Fragment-Based Test Generation For Web Apps

Rahulkrishna Yandrapally, *Student Member, IEEE,* and Ali Mesbah, *Member, IEEE*

◆

Abstract—Automated model-based test generation presents a viable alternative to the costly manual test creation currently employed for regression testing of web apps. However, existing model inference techniques rely on threshold-based whole-page comparison to establish state equivalence, which cannot reliably identify near-duplicate web pages in modern web apps. Consequently, existing techniques produce inadequate models for dynamic web apps, and fragile test oracles, rendering the generated regression test suites ineffective. We propose a model-based test generation technique, FRAGGEN, that eliminates the need for thresholds, by employing a novel state abstraction based on page fragmentation to establish state equivalence. FRAGGEN also uses fine-grained page fragment analysis to diversify state exploration and generate reliable test oracles. Our evaluation shows that FRAGGEN outperforms existing whole-page techniques by detecting more near-duplicates, inferring better web app models and generating test suites that are better suited for regression testing. On a dataset of 86,165 state-pairs, FRAGGEN detected 123% more near-duplicates on average compared to whole-page techniques. The crawl models inferred by FRAGGEN have 62% more precision and 70% more recall on average. FRAGGEN also generates reliable regression test suites with test actions that have nearly 100% success rate on the same version of the web app even if the execution environment is varied. The test oracles generated by FRAGGEN can detect 98.7% of the visible changes in web pages while being highly robust, making them suitable for regression testing.

Index Terms—Software Testing, Web Testing, Web Application Model Inference, Automatic Web App Exploration, Web Page - State Abstraction and Equivalence, State Abstraction, Web Application crawling, Test Generation

1 INTRODUCTION

Regression testing of modern web apps is a costly activity [1] in practice, which requires developers to manually create test suites, using a combination of programming and record/replay tools such as Selenium [2]. In addition, maintaining such test suites is known to be costly [3], [4] as even minor changes of the app can cause many tests to break; for example, according to a study at Accenture [4] even simple modifications to the user interface of apps result in 30–70% changes to tests, which costs $120 million per year to repair. When the test maintenance cost becomes overwhelming, whole test suites are abandoned [5].

Given the short release cycles of modern web apps and maintenance costs of manually written tests, automatic generation of regression test suites seems a viable alternative. However, the effectiveness of web test generation techniques [6], [7], [8] is limited by the ability to obtain an accurate and complete model of the app under test. Manual construction of such models for complex apps is not practical. Automated model inference techniques [9]

trigger user actions such as clicking on buttons and record corresponding transitions between states in the web app to build a graph-based model.

One particular challenge here is the presence of *near-duplicate* states in web apps, which can adversely impact the inferred model in terms of redundancy and adequacy [10]. Near-duplicates are states that are *similar* to each other in terms of functionality [11].

Another important challenge is the generation of test assertions. Two factors directly contribute to the challenge here, namely effectiveness and tolerance/robustness, i.e., regression test assertions should be able to detect unexpected app behavior, but at the same time, be tolerant to minor changes that do not affect the functionality.

Both model inference and test oracle generation thus require suitable abstractions to produce effective and robust test suites. Existing techniques generate regression tests that compare the *whole* page as seen during testing with an instance of the page recorded on a previous version [2], [7]. Our insight is that, such whole-page comparison techniques, although effective at detecting changes, are not tolerant enough to handle near-duplicates and make the test suites fragile. Test fragility is known to be a huge problem in web testing [12].

In this work, we propose a novel state abstraction that employs fine-grained fragments to establish functional equivalence. We conjecture that a web page is not a singular functional entity and thus partitioning it into separate fragments can help in determining functional equivalence when comparing different states. Using this novel state abstraction, we have developed a technique, called FRAGGEN. Our fragment-based analysis enables us to (1) prioritize available actions to diversify exploration, (2) accomplish state comparison without the need for manually selecting thresholds, a manual tedious fine-tuning process, required for all existing techniques [10]; our state comparison algorithm leverages both *structural* and *visual* properties of the page fragments to identify near-duplicate characteristics specific to the web app under test during model inference, and (3) generate test assertions that operate at the fragment level instead of the whole page level, and apply fragment *memoization* to make them much more robust to state changes that should not break regression tests.

Our empirical evaluation shows that FRAGGEN is able to outperform whole-page techniques in classifying state-pairs and identifying near-duplicates. On a dataset of 86,165 manually labelled state-pairs, FRAGGEN detected 123% more near-duplicates on average and 82% more than the best
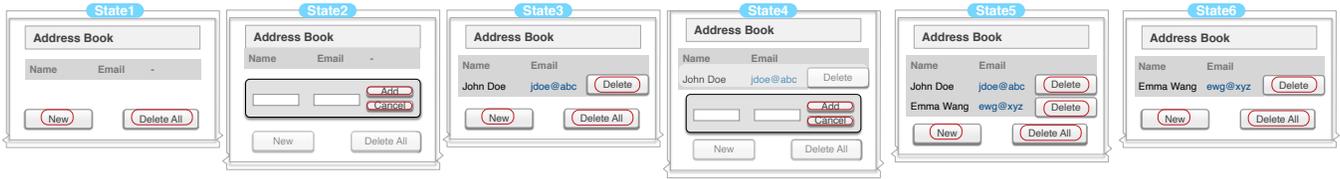
Fig. 1: Motivating example: app states with *actionables* highlighted.

performing existing technique. When employed to infer web app models, FRAGGEN is able to diversify exploration using fine-grained fragment analysis to produce models with 70% higher precision and 62% higher recall on an average compared to the state-of-the-art technique. In addition, our evaluation shows that FRAGGEN generated test suites that are better suited for regression testing. Where existing techniques generated brittle test suites with nearly 17% test actions that fail even on the same version of the web app, FRAGGEN generated reliable test suites with nearly 100% successful test actions on the same version of the web app, and detected more app changes with fewer false positives when run on different application versions. When evaluated through mutation analysis, FRAGGEN's test oracles could detect 98.7% of visible state changes while being tolerant enough to ignore 90.6% of equivalent mutants.

This paper makes the following contributions:

- A technique for determining state equivalence (i.e., distinct, clone, near-duplicate) at the fragments-level, without a need for setting thresholds.
- A state abstraction technique that uses the structural as well as visual properties of fragments for equivalence checking.
- A novel model inference approach that employs page fragments to explore the web application state-space.
- A fine-grained fragment-based test assertion generation for effective and reliable regression testing.
- The evaluation and implementation of FRAGGEN, which is publicly available [13].

2 BACKGROUND AND MOTIVATION

In this section, we provide the background information on web app testing and analysis, and introduce key terms and concepts that are used in the rest of the paper.

We use a running example, shown in Figure 1, based on one of our subject systems [14]. The example web app is a single page application which provides functionality to add, view, modify and delete addresses in a database through "actionable" web elements such as *buttons* and *links*, highlighted in their corresponding pages.

2.1 Automatic Test Generation for Web Apps

In this subsection, we describe automated model inference through state exploration and model-based test generation employed by existing techniques and their limitations.

Model inference. Automated model inference is an iterative process of exercising the functionality of a given web app by triggering events on actionable elements (α), such as *button clicks*, and capturing the resulting state transitions ($\mathcal{A}$) as a graph-based model ($\mathcal{M}$). Formally:
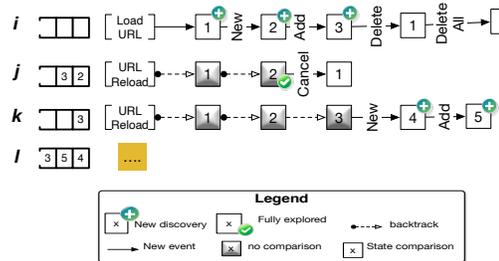

Fig. 2: Model inference for the motivating example.
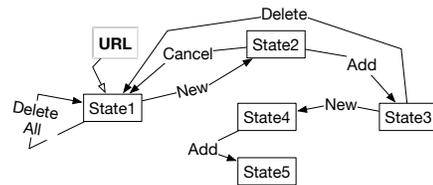
Fig. 3: Inferred model of the motivating example.

Definition 1 (**State Transition ($\mathcal{A}_x$)**). is a tuple ($\mathcal{S}_{src}$, α_x, $\mathcal{S}_{tgt}$) where exercising an actionable α_x in a state $\mathcal{S}_{src}$ produces a transition to state $\mathcal{S}_{tgt}$.

Definition 2 (**Application Model ($\mathcal{M}$)**). is a directed graph ($\{\mathcal{S}_1..\mathcal{S}_n\}$, $\{\mathcal{A}_1..\mathcal{A}_m\}$) with *app states* ($\mathcal{S}_a$) as nodes and *state transitions* ($\mathcal{A}_x$) as labelled directed edges between nodes.

Current model inference techniques rely on a state abstraction function (SAF), that determines similarity between two given states p_1 and p_2 in order to avoid redundancies in the captured model and duplication of exploration effort. Formally:

Definition 3 (**State Abstraction Function (SAF)**). is a pair ($dfunc$, t), where $dfunc$ is a similarity function that computes the distance between any two given web pages p_1, p_2, and t is a threshold defined over the output values of $dfunc$. SAF determines whether the distance between p_1 and p_2 falls below t.

$$SAF(dfunc, p_1, p_2, t) \begin{cases} true & : dfunc(p_1, p_2) < t \\ false & : otherwise \end{cases}$$

Figure 2 illustrates the steps of model inference for our motivating app, assuming a depth-first exploration strategy is followed. In the first iteration, labelled *i*, the technique loads the app in the browser using its URL, and stores the corresponding state as the root node in the model. Thereafter, each action performed on the web app can

Listing 1: Generated test case

```
def Test1():
    driver.loadURL("Base_URL")
    assert(isEqual(driver.currentState, state1)
    driver.findElement("List").click()
    assert(isEqual(driver.currentState, state2)
    driver.findElement("New").click()
    assert(isEqual(driver.currentState, state3)
    driver.findElement("add").click()
    assert(isEqual(driver.currentState, state4)
```

be either *exploration* step or a *back-tracking* step. In each exploration step, depicted in solid arrows, an unexplored actionable (α_{new}) from the current state is invoked and the resulting state is added to the model if it is deemed to be *different* from every existing state in the model by a SAF as defined in defintion 3. The observed state transition is then recorded as a directed edge between the source and target states in the model.

An iteration ends when the current state is fully explored and the next iteration starts by choosing an existing state in the model with unexplored actionables. In order to reach the selected state, *back-tracking* actions, indicated with broken arrows, are performed by using the transitions that are already recorded in the model. For example, iteration "i" ends upon reaching state $\mathcal{S}_1$, which is fully explored at that point. The next iteration "j" then starts by choosing and navigating to state $\mathcal{S}_2$, by using recorded transitions. Iteration "l" in Figure 2 continues exploration by choosing one of the unexplored states $\mathcal{S}_3$, $\mathcal{S}_5$ and $\mathcal{S}_4$.

Termination and stopping criteria. Model inference techniques provide options to configure *stopping criteria* such as exploration time limit in order to end the inference process. *Termination* on the other hand happens when the technique decides that no unexplored actionables are left to exercise.

Figure 3 shows the model inferred for our motivating example at the end of iteration "k". The model inference has not *terminated* at this point because there are unexplored actionables available.

Definition 4 (Path ($\mathcal{P}$)). A sequence of transitions $(\mathcal{A}_0...\mathcal{A}_n)$ is a $\mathcal{P}$ if for 0<=i<n, $\mathcal{A}_i$, $\mathcal{A}_{i+1} \in \mathcal{P} \implies \mathcal{A}_i(\mathcal{S}_{tgt}) = \mathcal{A}_{i+1}(\mathcal{S}_{src})$.

Test generation. Once a model of web app is available, model-based test generation provides a set of paths $\{\mathcal{P}_1, \mathcal{P}_m\}$, with adequate coverage of states and/or transitions in the model, where, each path $\mathcal{P}$ (formally defined in 4) is a sequence of recorded transitions. Listing 1 shows a test case $\mathcal{T}$ generated from the inferred model of Figure 3. Each test case starts by loading the URL of the app and verifying the browser state to be the $\mathcal{S}^0_{src}$, i.e. the source state of the first transition of the path. Thereafter, for each transition ($\mathcal{A}_x$), a *test action* is derived from the actionable α_x and a *test oracle* is added for the target state $\mathcal{S}^x_{tgt}$. The test case in Listing 1 is examining the path $\mathcal{P}$: $[\mathcal{S}_1, \alpha_{List}, \mathcal{S}_2, \alpha_{new}, \mathcal{S}_3, \alpha_{add}, \mathcal{S}_4]$.

2.2 Automatic Test Generation Challenges

The main challenge in model inference of web apps is the presence of a large number of nearly identical or near-duplicate web pages, which the existing techniques cannot identify effectively and as a result generate sub-optimal models.

While the presence of near-duplicates in the inferred web app model leads to generation of redundant test cases, it is also indicative of wasted exploration effort that could be spent discovering unseen states. Typically, exploration of similar actions in near-duplicate states does not improve functional coverage of the model but instead can lead to creation of even more near-duplicates, as can be seen even in our simple example app. In our example, the states $\mathcal{S}_3$, $\mathcal{S}_5$ and $\mathcal{S}_6$ shown in Figure 1 are all considered *functional near-duplicates* as they all offer similar web app functionality, and model inference may never *terminate* if similar actions α_{New} and α_{Add} are explored in each near-duplicate state.

In our previous work, we proposed [10] to categorize a given pair of web pages as either clone (Cl), near-duplicate (Nd) or distinct (Di) by labelling the observed changes between them. We formally defined *Near-duplicates* in web testing as:

Definition 5 (Functional Near-Duplicate (Nd)). A given state-pair (p_1, p_2), is considered to be a functional near-duplicate if the *changes* between the states do not alter the overall functionality of either state.

Near-duplicates are further divided in three categories based on the nature of changes between the two pages:

- Cosmetic (Nd_1-*data*): Changes such as different advertisements, that are irrelevant to functionality of web app.
- Dynamic Data (Nd_2-*data*): Changes are limited to data while the page structure remains the same. Example ($\mathcal{S}_3$, $\mathcal{S}_6$).
- Duplication (Nd_3-*struct*): Addition or removal of web elements equivalent to existing web elements. Example ($\mathcal{S}_3$, $\mathcal{S}_5$).

Researchers have employed various similarity functions and abstractions for web pages based on their DOM tree-structures [15], [16], and visual screenshots [17], [18], [19], [20] to perform state comparison. However, a study [10] on near-duplicates in web testing shows that the whole-page based SAFs currently being used by existing model inference techniques cannot reliably identify near-duplicates and as a result infer imprecise and incomplete models.

To illustrate the limitations of the SAFs currently being used, we analyzed three state-pairs of our example app in Table 1. For each of the three state-pairs, the table shows the distance between states computed by each of the ten state abstraction techniques and the *human* classification process followed in the study [10].

TABLE 1: Raw distances of state-pairs

state-pair	RTED [21]	Levenshtein [22]	TLSH [23]	SimHash [24]	HYST [25]	BlockHash [26]	PHASH [27]	PDIFF [28]	SIFT [29]	SSIM [30]	Human
	DOM				VISUAL						
$(\mathcal{S}_3, \mathcal{S}_6)$	0.0	0.002	1	0	91	0	0	0.0006	5	0.003	Nd_2
$(\mathcal{S}_1, \mathcal{S}_3)$	0.13	0.08	36	0	9751	4	2	0.008	14	0.04	Di
$(\mathcal{S}_3, \mathcal{S}_5)$	0.16	0.12	106	0	27683	5	2	0.045	14	0.05	Nd_3

Consider the states $\mathcal{S}_1$, $\mathcal{S}_3$, $\mathcal{S}_5$, and $\mathcal{S}_6$ from Figure 1, all of which display stored addresses. When examined manually, $\mathcal{S}_1$ is considered distinct (Di) from $\mathcal{S}_3$, $\mathcal{S}_5$ and $\mathcal{S}_6$ as it does not contain table row functionality to select an address entry. However, $\mathcal{S}_3$ and $\mathcal{S}_6$ differ only by the *data* in a table row, which does not alter the functionality and hence they are considered as functional near-duplicates. Further, even-though $\mathcal{S}_5$ contains extra table rows, they only duplicate the functionality of a single row in $\mathcal{S}_3$. Therefore, a human tester would label state-pairs such as ($\mathcal{S}_3$, $\mathcal{S}_5$) to be functional near-duplicates.

The 10 comparison techniques, as shown in the Table 1, consider the state-pair ($\mathcal{S}_3$, $\mathcal{S}_5$) to be the farthest apart of the four state-pairs even though they are functionally equivalent. If the distance thresholds were set to be higher, states such as $\mathcal{S}_3$ might be discarded from the model as they will be considered equivalent to $\mathcal{S}_1$, making the model *incomplete*. On the other hand, lower thresholds would make the model imprecise with the presence of states such as the state $\mathcal{S}_5$. This shows that threshold-based whole-page comparison cannot produce optimal web app models.

Test breakages. In addition to model inference, similarity between two given web page states is crucial in generating effective test oracles as well. If the state comparison techniques are sensitive to minor changes unrelated to functionality in modern apps, test oracles can break and result in a high number of false positive test failures. False positive test failures necessitate costly manual analysis and impact the effectiveness of automated regression testing techniques.

We identified duplicated functionality within a web page to be the root cause of Nd_3-*struct* near-duplicates that make existing model inference and test generation techniques impractical for modern web apps. This observation that "whole-page techniques" cannot detect Nd_3-*struct* near-duplicates motivated us to investigate the idea of decomposing a given web page into smaller fragments for a more fine-grained analysis.

2.3 Page Fragmentation

Page fragmentation, also known as page segmentation, is the decomposition of a given web page into smaller fragments or segments. In existing research, the most popular downstream tasks such as content extraction that apply page fragmentation relate to human consumption of web pages and focus on extracting textual semantics.

The "VIsion-based Page Segmentation algorithm" (VIPS) [31], proposed in 2003, is the de-facto standard for web page segmentation. A recent large scale empirical study [32] compares five page segmentation techniques using a dataset of 8,490 web pages and concludes that VIPS is still the overall best option for page fragmentation. VIPS employs a top-down approach, where, for each HTML node in the DOM [33], a DoC (Degree of Coherence) is assigned to indicate coherence of the content in the block based on visual perception. Then it tries to find the separators between these extracted blocks. Here, separators denote the horizontal or vertical lines in a webpage that visually cross with no blocks. Finally, based on these separators, it extracts a hierarchy of fragments. Each extracted fragment visually conforms to a rectangle in the page between two horizontal
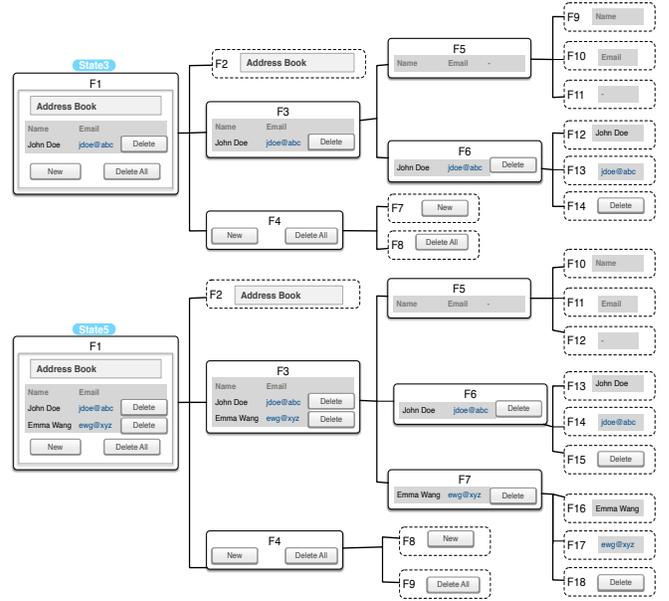

Fig. 4: Fragment-based state comparison in FRAGGEN

and vertical separators. All the DOM nodes that are part of the rectangle are then considered to be part of the fragment. We refer the interested reader to the original VIPS paper for more details of the algorithm [31].

In the next section, we explain how we leverage page fragmentation to overcome existing challenges in model inference and regression test generation for modern web apps.

3 APPROACH

At a high level, our approach, called FRAGGEN, relies on the insight that a web page is not a singular functional entity, but a set of functionalities, where each functionality may be available in more than one page. Based on this insight, we propose a novel state abstraction that defines a page as a *hierarchy of fragments*, where each fragment represents a semantic sub functionality. Our model inference technique then employs this state abstraction to detect near-duplicates and optimize the state space exploration by prioritizing actions that belong to unique page fragments in each page. From the inferred model, we subsequently generate test cases with robust assertions that rely on our fragment-based abstraction and are capable of reporting warnings in addition to errors by utilizing knowledge gained during model inference. Next, we describe our state abstraction, followed by our model inference, and test generation techniques.

3.1 Fragment-based State Abstraction

We consider a web page or a state as a hierarchy of fragments, where, each fragment is a portion of the state and represents functionalities offered by its child fragments.

Definition 6 (Application State ($\mathcal{S}$)). is a tuple ($\mathcal{D}$, $\mathcal{V}$, $\mathcal{F}_{root}$) where $\mathcal{D}$ is the dynamic DOM [33] of the page, $\mathcal{V}$ is the screenshot of the page and $\mathcal{F}_{root}$ is the root fragment.

Algorithm 1: Fragment-based classification

```
 1  Function Classify(F₁, F₂):
        // clone(Cl), distinct(Di), Nd₂-data, Nd₃-struct
        Output: class
 2      N_diff ⟵ treediff ( F₁.N , F₂.N )
 3      if N_diff = φ then                   /* Matching DOM */
 4          V_diff ⟵ imagediff ( F₁.V , F₂.V )
 5          if V_diff = φ then   /* Matching Screenshots */
 6              return Cl                     /* class: clone */
 7          else
 8              return Nd₂-data         /* class: Nd₂-data */
 9          end
10      else                      /* Check Child Fragments? */
11          return MapChildFragments(N_diff, F₁, F₂)
12      end
13  End Function

14  Function MapChildFragments(N_diff, F₁, F₂):
        // distinct(Di), Nd₃-struct
        Output: class
        /* Mapping child fragments for every
           changed node */
15      foreach n ∈ N_diff do
16          found ⟵ false
17          F_oth ⟵ ( n ∈ F₁.N ) ? F₂ : F₁
18          F_clo ⟵ closest ( n )
19          foreach F_chi ∈ F_oth.childen do
20              if Classify ( F_clo, F_chi ) != Di then
21                  found ← true         /* Found mapping */
22              end
23          end
24          if ! found then              /* No mapping found */
25              return Di                   /* class: Distinct */
26          end
27      end
28      return Nd₃-struct            /* class: Nd₃-struct */
29  End Function
```

The root fragment, $\mathcal{F}_{root}$, of a state $\mathcal{S}$ is the full page, and has no parent. It has all the nodes in the DOM tree $\mathcal{D}$ of $\mathcal{S}$. Therefore, comparing two states is the same as comparing their root fragments.

Definition 7 (Fragment ($\mathcal{F}$)). is a tuple ($\mathcal{N}$, V, $\{\mathcal{F}_{c_0}, \mathcal{F}_{c_1} \dots \}$) where $\mathcal{N}$ is the set of DOM nodes of $\mathcal{F}$, V is the screenshot of $\mathcal{F}$ and each $\mathcal{F}_{c_i}$ is a child fragment.

Comparing fragments. Using this fragment-based representation, we classify a given state-pair by comparing the fragments they are composed of. Figure 4 shows the fragment hierarchies for states $\mathcal{S}_3$ and $\mathcal{S}_5$ from our motivating example in Figure 1. Our fragment-based classification, which takes both structural and visual aspects into account is shown in Algorithm 1. The structural aspect (line 2) uses the nodes on the DOM subtree of the fragment after pruning textual content and attributes. The visual aspect (line 4) uses a localized screenshot of the fragment.

As algorithm 1 shows, we combine structural and visual analysis to identify near-duplicates. Using visual similarity instead of element attributes and textual data from DOM allows us to disregard changes in the DOM that have no visual impact. In addition, we found that visual comparison is effective in identifying changes in dynamic web elements such as carousels that use only JavaScript and CSS.

FRAGGEN classifies two fragments to be *clones* if their structural and visual properties are exactly the same. We employ APTED [34] and Color Histogram [25] to compare the structural and visual aspects of the fragments, respectively. These techniques have been employed individually as state abstractions for whole web pages in the literature [18], [10], but have not been combined to determine state equivalence. When the fragments differ visually, but are found to be structurally equivalent, FRAGGEN considers the fragments to be near-duplicates of the type Nd₂-*data* (see Definition 5).

In case they differ structurally, as root fragments of $\mathcal{S}_3$ and $\mathcal{S}_5$ do for example shown in Figure 4, the classification is performed by a mapping function (lines 14-28) that extracts all changed DOM nodes between the two fragments and maps the corresponding fragments in the hierarchy.

Given a pair of fragments $\mathcal{F}_1$ and $\mathcal{F}_2$ and a list of changed nodes between the two fragments $\mathcal{N}_{diff}$, the function $MapChildFragments$ gets the closest fragment $\mathcal{F}_{clo}$ for each changed DOM node (n). The closest fragment for a DOM node is the smallest child fragment in the fragment hierarchy containing n. Thereafter, in lines 19-23, it attempts to find an equivalent fragment for the closest fragment in the fragment hierarchy of the other fragment. For example, if n_1 is a changed node that belongs to $\mathcal{F}_1$, with $\mathcal{F}_{clo}$ as its closest fragment, then $\mathcal{F}_{oth}$ is $\mathcal{F}_2$ and we attempt to find if any of the child fragments of $\mathcal{F}_2$ are equivalent to $\mathcal{F}_{clo}$. If the closest fragment for any of the changed DOM nodes cannot be mapped to a child fragment of the other fragment (lines 24-26), we declare the two given fragments $\mathcal{F}_1$ and $\mathcal{F}_2$ to be distinct (Di). Otherwise, the two fragments are considered to be near-duplicates of type Nd₃-*struct*.

In the example shown in Figure 4, when classifying the root fragments for $\mathcal{S}_3$ and $\mathcal{S}_5$, $MapChildFragments$ would be called to classify $\mathcal{F}_3$ of states $\mathcal{S}_5$ and $\mathcal{S}_3$. All the changed DOM nodes ($\mathcal{N}_{diff}$) between the two fragments belong to $\mathcal{F}_7$ in $\mathcal{S}_5$, As the function iterates (line 19) through child fragments to find a mapping, eventually $\mathcal{F}_7$ of $\mathcal{S}_5$ will be found to be equivalent to $\mathcal{F}_6$ of $\mathcal{S}_3$. As a result, the fragment pair would be classified as Nd₃-*struct*. As the rest of the fragments are equivalent, the overall state-pair ($\mathcal{S}_3$, $\mathcal{S}_5$) would be classified as Nd₃-*struct* near-duplicates as well.

When the fragments containing changed DOM nodes cannot be mapped, FRAGGEN considers the two states to be *distinct*. One such example is the state-pair ($\mathcal{S}_1$, $\mathcal{S}_3$), which is classified to be *distinct* because the table row in $\mathcal{S}_3$ contains no equivalent fragment in $\mathcal{S}_1$. On the other hand, ($\mathcal{S}_3$, $\mathcal{S}_6$) are classified as Nd₂-*data* because the abstracted DOM hierarchy is equal, making them structurally similar, while the data changes inside the table make them visually dissimilar.

3.2 Fragment-based Model Inference

FRAGGEN infers the model of a given web app by iteratively triggering user interactions on actionables and recording the corresponding state transitions much like the existing techniques (described in subsection 2.1). After an action (α) is performed on a source state, the resulting browser state is compared to all the existing states in the model using the $Classify$ function algorithm 1. If the classification is clone or Nd₂-*data* for any of the existing states, the new state is discarded, otherwise, it is added to the model.

However, existing techniques assign equal importance to every actionable of a newly discovered state, often wasting

exploration effort exercising similar actionables. In contrast, FRAGGEN identifies similar actionables through fragment analysis to diversify the exploration. In this section, we describe how our fragment analysis is used to 1) diversify state exploration to discover unique states faster, and 2) identify *data-fluid* fragments to generate effective test oracles.

3.2.1 Exploration strategy

FRAGGEN ranks the actionables and states using Equation 2 and Equation 3, respectively, using the fragment comparisons in Equation 1 to determine the equivalence of actionables. Actionables are special DOM nodes such as buttons which are used for user interaction.

Where a function $Xpath(\alpha, \mathcal{F})$ provides a relative XPath expression [35] for α within $\mathcal{F}$, we decide the equivalence of two actionables $\alpha_x \in \mathcal{F}_x, \alpha_y \in \mathcal{F}_y$ using:

$$\alpha_x \equiv \alpha_y \Longleftrightarrow (\text{Classify}(\mathcal{F}_x, \mathcal{F}_y) = \text{Clone} \vee \text{Nd}_2\text{-}data)$$
$$\wedge (Xpath(\alpha_x, \mathcal{F}_x) == Xpath(\alpha_y, \mathcal{F}_y)) \quad (1)$$

Given a constant c_0 where $0 < c_0 <= 1$, and α_{eq} is an equivalent actionable determined using Equation 1, the score for an actionable α is computed as:

$$score(\alpha) = \begin{cases} -1, & \text{if } \alpha.\text{explored.} \\ 0, & \text{if } \exists\, \alpha_{eq} \mid \alpha_{eq}.\text{explored.} \\ c_0 * size(\{\alpha_{eq}\ldots\}), & \text{otherwise.} \end{cases} \quad (2)$$

The score for a state $\mathcal{S}$ is computed using the scores of actionables in the state as:

$$score(\mathcal{S}) = \sum_{n=1}^{size(\{\alpha\})} score(\alpha) \quad (3)$$

Using the Equation 3, FRAGGEN chooses the next state to explore based on the total score of the actionables in each state. The score for each actionable (α) is assigned by Equation 2 based on its equivalent actionables (α_{eq}) across states. Once a state is chosen, again the actionable with a higher score is chosen for exploration, which helps prioritize unexplored actionables that have high repetition such as navigation links. Equivalence of actionables is established using the fragments that contain them using the Equation 1. Through these equations, every time an actionable is explored, FRAGGEN de-prioritizes all of its equivalent actionables to diversify the exploration. As a result of Equation 1, this de-prioritization often reduces redundant exploration effort spent in already seen functionality.

Figure 5 shows the model inference performed by FRAGGEN for our running example. As it can be seen, FRAGGEN can generate a more precise model, exercising actionables that are considered unique by our fine-grained fragment analysis. For example, FRAGGEN can identify α_{New} in $\mathcal{S}_3$ to be equivalent to that of $\mathcal{S}_1$ and doesn't consider it to be unexplored, thus *terminating* the exploration. On the other hand, as shown in Figure 3, existing techniques exercise α_{New} on $\mathcal{S}_3$ as well, leading to the creation of $\mathcal{S}_4$ and $\mathcal{S}_5$ and necessitating further exploration of the same actionables in the newly added states as well. In fact, without identifying the equivalence of such repeated actionables, state exploration will never *terminate* even for this simple example app we present in the paper. Existing techniques would require a stopping criterion such as a time
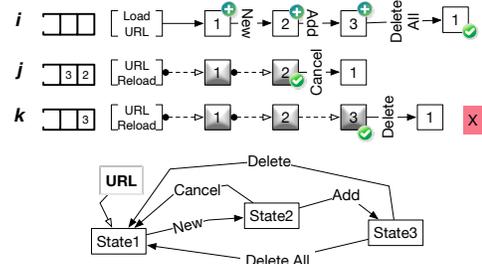

Fig. 5: Model inferred by FRAGGEN

limit to end the inference process and will likely generate an imprecise app model with a high number of near-duplicates.

Configuration options. *constant c_0 in Equation 2*: c_0 determines the weight given to the presence of duplicates for an action that is yet to be explored. In our experiments, we used the default value which is set to 1. We provide the ability to tune it if necessary for specific web applications. The value of constant c0 in Equation 2 can impact prioritization of states, where, a high value could delay exploration of a potentially more interesting newly discovered state. This scenario is possible if the number of duplicates for a single unexplored actionable outnumber the total actionables in a newly discovered state. In this scenario, the older states containing this particular unexplored action would get a higher priority even if a majority of the actions in the state itself are already explored. However, once any one of the instances of the particular actionable is exercised, the value of constant c_0 becomes irrelevant.

Termination: We designed FRAGGEN to be configured to limit the amount of duplication in order to avoid generating inflated models. By default, FRAGGEN does not exercise any actionable which is found to be similar to an already explored actionable. It will *terminate* once all unique actionables are exercised. However, this duplication in exploration effort can sometimes be necessary to fully explore app functionality. Therefore, we provide configuration to continue exploration using unexplored actionables that were previously skipped because they were similar to already explored actionables. FRAGGEN would continue to use Equation 2 and Equation 3, for prioritization.

3.2.2 Memoization and data-fluid fragments

During the model inference process, recall that after every exploration step, FRAGGEN decides if the resulting browser state should be retained in the model by comparing it to the existing states using the $Classify$ function. Our exploration strategy also relies heavily on the $Classify$ function to search for diverse actions as seen in Equation 1.

From algorithm 1, it can be seen that $Classify$ for a given pair of fragments is recursive in nature that may require comparing the corresponding child fragments. To improve its performance, we apply a classic technique for recursive algorithms called memoization, which aims to avoid recalculation of results for sub-problems. Here, instead of storing just the comparison results for every pair

of fragments, we employ a map-based implementation with unique fragments as map keys and a list of their duplicates as the values.

Whenever a new state is added to the model, we update the map by comparing each of the fragments in the new state to existing unique fragments. If any of the fragments in the new state is not a clone (Cl) of existing unique fragments, it is added to the map as a new unique fragment. Otherwise, it is added as a duplicate of the unique fragment it is found to be a clone of.

Further, we utilize this memoization technique using the map of fragments to identify *data-fluid* fragments. A fragment is data-fluid if it is found to be an Nd_2-*data* near-duplicate of at-least one other fragment in the map. Intuitively, we are trying to recognize the parts of the dynamic DOM, specific to the web app being tested, that are likely to have data changes. Any clone of a data-fluid fragment is also data-fluid. For example, in Figure 4, $\mathcal{F}_6$ in $\mathcal{S}_3$ and $\mathcal{F}_6$, $\mathcal{F}_7$ in $\mathcal{S}_5$ are identified to be *data-fluid*.

In practice, there could be multiple causes for changing data in fragments; such as a DOM element that displays current time. Indeed, this ability to mutate dynamic DOM in real-time is the primary reason that developers are able to create highly interactive and responsive web apps. However, such changes also pose a challenge to web testing in the form of fragility of the automatically generated test oracles which fail in response to every change in the dynamic DOM. In the next section, we describe how FRAGGEN mitigates this challenge by making use of memoization and data-fluid fragments to generate robust test oracles.

3.3 Test Generation

UI Test cases, such as the one shown in Listing 1, are sequences of UI events derived from inferred web app models as discussed in subsection 2.1. We designed FRAGGEN to generate tests using *exploration paths* which are essentially inferred model iterations as shown in Figure 5, and contain at least one unexplored action in each new path. Recall (from subsection 2.1) that a model iteration starts by reloading the URL and ends when the browser state does not have any unexplored actions remaining. As the exploration paths or model iterations cover all states and transitions, test cases generated by FRAGGEN also cover all the states and transitions in the model.

Each exploration path translates to a test case that starts with loading URL, similar to the test case shown in Listing 1. For each transition ($\mathcal{A}$ -> ($\mathcal{S}_{src}$, α, $\mathcal{S}_{tgt}$)) in an exploration path ($\mathcal{P}$, definition 4), we generate a test step to perform the action (α) and generate assertions that compare corresponding recorded model states with the browser states by using the function $Classify$ (algorithm 1).

In the remainder of this section, we describe the challenges in generating robust test assertions and how FRAGGEN utilizes fragment analysis to tackle this challenge.

3.3.1 Test assertions

An important but often neglected aspect of regression testing is the creation of effective test assertions or test oracles, which are key in detecting app regressions. Existing test generation techniques generate fragile/brittle test assertions [36], which results in many false positive test failures.
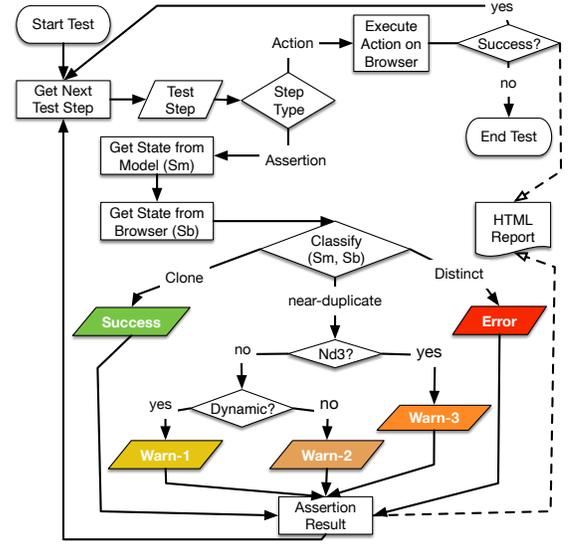

Fig. 6: Test execution flowchart. The function $Classify$ is defined in algorithm 1

Generating web test assertions is considered difficult [37], primarily because they require a reliable state comparison to handle near-duplicates. Previously, researchers have relied on manual specification of DOM invariants [7] specific to a web app. Although effective, such a procedure is not scalable for complex large applications and requires significant human effort as well as domain knowledge.

Our test assertions use the fragment-based state abstraction to identify changes between recorded states in the model and the states observed during test execution. We designed FRAGGEN's test execution to utilize data-fluid fragments identified during model inference to assign importance to the changes detected by our state comparison. As a result, FRAGGEN can produce fine-grained warnings with different levels of severity based on the characteristics of detected changes. Test failures can then be declared based on the severity of warnings, reducing false positive test failures or test breakages caused by unimportant changes.

3.3.2 Test execution

Figure 6 shows the flowchart of our test execution. For each action step, that exercises actionable (α) in the transition ($\mathcal{A}$), the test execution continues if successful or stops otherwise. For each assertion step, we invoke our fragment-based comparison ($Classify$ in algorithm 1) between the browser state ($\mathcal{S}_b$) and the model state ($\mathcal{S}_m$) specified for the particular test step. Based on the output of $Classify$ and using *memoization*, FRAGGEN can create three levels of warnings in addition to success and error for each assertion. Warn-3 has the highest severity while Warn-1 is the lowest. To the best of our knowledge no prior technique exists to generate such warnings for UI test cases. Existing techniques are only capable of creating a binary decision for a given test oracle and generate either an error or a success status.

Using data-fluid fragments. We lower the severity of a warning from Warn-2 to Warn-1 when we detect that the

7

changes observed belong to a fragment that is data-fluid. For example, a web element displaying current time on a web page changes in every concrete instance of the web page and should be given a lower severity warning compared to a change to the title of a web page, which has never been found to change during model inference. It is important to recognize that both these changes would just be textual changes in the DOM for any state comparison technique.

Configuration. In our experiments, we consider any warning other than Warn-1 to be an assertion failure. However, we provide the capability to configure the assertion failures based on warning levels. Our empirical evaluation assesses how the usage of memoization and identification of data-fluid fragments can help make our assertions robust to near-duplicates while retaining the ability to detect bugs.

Report. FRAGGEN also generates an HTML test report that shows the test execution results for easier manual analysis. The report provides a visualization of changes and their severity. Our replication package contains a test report for every test execution in the evaluation.

3.4 Implementation

For page fragmentation, existing implementations of VIPS rely on web rendering frameworks that are no longer maintained and render modern pages incorrectly. Therefore, we ported one implementation [38] to WebDriver API in Java, which is also used by CRAWLJAX so that pages rendered on modern browsers can be fragmented.

We use our ported version of the VIPS [31] algorithm for decomposing a web page into fragments, and compare the structural and visual aspects of the fragments using APTED [34] and Histogram [25] respectively.

We modified the latest version of CRAWLJAX [9] to employ our state abstraction for state equivalence and use our exploration strategy. Finally, our test generator is implemented as a CRAWLJAX plugin and generates JUnit test cases using the WebDriver API. Our tool, FRAGGEN, is publicly available [13].

4 EVALUATION

Our empirical evaluation aims to assess FRAGGEN through (1) its ability to detect near-duplicates, (2) the adequacy of its inferred web app models, and (3) the suitability of its generated tests for regression testing. We do so by answering the following research questions.

RQ$_1$: *How effective is FRAGGEN in distinguishing near-duplicates from distinct states compared to current whole-page techniques?*

RQ$_2$: *How do the models generated by FRAGGEN compare to the models generated by current techniques?*

RQ$_3$: *Are the tests generated by FRAGGEN suitable for regression testing?*

- **RQ$_{3a}$:** *How do the generated test cases perform in regression testing scenarios?*
- **RQ$_{3b}$:** *How effective and tolerant are the generated test oracles?*

In RQ$_1$, we assess the effectiveness of competing techniques in detecting near-duplicates through state-pair classification. In RQ$_2$, we measure the model quality which

TABLE 2: Experimental subjects

App	v_0	v_1	Framework	LOC
addressbook [14]	*8.2.5*	*8.2.5.1*	*PHP, JavaScript*	32K
petclinic [39]	*6010d5*	*4aa89ae*	*Java, Spring MVC*	6K
ppma [40]	*0.6.0*	*0.5.2*	*Yii, JavaScript*	556K
dimeshift [41]	*261166d*	*44089e*	*Backbone.js, JQuery*	10K
claroline [42]	*1.11.10*	*1.11.9*	*PHP, JavaScript*	340K
phoenix-trello [43]	*60c874d*	*c1cdf30*	*Phoenix, Elixir, ReactJS*	5K
pagekit [44]	*1.0.16*	*1.0.14*	*Symfony, Vue.js*	275K
mantisbt [45]	*1.1.8*	*1.2.1*	*PHP, JavaScript*	120K

directly influences the completeness and redundancies in the generated tests. In RQ$_3$, we assess the suitability of generated test suites in regression test scenarios by measuring their reliability in addition to the effectiveness in detecting app changes. For RQ$_{3a}$, we execute the generated tests on the same web app version but different platform/browser versions to analyze their usability in regression test scenarios, and then evaluate their effectiveness in detecting app changes for a different version of the web app. Finally, for RQ$_{3b}$, we perform mutation analysis of recorded web states to assess the effectiveness and robustness of generated test oracles.

4.1 Subject Systems

To address our research questions, we needed to manually analyze captured application states which requires a certain level of control over app behaviour and ability to replicate app states under similar experimental conditions. In addition, for RQ$_1$ and RQ$_2$, we require a manually labelled ground-truth for state-pair classification and unique states in a web app respectively. To this end, we selected eight open-source web apps with an available ground-truth [46] and used in prior web testing research [47], [48], [6], [49], [10]. Our eight subjects shown in Table 2 cover a diverse set of several popular back-end and front-end web app frameworks such as *Symfony, Yii, Spring MVC, Backbone.js, Vue.js, Phoenix/React, JQuery, Bootstrap* and *AngularJS*. Some of our subjects such as *MantisBT, Claroline, PageKit* are quite complex and immensely popular with a sizeable active user base. For example, *Claroline* is an award winning learning management system (LMS) used in more than 100 countries and is available in 35 languages; *MantisBT* is one of the most popular open source issue tracking systems in active use.

4.2 Competing techniques

Based on the results of our recent empirical study [10], we choose two of the best whole-page techniques, (1) RTED [21], a DOM tree differencing technique which inferred the best models on an average, and, (2) Histogram [25], which was the best performing visual algorithm outperforming RTED on several subjects. From here on, we refer to RTED as *Structural* and Histogram as *Visual* when presenting our evaluation.

5 STATE-PAIR CLASSIFICATION (RQ$_1$)

To address RQ$_1$, we compare FRAGGEN to whole-page techniques (subsection 4.2) in terms of their ability to classify a given state-pair as either clone, near-duplicate, or distinct.

TABLE 3: Manually classified state-pair dataset

Subject	Pairs	Clones	Near-Duplicates			Distinct
			Nd_2	Nd_3	Total	
addressbook	8515	26	52	2295	2347	6142
petclinic	11175	2	1433	180	1613	9411
claroline	17766	2707	69	2	71	14988
dimeshift	11628	375	570	0	570	10683
pagekit	9730	0	904	3044	3948	5782
phoenix	11175	1	25	4580	4605	6569
ppma	4851	64	467	0	467	4320
mantisbt	11325	2	1117	0	1117	10206
Total	86165	3177	4637	10101	14738	68101

(a) Finding optimal thresholds using classification F_1

(b) Correctly classified state-pairs for optimal thresholds

Fig. 7: State-pair classification results on the dataset

5.1 Procedure and Metrics

Dataset. We use an existing dataset (Table 3) of 86,165 state-pairs that are manually labeled as either clone (Cl), Nd_2-*data*, Nd_3-*struct*, or distinct (Di) as ground truth. It contains 4,637 Nd_2-*data* and 10,101 Nd_3-*struct* near-duplicates along with 3,177 clones and 68,101 distinct state-pairs belonging to eight subjects used in this work. For the two competing whole-page techniques, *structural* and *visual*, the dataset also includes the computed distance between the two states for each of the 86,165 state-pairs.

Comparison metrics. We use two metrics to compare the competing techniques, 1) the classification F_1 score and 2) the number of detected near-duplicates. We designed our experiment so that, given a state-pair, each technique classifies it to be either clone (Cl), near-duplicate (Nd) or distinct (Di). Using the manually classified ground-truth, we then compute the multi-class classification F_1 score which is the average of F_1 over all three classes (Cl, Nd, Di).

Whole-page technique configuration. The whole-page techniques by default only output a distance between the two states of a given state-pair. In order to compute the output *class* for such techniques, we follow previous work [10] which defined a function Γ as :

$$\Gamma(S_1, S_2, \mathcal{W}, t_c, t_n) \begin{cases} Cl & : \mathcal{W}(S_1, S_2) < t_c \\ D & : \mathcal{W}(S_1, S_2) > t_n \\ Nd & : otherwise \end{cases}$$

Γ takes as inputs a whole-page comparison technique $\mathcal{W}$ which computes the distance between two given states in a state-pair (S_1, S_2) and outputs a *class* using a pair of thresholds t_c, t_n. Γ classifies a state-pair to be a clone if the computed distance falls below the threshold t_c, distinct if it is above t_n or near-duplicate otherwise. Therefore, the choice of thresholds t_c, t_n, can play a huge role in the overall effectiveness of the techniques.

Threshold determination for whole-page techniques. To obtain optimal thresholds that maximize the classification scores of the two whole-page techniques, using the labelled dataset (shown in Table 3) as ground-truth, we employ bayesian optimization [50] to search for thresholds that maximize the multi-class classification F_1 score for Γ. We ran the optimization technique for 10,000 iterations or trials and retrieved the thresholds that provided best F_1 score to be optimal thresholds. In each trial, the optimizer chooses a pair of thresholds from the sample space of distances possible for the corresponding whole-page technique and computes the F_1 score.
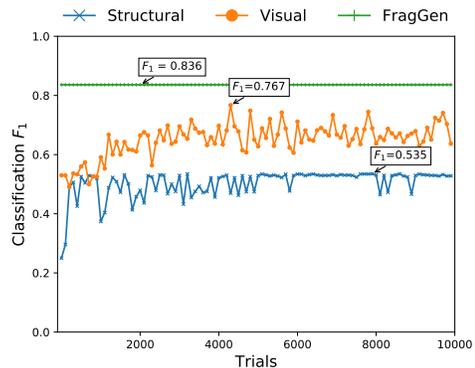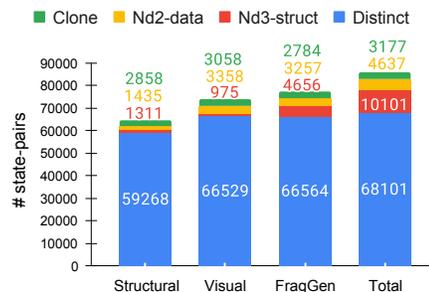
FRAGGEN configuration. For each state-pair in the dataset, we invoke the $Classify$ function (algorithm 1) which outputs one of the four classes (Di, Cl, Nd_2-*data*, Nd_3-*struct*). For this experiment, we combine the two outputs Nd_2-*data* and Nd_3-*struct* into a single class Nd to compare with the output of Γ.

5.2 Results

Figure 7a shows the classification F_1 being optimized for the structural and visual whole-page techniques by a bayesian optimizer in 10,000 trials. To make the data comprehensible, we divided the 10,000 trials into 100 intervals and plotted a single data-point per interval in Figure 7a. We chose the maximum F_1 found in the corresponding interval of 100 trials as the representative data point of the interval.

As seen in the Figure 7a, FRAGGEN with an F_1 score of 0.836 performs better than the two whole-page techniques. FRAGGEN's F_1 score is 55% better than structural technique and 9% better than visual technique, which have scores of 0.535 and 0.767 respectively.

Figure 7b shows the correctly classified state-pairs for each of the four state-pair classes in the dataset for the optimal thresholds. FRAGGEN correctly classifies the highest number of Nd_3-*struct* as well as distinct (Di) state-pairs, while the visual technique is the best at classifying clones (Cl) and Nd_2-*data* state-pairs. Structural is the worst performer overall with a slightly better Nd_3-*struct* detection

9

TABLE 4: F_1 of inferred models for 60 minute crawls

	addressbook	claroline	dimeshift	mantisbt	pagekit	petclinic	phoenix	ppma	Average
Structural	0.44	0.83	0.26	0.42	0.40	0.79	0.38	0.23	0.47
Visual	0.42	0.71	0.50	0.28	0.47	0.93	0.19	0.21	0.46
FRAGGEN	0.89	0.82	0.74	0.90	0.74	1.00	0.59	0.71	0.80
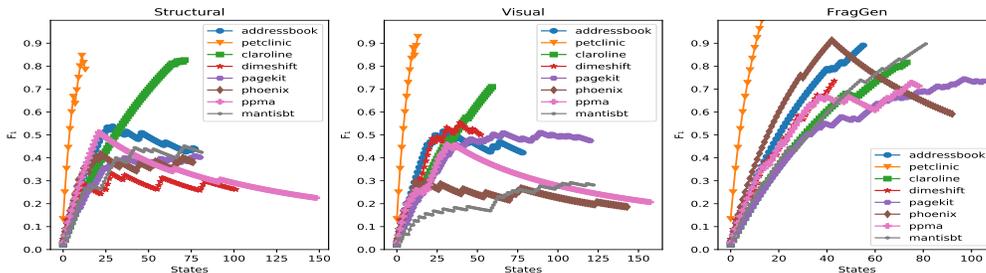
Fig. 8: F_1 as states are detected and added to the model.

TABLE 5: Comparison of inferred models (*eight subjects*)

SAF	Model Quality			Labelled Web States				Termination
	Pr	Re	F_1	Unique	Near-Duplicates			
					Nd_2	Nd_3	All	
Structural	0.46	0.51	0.47	249	44	282	326	2
Visual	0.45	0.52	0.46	244	325	116	441	1
FRAGGEN	0.79	0.83	0.80	411	90	29	119	4

compared to visual. Overall, FRAGGEN correctly classifies 77,261 state-pairs while the visual and structural techniques correctly classify 73,920 and 64,872 state-pairs respectively.

Out of the 14,738 near-duplicates in the dataset, FRAGGEN detects 7,913 which is 188% and 82% better than the structural and visual techniques that detect 2,746 and 4,333 near-duplicates respectively. The result is significant because the whole-page techniques performed worse despite having an obvious advantage of classifying the same dataset that is being used for optimization. As the previous study [10] shows, given a set of state-pairs from an unseen web app, the whole-page techniques are unlikely to perform at a similar level using the same thresholds.

Indeed, the study delves deeper into the threshold selection for generating optimal models; it concludes that 1) the thresholds should be chosen specific to each web app and 2) a ground-truth should be created for any new web app in order to obtain optimal thresholds. Even after such an optimization per web app, as we illustrate with an example in subsection 2.2, reliance on thresholds creates an inherent limitation for whole-page techniques in separating the Nd_3-*struct* near-duplicates from distinct state-pairs.

In RQ_2, we investigate the quality of models inferred by each of the competing techniques configured with the optimal thresholds. We obtained optimized thresholds that generate the best model for each of the subjects from the empirical study [10].

6 MODEL INFERENCE COMPARISON (RQ₂)

6.1 Procedure and Metrics

Model quality. We measure the quality of an inferred model in terms of its coverage of the app state-space, recall (Re), and amount of duplication, precision (Pr), by manually analyzing it with reference to a ground truth model for the web app using a methodology established in prior research [10]. The ground truth models for our subject apps taken from a published dataset [46] represent the functionality of a given web app using a minimal set of states and transitions. A ground truth model is a set of unique states. To analyze a given inferred model, each state in the inferred model is manually *mapped* to one of the ground truth states. A state (S_m) in the inferred model is mapped to a state (S_g) in the ground truth if the manual classification of the state-pair (S_m, S_g) is either clone, Nd_2-*data* or Nd_3-*struct*. Every ground truth state that is mapped to at-least one state in the inferred model is *covered* by the inferred model. We then compute Pr of the inferred model as the ratio of *covered* ground truth states to the total states in the inferred model, and Re as the ratio of *covered* ground truth states to the total number of ground truth states.

Experiment set-up. For each of our subject apps, we generate models using each technique by setting a maximum exploration time of 1 hour to be the *stopping criteria*. We use *Google Chrome (v82.0)* browser, and reset the subject app after every crawl to remove any back-end changes done by previous run. For a fair comparison, we configure each technique to use exactly the same crawl rules (e.g., form fill data). A technique can also *terminate* exploration before the stopping criterion is invoked by the crawler. Table 5 shows the number of times a technique terminated on its own.

6.2 Results

Table 5 shows the overall statistics for all the eight subject apps for FRAGGEN compared to structural and visual techniques. For recall, on average, FRAGGEN covered 83% of the state space, which is 60% higher than visual, the next best technique. For precision, at 79%, FRAGGEN produced models with 71% higher precision than structural, which itself
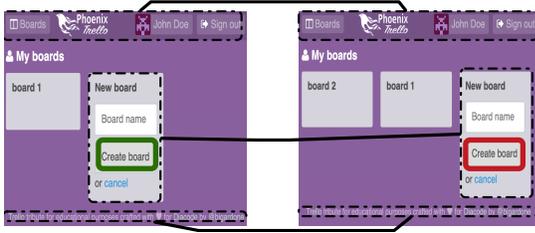
Fig. 9: Detected near-duplicates in phoenix

TABLE 6: Regression test execution set-up

Execution	app_version	Platform	Browser
Crawl	v_0	MacOS-14 [51]	chrome-82
TestSuite ϵ_1	v_0	MacOS-14	chrome-83
TestSuite ϵ_2	v_0	RHEL-7 [52]	chrome-84
TestSuite ϵ_3	v_1	RHEL-7	chrome-84

Fig. 10: Test breakages on the same app version (ϵ_1, ϵ_2)

performed slightly better than visual. In total, FRAGGEN added only 119 near-duplicates in all models, compared to the 326 by structural and 441 by visual, which are nearly 3 and 4 times higher, respectively. FRAGGEN also discovered 411 unique states overall, while the existing whole-page techniques detected 37% less app states in aggregate terms.

Overall, the F1 measure of FRAGGEN is 0.80, while that of structural and visual whole-page techniques are 0.47 and 0.46, respectively. Table 4 shows the details for all the subject systems. FRAGGEN consistently produced models with better F_1 except for Claroline where structural technique's model is marginally better.

When F_1 of the models is plotted against the states being added to the model, as shown in Figure 8, it can be seen that the F_1 score for existing techniques does not improve after an initial exploration period. As a result, the final models generated by existing techniques can sometimes deteriorate over time as more and more near-duplicates are added. FRAGGEN, however, keeps improving the model quality when given more time as it diversifies the exploration to discover unseen states while avoiding the addition of near-duplicates to the model.

One of the main reasons for this trend is that existing techniques exercise similar actions repeatedly, and in dynamic web apps, this often results in the creation of near-duplicates and infinite loops, as mentioned in subsection 2.2. Consider a real example from our subject Phoenix, where the action "create board" is available in two model states as shown in Figure 9. In both states, the action is functionally similar as it creates another board and therefore, need not be exercised more than once. However, existing techniques, which rely on whole-page comparison, cannot infer such similarities and keep creating boards and Nd₃-*struct* near-duplicate states as they repeatedly exercise the "create board" action.

On the other hand, FRAGGEN is able to identify similar actionables through fine-grained fragment-based analysis to avoid creating Nd₃-*struct* states and successfully diversify the exploration. Indeed, as the Table 5 shows, FRAGGEN added just 29 Nd₃-*struct* states overall while *structural* and *visual* techniques added 282 and 116 respectively.

The trend is prominently noticeable in Figure 8, where FRAGGEN is able to improve the model continuously by avoiding addition of Nd₃-*struct* near-duplicates to the model and seeking out unique actionables to exercise. As a result, FRAGGEN also terminates exploration for 4 out of 8 subjects apps as shown in Table 5, whereas structural and visual techniques terminate only 2 and 1 times respectively.
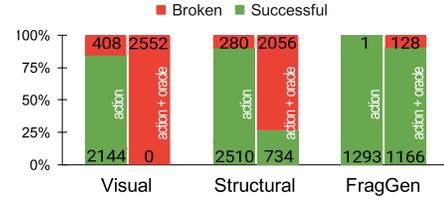
7 REGRESSION TESTING SUITABILITY (RQ₃)

As part of assessing the suitability of generated tests in regression scenarios, in RQ₃ₐ, we execute generated tests on different browser/platform and app versions, in order to evaluate the model inaccuracies and fragility of test cases in addition to the effectiveness of techniques in detecting application changes. In RQ₃ᵦ, we further evaluate the robustness and effectiveness of test oracles by simulating evolution through mutation of recorded web states.

7.1 Test Breakages (RQ₃ₐ)

7.1.1 Procedure and Metrics

The goal of our evaluation in RQ3a is to assess how FRAGGEN compares against existing techniques that generate regression tests automatically. There are not many techniques available that generate regression tests for web applications currently. We compare tests generated by FRAGGEN against tests generated by CRAWLJAX, which also employs exploration paths to generate tests from the crawl model. While FRAGGEN uses its fine-grained fragment analysis to generate test oracles, for the whole-page techniques, CRAWLJAX is configured to generate test oracles that use the same whole-page comparison as the one used for model inference, namely, structural (RTED) and visual (Histogram).

Table 6 shows the three regression test scenarios used in our evaluation, where in ϵ_1 and ϵ_2, we execute tests on the same app version but vary browser/platforms from the crawl, allowing us to evaluate the validity and robustness of the generated test suite. We then execute tests (ϵ_3) on a different version of the web app to determine the effectiveness of the test suites in detecting real application changes.

7.1.2 RQ₃ₐ Results

As seen in Table 7, whole-page techniques generated test actions that succeeded only 86.2% on an average. A potential cause for failure of remaining nearly 14% actions on the same version could be limitations in handling the near-duplicates for existing techniques. As Figure 10 show, breakage of these test actions resulted in breakage of 16% and 10% of tests in visual and structural test suites respectively without even considering the test oracle fragility. When test oracles are considered in declaring test breakages,

TABLE 7: Regression test run results

Test Execution	Whole-page techniques											FRAGGEN			
	Visual				Structural				Average						
	ϵ_1	ϵ_2	ϵ_3	*All*	ϵ_1	ϵ_2	ϵ_3	*All*	v_0	v_1	*All*	ϵ_1	ϵ_2	ϵ_3	*All*
Action Success %	83.6	83.2	63.0	76.6	89.1	88.8	69.0	82.3	86.2	66.0	79.5	99.5	99.9	93.3	97.6
Oracle Success %	0.0	0.0	0.0	0.0	53.1	51.3	22.8	42.4	26.1	11.4	21.0	98.6	97.5	64.8	87.0

all visual test oracles fail during test executions causing 100% test breakage, while 52% structural oracles fail breaking nearly 74% of tests in the same app version.

In contrast, FRAGGEN can execute test actions on the same version of the web app with nearly 100% success rate. As Table 7 shows, in ϵ_1, ϵ_2, FRAGGEN's test oracles also have a 98% success rate on average, showing greater adaptability to changing execution environments compared to the whole-page techniques.

(a) Test action failures (ϵ_3)

(b) Random test oracle failures (ϵ_3)
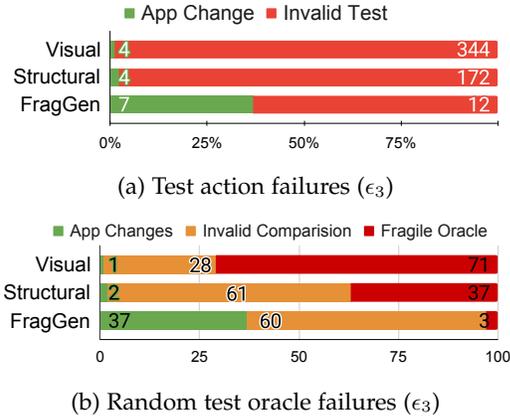
Fig. 11: Manual analysis of failed test actions and oracles on a different version of the web apps

Next, we evaluate test suite effectiveness in detecting app changes on a different version of web app using the test execution (ϵ_3). We follow the standard web testing practice where tests created for an existing version (v_0) are executed on a new version (v_1) of web app in order to detect faults through manual analysis of test failures. We mitigate the potential bias in manual analysis of test failures by labelling *all* observed application changes to be faults.

When we manually analyzed all test failures in ϵ_3 as shown in Figure 11a, we found that both structural and visual test suites detected 4 app changes each and produced 172 and 344 false positives respectively. In contrast, FRAGGEN was able to detect 7 app changes with only 10 false positives, significantly reducing the manual effort required to identify app changes and maintain test suites.

We then manually analyzed 100 randomly selected oracle failures for each technique to label the cause of failure. In this manual analysis, if the test state matches the expected state, we categorize the failure as either an *app change* if we notice a change in application behaviour or a *fragile oracle* as the failure represents a limitation of state comparison. Our analysis depicted in Figure 11b, shows that 71 visual and 37 structural oracles failed because of fragility in state comparison, whereas FRAGGEN generated only 3 fragile test oracles. Of the 100 examined failures, while the structural and visual oracles only detected 2 and 1 app changes respectively, FRAGGEN could detect 37 app changes.

When the test state and expected state do not match for the failure being analyzed, we categorize it as an invalid comparison, where typically test execution does not reach the target app state due to failure of earlier test actions or invalid test sequences due to changed app behaviour. Further root cause analysis of such failures is difficult even with domain knowledge, which makes it challenging to directly evaluate test oracle robustness through regression test executions. We mitigate this problem in RQ$_{3b}$, where the test oracles are compared on a synthetic dataset generated through mutation of recorded web pages.

These results show that the tests generated by FRAGGEN are significantly more robust than existing techniques in regression testing scenarios involving the same version of the web app, and show greater efficacy in detecting evolutionary app changes with significantly lower manual effort.

7.2 Effectiveness and Robustness of Test Oracles (RQ$_{3b}$)

In RQ$_{3b}$, we simulate evolutionary app changes using mutation analysis to evaluate the suitability of generated test oracles for regression testing.

7.2.1 Mutation Analysis Methodology

We define four mutation operators for DOM nodes as shown in Table 8 in order to generate mutant web pages. Similar mutation operators for GUI artifacts have been proposed in prior web [53], [54], [55], Android [56] and GUI [57], [58] mutation analysis research. Since RQ$_{3b}$ only aims to evaluate test oracles, we find this approach of mutating recorded web pages to be adequate and more flexible, instead of mutating the actual web app source code as prior work does.

Given a crawl model and a test execution trace, our mutation analysis tool applies a random mutation to one of the model states and compares the mutant to the corresponding states in the test trace using each of the three competing techniques.

We partition the generated mutants into either app changes or irrelevant to functionality based on their visibility in the GUI. We categorize (1) the mutants not visible on the page as equivalent mutants that test oracles should tolerate and ignore, and (2) all visible mutations as app changes that should be detected by effective test oracles. As Table 9 shows, we use 3,042 Visible mutations to compute effectiveness. Robustness is computed using 6,013 equivalent mutants that are of three kinds : *None* - where no mutation has been applied on the state; *AttributeMutator* -

TABLE 8: Mutation operators for DOM nodes

Type	Tags or Attributes	Description	Example *(original : mutant)*
Attribute	{id, class, title}	Modifies mentioned attribute value of any node.	[<div id="a">] : [<div id="aMut">]
Tag	{span, h1-h6, p}	The tag name is changed to a similar tag.	[xyz] : [xyz]
Subtree	{div, table, tr, td, ul, li, p}	Deletes all children of selected container node.	[<tr><td>xyz</td></tr>] : [<tr></tr>]
Text	{h1-h6, p, b, i, }	The text content of selected leaf node is changed.	[<h1>abc</h1>] : [<h1>abcMut<h1>]

TABLE 9: Effectiveness and Robustness of test oracles

Score	Mutation		Structural	Visual	FRAGGEN	
	Type	#			*No-Mem*	*With-Mem*
	Visible	*3042*	2235	2908	3001	3001
Effectiveness			**73.5**	**95.6**	**98.7**	*98.7*
	None	*4620*	198	1964	921	136
	Attribute	*770*	10	585	166	157
	Invisible	*623*	553	577	275	274
	Total	*6013*	761	3126	1362	567
Robustness			**87.3**	*48.0*	**77.3**	**90.6**

Fig. 12: Effectiveness vs Robustness of test oracles

an attribute is mutated; and, *Invisible* - where a hidden web element is mutated.

Note that we recognize the possibility of invisible DOM or attribute changes affecting web page functionality. However, we expect that such changes manifest in failure of test actions, which is outside the scope of RQ$_{3b}$, as our set-up does not execute test actions on these mutated states. Equivalent mutants of the kind *None* may still have DOM changes that result from a variety of reasons such as auto-generation of pages in back-end. We still consider them to be equivalent because the changes do not relate to either 1) change in source code which was unmodified or 2) change in test execution which remained consistent.

Finally, we consider a test oracle to be suitable for regression testing if it has high robustness as well as effectiveness.

FRAGGEN configuration. As described in subsubsection 3.2.2, during model inference, FRAGGEN identifies data-fluid fragments and employs this knowledge to generate test oracles. In Table 9, *No-Mem* refers to test oracles if only fragment-based comparison is used; and *With-Mem* shows the test oracles that employ knowledge of data-fluid fragments. As shown in Figure 6, without the knowledge of data-fluid fragments, we cannot lower the severity of warnings for changes that are typical for the web app under test and therefore may not necessarily be application bugs.

7.2.2 Results

As Table 9 shows, our mutation analysis experiment resulted in a total of 3,042 randomly generated visible mutations that should be detected by the oracles and 6,013 equivalent mutants that the oracles should be tolerant to.

Overall, amongst the existing techniques, *visual* test oracles detect 95.6% of visible mutations while *RTED*, the structural comparison technique which only considers HTML tags for equivalence, is tolerant to 87.3% of equivalent mu-

tants encountered during the experiment. However, visual oracles are fragile, with a robustness score of only 48%, whereas structural oracles fail to detect 26.5% of the visible mutants.
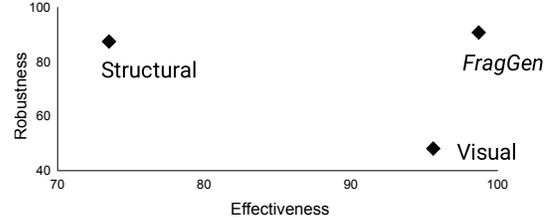
In contrast, test oracles generated by FRAGGEN are able to detect 98.7% of visible mutants in both the configurations. When *memoization* knowledge is applied in the *With-Mem* configuration, FRAGGEN has the highest robustness of all the competing techniques by detecting 90.6% of the equivalent mutants. In the *No-Mem* configuration, when only fragment-based state comparison is used, the robustness drops to 77.3%, which is still 61% more than visual. The difference in numbers between the two configurations is because of handling the *None* category, where no mutation has been applied to the recorded web pages. It can be concluded therefore, that FRAGGEN is able to make use of memoization and data-fluid fragments to improve the robustness of test oracles.

As the Figure 12 shows, existing whole-page oracles can either be highly effective or highly tolerant but cannot balance both the aspects, causing either a large number of false positives or false negatives respectively. FRAGGEN's test oracles on the other hand, use fine-grained fragment analysis with memoization to outperform existing techniques in both effectiveness and robustness at the same time, making them suitable for regression testing of modern web apps.

8 DISCUSSION

Granularity of fragments. During our experiments, we found that determining equivalency of certain fragments that are too small such as $\mathcal{F}_3$, shown in dotted lines in Figure 4, using DOM or visual characteristics results in drawing false equivalence between semantically different web pages. In order to avoid this, using a predefined threshold, FRAGGEN prunes out the smaller fragments and does not use their equivalence in determining page similarity.

One such example is shown in Figure 13, where each individual tag can have two states – selected or not. Therefore each tag can be semantically equivalent to other depending on this state. However, as we avoid comparing the smaller
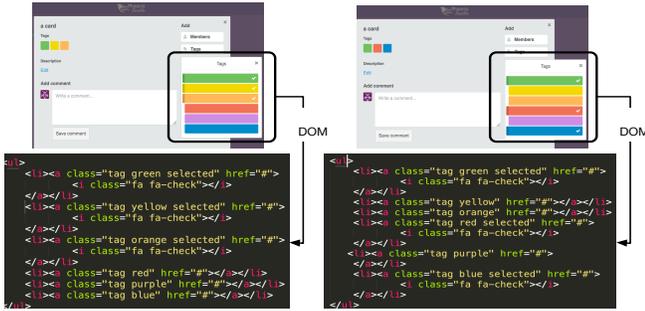
Fig. 13: Undetected near-duplicate fragments in phoenix

fragments of individual tags, we cannot determine equivalency of parent fragments which differ structurally upon selection of different tags. The number of possible near-duplicates we cannot detect in such cases can explode in number, and is the reason for the reduced precision of the web app model for *phoenix*.

Even-though smaller fragments with even one web element can be separate functional entities, such fine-grained comparison needs semantic inference beyond simple DOM and visual characteristics, which we leave for future work.

Test oracle generation. FRAGGEN currently identifies data-fluid fragments assuming that during model inference, only the events fired by the crawler can cause persistent changes. Such assumption can fail on live web applications where multiple users can concurrently induce back-end changes. However, since our technique is meant for test environments, we consider our assumption reasonable.

Model Inference in other domains. While our FRAGGEN implementation is specific to web apps, conceptually our approach to model generation based on fragments (instead of whole pages or whole screens) can be applied to other domains such as mobile apps or desktop applications.

Threats to validity. Using a limited number of web apps in a controlled setting in our evaluation poses an external validity threat and further experiments are needed to fully assess the generalizability of our results; we have chosen eight subject apps used in previous web testing research, pertaining to different domains in order to mitigate the threat. Threats to internal validity come from the manual labelling of web pages and changes, which was unavoidable because no automated method could provide us with the required ground truth. To mitigate this threat, we performed the labelling by following a process established in prior work. For reproducibility of our findings, we made our tool publicly available [13] along with usage instructions and used subject systems.

9 RELATED WORK

In web model inference, CRAWLJAX [9] explores dynamic web apps with client-side JavaScript by triggering actions and analyzing the DOM. WebMate [59] crawls web apps using a state equivalence based on actionables in web pages. FRAGGEN is built upon CRAWLJAX framework by adding a new fragment-based state abstraction and modifying several core components that drive state exploration.

For regression test suite generation from web app models, Mesbah et al. proposed test generation based on model coverage [9] while Marchetto et al. [60] used coverage criteria based on semantically interacting events. FRAGGEN uses the exploration paths that cover all the events and states in the inferred model for test generation.

Instead of crawl models, Biagiola et al. use *page objects* derived from crawl models, and employed search-based [61] techniques, diversification [6] of test events to generate test cases. These page objects used in existing work are derived from crawl models generated by the baseline Crawjax that was manually configured with thresholds by the researchers for their experiments. In contrast, we generate test cases directly from crawl models without any manual threshold selection using the paths followed by our exploration strategy during model inference. Although the generation of page objects from crawl models is largely automated, some manual effort such as specifying guards to avoid test dependencies is required for optimal results. Indeed, the superior crawl models generated by FRAGGEN should improve the test cases generated by the existing techniques based on page objects.

Prior research used histogram [17] to detect cross browser differences and manually specified DOM invariants [15] to create robust test oracles. FRAGGEN, however, automatically generates robust fine-grained test oracles at fragment level by leveraging model inference.

To assess the quality of web test suites, prior work mutated source code artifacts such as JSP web pages [55], [62], and client-side JavaScript [54], [53]. We assess the quality of test oracles by mutating the recorded web pages.

Sun et al. [63] employed page fragment-based exploration for efficient information retrieval while we use fragmentation for model inference and test oracle generation.

FeedEx [64] employs a combination of client-side JavaScript code coverage, DOM and Path diversity to prioritize re-exploration of already explored actions. WebExplor [65] by Zheng et al. also proposes a method to prioritize web elements that were already explored by rewarding discovery of new states in previous result of the same action. Both these existing techniques assign the same priority to all unexplored actions. State exploration strategy is a concern for UI testing in general and Degott et al. [66] also prioritize exploratory actions by first exploring all available actions in a mobile app and analyzing the results of execution. In contrast, FragGen is able to make use of fragment based equivalence of unexplored actions to diversify exploration.

To the best of our knowledge, we are the first to employ page fragmentation to establish state equivalence for web app testing. Our novel state comparison combines both structural and visual analysis of the identified fragments to effectively detect near-duplicates. Our tool, FRAGGEN, uses this novel state comparison to infer precise models, as well as generate effective and reliable regression test suites.

10 CONCLUSIONS AND FUTURE WORK

Automated model inference and test generation for complex dynamic modern web apps is a challenging problem because of the presence of near-duplicates that cannot be

detected by whole-page analysis employed by existing techniques. We developed a novel technique, FRAGGEN, which uses smaller page fragments to detect near-duplicates and diversify web app exploration to generate precise models while covering a high percentage of application state space. FRAGGEN is also able to generate oracles that are suitable for regression testing as they are highly effective in detecting visible app changes while being tolerant to minor changes unrelated to functionality.

As part of the future work, we plan to improve our state comparison technique through inference of web page semantics when structural and visual characteristics are inadequate.

REFERENCES

[1] S. Thummalapenta, K. V. Lakshmi, S. Sinha, N. Sinha, and S. Chandra, "Guided test generation for web applications," in *2013 35th International Conference on Software Engineering (ICSE)*, May 2013, pp. 162–171.

[2] D. Roest, A. Mesbah, and A. v. Deursen, "Regression testing Ajax applications: Coping with dynamism," in *International Conference on Software Testing, Verification and Validation (ICST)*, 2010, pp. 127–136.

[3] A. M. Memon and M. L. Soffa, "Regression testing of guis," *SIGSOFT Softw. Eng. Notes*, vol. 28, no. 5, p. 118–127, Sep. 2003. [Online]. Available: https://doi.org/10.1145/949952.940088

[4] M. Grechanik, Q. Xie, and C. Fu, "Maintaining and evolving GUI-directed test scripts," in *Proceedings of 31st International Conference on Software Engineering*, ser. ICSE 2009. IEEE, 2009, pp. 408–418.

[5] L. Christophe, R. Stevens, C. D. Roover, and W. D. Meuter, "Prevalence and maintenance of automated functional tests for web applications," in *2014 IEEE International Conference on Software Maintenance and Evolution*, 2014, pp. 141–150.

[6] M. Biagiola, A. Stocco, F. Ricca, and P. Tonella, "Diversity-based web test generation," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 142–153. [Online]. Available: https://doi.org/10.1145/3338906.3338970

[7] A. Mesbah, A. van Deursen, and D. Roest, "Invariant-based automatic testing of modern web applications," *IEEE Transactions on Software Engineering*, vol. 38, no. 1, pp. 35–53, 2012.

[8] S. Mirshokraie and A. Mesbah, "Jsart: JavaScript assertion-based regression testing," in *International Conference on Web Engineering (ICWE)*, 2012, pp. 238–252.

[9] A. Mesbah, A. van Deursen, and S. Lenselink, "Crawling ajax-based web applications through dynamic analysis of user interface state changes," *ACM Transactions on the Web*, vol. 6, no. 1, pp. 3:1–3:30, 2012.

[10] R. K. Yandrapally, A. Stocco, and A. Mesbah, "Near-duplicate detection in web app model inference," in *Proceedings of the ACM/IEEE International Conference on Software Engineering (ICSE)*. ACM, 2020, p. 12 pages.

[11] M. Henzinger, "Finding near-duplicate web pages: A large-scale evaluation of algorithms," in *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '06. ACM, 2006, pp. 284–291.

[12] M. Leotta, A. Stocco, F. Ricca, and P. Tonella, "Using multi-locators to increase the robustness of web test cases," in *Proceedings of 8th International Conference on Software Testing, Verification and Validation*, ser. ICST 2015. IEEE, 2015, pp. 1–10.

[13] anon, "FragGen VM," https://doi.org/10.5281/zenodo.4007539, 2020, zenodo Software Upload.

[14] PHP AddressBook, "Simple, web-based address & phone book," http://sourceforge.net/projects/php-addressbook, 2015, accessed: 2018-10-01.

[15] A. Mesbah and A. van Deursen, "Invariant-based automatic testing of ajax user interfaces," in *Proceedings of the 31st International Conference on Software Engineering*, ser. ICSE '09. USA: IEEE Computer Society, 2009, p. 210–220. [Online]. Available: https://doi.org/10.1109/ICSE.2009.5070522

[16] G. A. Di Lucca, M. Di Penta, A. R. Fasolino, and P. Granato, "Clone analysis in the web era: an approach to identify cloned web pages," in *Proceedings of the International Workshop of Empirical Studies on Software Maintenance - November 2001 - Florence - Italy*, 2001, pp. 107–113.

[17] S. R. Choudhary, M. R. Prasad, and A. Orso, "CrossCheck: Combining Crawling and Differencing to Better Detect Cross-browser Incompatibilities in Web Applications," in *Proc. of the 2012 IEEE Fifth International Conference on Software Testing, Verification and Validation*, ser. ICST '12. IEEE, 2012, pp. 171–180.

[18] S. Roy Choudhary, M. R. Prasad, and A. Orso, "X-PERT: Accurate Identification of Cross-browser Issues in Web Applications," in *Proc. of the 2013 International Conference on Software Engineering*, ser. ICSE '13. IEEE Press, 2013, pp. 702–711.

[19] S. Mahajan and W. G. J. Halfond, "Detection and Localization of HTML Presentation Failures Using Computer Vision-Based Techniques," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation*, ser. ICST '15, April 2015, pp. 1–10.

[20] S. Mahajan and W. G. Halfond, "Finding HTML Presentation Failures Using Image Comparison Techniques," in *Proc. of the 29th ACM/IEEE International Conference on Automated Software Engineering*, ser. ASE '14. ACM, 2014, pp. 91–96.

[21] M. Pawlik and N. Augsten, "Efficient computation of the tree edit distance," *ACM Trans. Database Syst.*, vol. 40, no. 1, pp. 3:1–3:40, Mar. 2015.

[22] V. Levenshtein, "Binary Codes Capable of Correcting Deletions, Insertions and Reversals," *Soviet Physics Doklady*, vol. 10, p. 707, 1966.

[23] J. Oliver, C. Cheng, and Y. Chen, "TLSH – A Locality Sensitive Hash," in *2013 Fourth Cybercrime and Trustworthy Computing Workshop*, Nov 2013, pp. 7–13.

[24] M. S. Charikar, "Similarity estimation techniques from rounding algorithms," in *Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing*, ser. STOC '02. New York, NY, USA: ACM, 2002, pp. 380–388. [Online]. Available: http://doi.acm.org/10.1145/509907.509965

[25] M. J. Swain and D. H. Ballard, "Indexing via color histograms," in *Active Perception and Robot Vision*, A. K. Sood and H. Wechsler, Eds. Springer Berlin Heidelberg, 1992, pp. 261–273.

[26] B. Yang, F. Gu, and X. Niu, "Block mean value based image perceptual hashing," in *2006 International Conference on Intelligent Information Hiding and Multimedia*, 2006, pp. 167–172.

[27] C. Zauner, "Implementation and benchmarking of perceptual image hash functions," Ph.D. dissertation, 2010.

[28] H. Yee, S. Pattanaik, and D. P. Greenberg, "Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments," *ACM Trans. Graph.*, vol. 20, no. 1, pp. 39–65, Jan. 2001.

[29] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of Seventh IEEE International Conference on Computer Vision*, vol. 2, 1999, pp. 1150–1157.

[30] and A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.

[31] D. Cai, S. Yu, J.-R. Wen, and W.-Y. Ma, "Vips: a vision-based page segmentation algorithm," Tech. Rep. MSR-TR-2003-79, November 2003. [Online]. Available: https://www.microsoft.com/en-us/research/publication/vips-a-vision-based-page-segmentation-algorithm/

[32] J. Kiesel, L. Meyer, F. Kneist, B. Stein, and M. Potthast, "An empirical comparison of web page segmentation algorithms," in *ECIR (2)*, 2021, pp. 62–74. [Online]. Available: https://doi.org/10.1007/978-3-030-72240-1_5

[33] "Document Object Model (DOM)," https://en.wikipedia.org/wiki/Document_Object_Model, accessed: 2022-01-07.

[34] M. Pawlik and N. Augsten, "Tree edit distance: Robust and memory-efficient," *Inf. Syst.*, vol. 56, pp. 157–173, 2016.

[35] "XPath, XML Path Language," https://en.wikipedia.org/wiki/XPath#Abbreviated_syntax, accessed: 2022-01-07.

[36] A. Torsel, "Automated Test Case Generation for Web Applications from a Domain Specific Model," in *2011 IEEE 35th Annual Computer Software and Applications Conference Workshops*, Jul. 2011, pp. 137–142.

[37] E. T. Barr, M. Harman, P. McMinn, M. Shahbaz, and S. Yoo, "The oracle problem in software testing: A survey," *IEEE Transactions on Software Engineering*, vol. 41, no. 5, pp. 507–525, 2015.

[38] T. Popela, "Implementace algoritmu pro vizuální segmentaci www stránek," Master's thesis, Brno University of Technology, Faculty of Information Technology, 2012. [Online]. Available: https://www.fit.vut.cz/study/thesis/14163/

[39] "Angular version of the Spring PetClinic web application," https://github.com/spring-petclinic/spring-petclinic-angular, 2018.

[40] " PHP Password Manager," https://github.com/pklink/ppma, 2018.

[41] " DimeShift: easiest way to track your expenses," https://github.com/jeka-kiselyov/dimeshift, 2018.

[42] "Claroline. Open Source Learning Management System." https://sourceforge.net/projects/claroline/, 2015.

[43] " Phoenix: Trello tribute done in Elixir, Phoenix Framework, React and Redux." https://github.com/jesuvalan/phoenix-trello, 2017.

[44] " Pagekit: modular and lightweight CMS. ," https://github.com/pagekit/pagekit, 2018.

[45] " Mantis Bug Tracker," https://github.com/mantisbt/mantisbt, 2018.

[46] N. D. Study, "Near duplicate study crawls," Sep. 2019. [Online]. Available: https://doi.org/10.5281/zenodo.3385377

[47] A. Stocco, M. Leotta, F. Ricca, and P. Tonella, "Clustering-aided page object generation for web testing," in *Proceedings of 16th International Conference on Web Engineering*, ser. ICWE 2016. Springer, 2016, pp. 132–151.

[48] ——, "APOGEN: Automatic Page Object Generator for Web Testing," *Software Quality Journal*, vol. 25, no. 3, pp. 1007–1039, Sep. 2017.

[49] M. Biagiola, A. Stocco, A. Mesbah, F. Ricca, and P. Tonella, "Web test dependency detection," in *Proceedings of 27th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2019. ACM, 2019, p. 12 pages.

[50] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 2951–2959. [Online]. Available: http://papers.nips.cc/paper/4522-practical-bayesian-optimization-of-machine-learning-algorithms.pdf

[51] " MacOS Mojave ," https://en.wikipedia.org/wiki/MacOS_Mojave, 2020.

[52] " RedHat Linux ," https://en.wikipedia.org/wiki/Red_Hat_Enterprise_Linux#RHEL_7, 2018.

[53] K. Nishiura, Y. Maezawa, H. Washizaki, and S. Honiden, "Mutation analysis for javascript web applications testing," vol. 2013, 01 2013.

[54] S. Mirshokraie, A. Mesbah, and K. Pattabiraman, "Guided mutation testing for javascript web applications," *IEEE Transactions on Software Engineering*, vol. 41, no. 5, pp. 429–444, May 2015.

[55] U. Praphamontripong and J. Offutt, "Applying mutation testing to web applications," in *2010 Third International Conference on Software Testing, Verification, and Validation Workshops*, April 2010, pp. 132–141.

[56] M. Linares-Vásquez, G. Bavota, M. Tufano, K. Moran, M. Di Penta, C. Vendome, C. Bernal-Cárdenas, and D. Poshyvanyk, "Enabling mutation testing for android apps," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2017. New York, NY, USA: Association for Computing Machinery, 2017, pp. 233–244. [Online]. Available: https://doi.org/10.1145/3106237.3106275

[57] R. A. P. Oliveira, E. AlÃ©groth, Z. Gao, and A. Memon, "Definition and evaluation of mutation operators for gui-level mutation analysis," in *2015 IEEE Eighth International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*, 2015, pp. 1–10.

[58] E. Alegroth, Z. Gao, R. Oliveira, and A. Memon, "Conceptualization and evaluation of component-based testing unified with visual gui testing: An empirical study," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, 2015, pp. 1–10.

[59] V. Dallmeier, M. Burger, T. Orth, and A. Zeller, "Webmate: A tool for testing web 2.0 applications," in *Proceedings of the Workshop on JavaScript Tools*, ser. JSTools '12. New York, NY, USA: Association

[60] A. Marchetto, P. Tonella, and F. Ricca, "State-based testing of ajax web applications," in *2008 1st International Conference on Software Testing, Verification, and Validation*, 2008, pp. 121–130.

[61] M. Biagiola, F. Ricca, and P. Tonella, "Search based path and input data generation for web application testing," 08 2017, pp. 18–32.

[62] U. Praphamontripong, J. Offutt, L. Deng, and J. Gu, "An experimental evaluation of web mutation operators," in *2016 IEEE Ninth International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*, 2016, pp. 102–111.

[63] Y. Sun, P. Jin, and L. Yue, "A framework of a hybrid focused web crawler," in *2008 Second International Conference on Future Generation Communication and Networking Symposia*, vol. 2, Dec 2008, pp. 50–53.

[64] A. Milani Fard and A. Mesbah, "Feedback-directed exploration of web applications to derive test models," in *Proceedings of the International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 2013, pp. 278–287.

[65] Y. Zheng, Y. Liu, X. Xie, Y. Liu, L. Ma, J. Hao, and Y. Liu, "Automatic web testing using curiosity-driven reinforcement learning," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, 2021, pp. 423–435.

[66] C. Degott, N. P. Borges Jr., and A. Zeller, *Learning User Interface Element Interactions*. New York, NY, USA: Association for Computing Machinery, 2019, p. 296–306. [Online]. Available: https://doi.org/10.1145/3293882.3330569

RahulKrishna Yandrapally received his Bachelors and Masters in Computer Science and Engineering from Indian Institute of Technology, Kanpur (IIT Kanpur) in 2012. Thereafter, he joined the Software Engineering group at IBM Research, India as a blue scholar and developed tools for web and mobile testing. Currently, he is a PhD candidate at the University of British Columbia. His main area of research is in software engineering, with emphasis on web applications, reverse engineering, and software testing. He was awarded the four-year fellowship in 2017. He is a student member of the IEEE Computer Society and served as a part of Virtualization Team in MobileSoft 2020.

Ali Mesbah is a professor at the University of British Columbia (UBC) where he leads the Software Analysis and Testing (SALT) research lab. His main area of research is in software engineering and his research interests include software analysis and testing, web and mobile-based applications, software maintenance and evolution, debugging and fault localization, and automated program repair. He has published over 60 peer-reviewed papers and received numerous best paper awards, including two ACM Distinguished Paper Awards at the International Conference on Software Engineering (ICSE 2009 and ICSE 2014). He is the recipient of a Killam Accelerator Research Fellowship (KARF) award (2020), a Killam Faculty Research Prize (2019) at UBC, and was awarded the NSERC Discovery Accelerator Supplement (DAS) award in 2016. He is currently an associate editor of the IEEE Transactions on Software Engineering (TSE) and regularly serves on the program committees of numerous software engineering conferences such as ICSE, FSE, ASE, ISSTA, and ICST.