability detection on JavaScript libraries with known vulnerabilities. Some of the payloads are injected into the fragment identifier part of the URLs: `jQuery 1.11.1` does not require any specific payload to reach the vulnerable sink; `<img src=/onerror=alert(1)>` is used for `jQuery 1.6.1` and `jQuery-migrate 1.1.1` and `x"onerror=alert(1) nothing` for `handlebars 1.0.0.beta.2` and `mustache 0.3.0`. The payload for `dojo` is injected as a URL query parameter and needs a specific key: `?theme=payload`, which is automatically generated by GELATO.

Benchmark	GELATO	DexterJS	CTT
jQuery 1.11.1 [5]	✓	×	×
jQuery 1.6.1 [4]	✓	×	✓
jQuery-migrate 1.1.1 [16]	✓	×	✓
handlebars 1.0.0.beta.2 [12]	✓	×	✓
mustache 0.3.0 [21]	✓	×	✓
dojo 1.4.1 [7]	✓	×	×

state-aware crawler, successfully exploring the application and triggering the vulnerable path.

TABLE VI: Open-source deliberately vulnerable applications.

Benchmark	GELATO	DexterJS	CTT
Juice-shop 8.3.0 [18]	✓	×	×
Damn Vulnerable Web App (DVWA) [6]	✓	×	×

3) *Effectiveness of guided crawler and data input generation for DOM-based XSS detection*: We evaluate the effectiveness of guided crawling strategy for DOM-based XSS detection analysis by comparing it with the random strategy. The target locations in this experiment are DOM manipulation operations in the program. Table VII shows that for both of our microbenchmarks, the guided crawling helps find the DOM-based XSS vulnerabilities faster than a crawler that uses a random strategy.

```

1 <!DOCTYPE html>
2 <html>
3 <body>
4 <script>
5   if (window.location.href.indexOf("?") > -1) {
6     loc = window.location.href
7     var str = loc.substr(loc.indexOf("?") +
→ 1).split(/#/);
8     var ary = str[0].split(/&/);
9     for (var i = 0; i < ary.length; i++) {
10      var split = ary[i].split(/=/),
11        key = split[0],
12        value = split[1];
13      switch (key) {
14        case "theme":
15          theme = value;
16          break;
17      }
18    }
19  }
20  if (theme) {
21    document.write('<link rel="stylesheet"
→ type="text/css" href="' + theme + '>');
22  }
23 </script>
24 </body>
25 </html>

```

Listing 5: A simplified snippet from the Dojo version 1.4 library that is vulnerable to DOM-based XSS, where attacker-controllable input enters via `theme` URL query parameter through the `window.location.href` property and is written to DOM at line 21.

#### D. Threats to validity

While we aimed to select representative applications and benchmarks that use modern technologies, the choice of benchmarks might have affected the validity of the experiments

TABLE VII: Comparing our guided crawling strategy (ACG) with random strategy (Random) for DOM-based XSS detection.

Benchmark	Analysis Step	ACG	Random
Firing Range	Link Extraction	19m, 51s	19m, 39s
	State crawling	3m, 35s	10m, 12s
	Taint inference	10m, 39s	40m, 18s
	Total	34m, 31s	70m, 27s
IBM	Link Extraction	16m, 41s	16m, 56s
	State crawling	7m, 6s	6m, 17s
	Taint inference	14m, 42s	52m, 42s
	Total	38m, 58s	76m, 55s

presented in this paper. In the experiments we showed that the DOM-based XSS detection in GELATO has a high accuracy for the analyzed applications and libraries. However, depending on the complexity of the taint manipulation operations in the given program, the accuracy can vary.

## VII. CONCLUSION

In this paper, we proposed GELATO, a dynamic analysis tool that detects vulnerabilities in modern and complex client-side JavaScript applications, which are often built upon libraries and frameworks. We studied the state-of-the-art tools and presented the most crucial features a security-aware client-side analysis should be supporting. We proposed the first security-guided client-side analysis that closes the gap between state-aware crawling and client-side security analysis by taking a feedback-driven approach that combines static and dynamic analysis to analyze complex frameworks automatically, and increase the coverage of security-sensitive parts of the program efficiently. We evaluated GELATO on various applications and benchmarks with different levels of complexity, and showed it outperforms the existing crawlers. Finally, we proposed a new lightweight non-intrusive client-side taint analysis that reports non-trivial taint flows on modern JavaScript applications, and has higher accuracy compared to the existing dynamic client-side taint analysis tools.

## REFERENCES

- [1] Alexa top. <https://www.alexa.com/topsites>.
- [2] AngularJS. <https://angularjs.org/>.
- [3] Arachni Framework v1.4. <http://www.arachni-scanner.com/blog/tag/crawl/>.
- [4] CVE-2011-4969. <http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2011-4969>.
- [5] CVE-2015-9251. <https://nvd.nist.gov/vuln/detail/CVE-2015-9251>.
- [6] Damn Vulnerable Web Application (DVWA). <http://dvwa.co.uk/>.
- [7] DOM-Based XSS in Dojo Toolkit SDK. <https://bugs.dojotoolkit.org/ticket/10773>.
- [8] Firing Range: Address DOM-based XSS. <https://public-firing-range.appspot.com/address/index.html>.
- [9] Firing Range: DOM tests (toxicdom). <https://public-firing-range.appspot.com/dom/index.html>.
- [10] Firing Range Test Bed. <https://public-firing-range.appspot.com/>.
- [11] Firing Range: URL-based DOM-based XSS. <https://public-firing-range.appspot.com/urlldom/index.html>.
- [12] Handlebars.js GitHub Issue #68. <https://github.com/wycats/handlebars.js/pull/68>.
- [13] Htcap v1.0.1. <https://github.com/segment-srl/htcap/releases>.
- [14] Jalangi2. <https://github.com/Samsung/jalangi2>.
- [15] jQuery. <https://jquery.com/>.
- [16] jQuery Bug Tracker Issue #11290. <https://bugs.jquery.com/ticket/11291>.
- [17] jsTaint. <https://github.com/idkwim/jsTaint>.
- [18] Juice-shop 8.3.0. <https://github.com/bkimminich/juice-shop>.
- [19] Knockout.js. <https://knockoutjs.com/>.
- [20] LaBaSec: Language-based Security. [http://m.ibm.com/http/researcher.ibm.com/researcher/view\\_group\\_subpage.php?id=1598](http://m.ibm.com/http/researcher.ibm.com/researcher/view_group_subpage.php?id=1598).
- [21] Mustache.js GitHub Issue #112. <https://github.com/janl/mustache.js/issues/112>.
- [22] OJET. <https://www.oracle.com/webfolder/technetwork/jet/index.html>.
- [23] OWASP DOM-Based XSS. [https://www.owasp.org/index.php/DOM-Based\\_XSS](https://www.owasp.org/index.php/DOM-Based_XSS).
- [24] OWASP ZAP. [https://www.owasp.org/index.php/OWASP\\_Zed\\_Attack\\_Proxy\\_Project](https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project).
- [25] Pypeteer. <https://github.com/miyakogi/pypeteer>.
- [26] React.js. <https://reactjs.org/>.
- [27] Retire.js. <http://retirejs.github.io/retire.js/>.
- [28] w3af v1.6.49. <http://w3af.org/plugins/crawl>.
- [29] WebScanTest. <https://www.webscantest.com/>.
- [30] Esben Andreasen, Gong Liang, Michael Pradel, Tu Darmstadt, Marija Selakovic, and Koushik Sen. A Survey of Dynamic Analysis and Test Generation for JavaScript. *ACM Computing Surveys*, 2017.
- [31] Shay Artzi, Julian Dolby, Simon Holm Jensen, Anders Møller, and Frank Tip. A Framework for Automated Testing of JavaScript Web Applications. In *ICSE*, 2011.
- [32] Kamara Benjamin, Gregor Von Bochmann, Mustafa Emre Dincturk, Guy-Vincent Jourdan, and Iosif Viorel Onut. A Strategy for Efficient Crawling of Rich Internet Applications. In *ICWE*, 2011.
- [33] Marcel Böhme, Van-Thuan Pham, Manh-Dung Nguyen, and Abhik Roychoudhury. Directed Greybox Fuzzing. In *CCS*, 2017.
- [34] Laurent Christophe, Elisa Gonzalez Boix, Wolfgang De Meuter, and Coen De Roover. Linvail: A general-purpose platform for shadow execution of JavaScript. In *SANER*, 2016.
- [35] Asger Feldthaus, Max Schäfer, Manu Sridharan, Julian Dolby, and Frank Tip. Efficient Construction of Approximate Call Graphs for JavaScript IDE Services. In *ICSE*, 2013.
- [36] François Gauthier, Behnaz Hassanshahi, and Alexander Jordan. AF-FOGATO: Runtime Detection of Injection Attacks for Node.js. In *SOAP*, 2018.
- [37] Salvatore Guarnieri, Marco Pistoia, Omer Tripp, Julian Dolby, Stephen Teilhet, and Ryan Berg. Saving the World Wide Web from Vulnerable JavaScript. In *ISSSTA*, 2011.
- [38] Matthias Hörschele and Andreas Zeller. Mining Input Grammars with AUTOGRAM. In *ICSE-C*, 2017.
- [39] Casper Svenning Jensen, Anders Møller, and Zhendong Su. Server interface descriptions for automated testing of JavaScript web applications. In *ESEC/FSE*, 2013.
- [40] Simon Holm Jensen, Anders Møller, and Peter Thiemann. Type Analysis for JavaScript. In *SAS*, 2009.
- [41] Hongki Lee, Sooncheol Won, Joonho Jin, Junhee Cho, and Suyoung Ryu. SAFE: Formal specification and implementation of a scalable analysis framework for ECMAScript. In *FOOL*, 2012.
- [42] Sebastian Lekies, Ben Stock, and Martin Johns. 25 million flows later: large-scale detection of DOM-based XSS. In *CCS*, 2013.
- [43] William Melicher, Anupam Das, Mahmood Sharif, Lujia Bauer, and Limin Jia. Riding out DOMsday: Towards Detecting and Preventing DOM Cross-Site Scripting. In *NDSS*, 2018.
- [44] Ali Mesbah, Arie van Deursen, and Stefan Lenselink. Crawling Ajax-Based Web Applications Through Dynamic Analysis of User Interface State Changes. *ACM Trans. Web*, 2012.
- [45] Ali Mesbah, Arie van Deursen, and Danny Roest. Invariant-Based Automatic Testing of Modern Web Applications. *IEEE Trans. Softw. Eng.*, 2012.
- [46] Amin Milani Fard and Ali Mesbah. Feedback-directed exploration of web applications to derive test models. In *ISSRE*, 2013.
- [47] S Needleman and C Wunsch. A general method applicable to the search for similarities in the amino acid sequences of two proteins. In *J. Mol. Biol.* 48, 1970.

- [48] Christopher Olston and Marc Najork. Web Crawling. *Found. Trends Inf. Retr.*, 2010.
- [49] Inian Parameshwaran, Enrico Budio, Shweta Shinde, Hung Dang, Atul Sadhu, and Prateek Saxena. Auto-patching DOM-based XSS at scale. In *FSE*, 2015.
- [50] Giancarlo Pellegrino, Constantin Tschürtz, Eric Bodden, and Christian Rossow. jÄk: Using Dynamic Analysis to Crawl and Test Modern Web Applications. In *RAID*, 2015.
- [51] Sanjay Rawat, Vivek Jain, Ashish Kumar, Lucian Cojocar, Cristiano Giuffrida, and Herbert Bos. VUzzer: Application-aware Evolutionary Fuzzing. In *NDSS*, 2017.
- [52] Prateek Saxena, Devdatta Akhawe, Steve Hanna, Feng Mao, Stephen McCamant, and Dawn Song. A Symbolic Execution Framework for JavaScript. In *SP*, 2010.
- [53] Koushik Sen, Swaroop Kalasapur, Tasneem Brutch, and Simon Gibbs. Jalangi: A Selective Record-replay and Dynamic Analysis Framework for JavaScript. In *ESEC/FSE*, 2013.
- [54] Ben Spencer, Michael Benedikt, Anders Møller, and Franck van Breugel. ArtForm: A Tool for Exploring the Codebase of Form-based Websites. In *ISSTA*, 2017.
- [55] Suresh Thummalapenta, K. Vasanta Lakshmi, Saurabh Sinha, Nishant Sinha, and Satish Chandra. Guided Test Generation for Web Applications. In *ICSE*, 2013.
- [56] Philipp Vogt, Florian Nentwich, Nenad Jovanovic, Engin Kirda, Christopher Krügel, and Giovanni Vigna. Cross Site Scripting Prevention with Dynamic Data Tainting and Static Analysis. In *NDSS*, 2007.