Enhancing Mobile App Bug Reporting via Real-time Understanding of Reproduction Steps

Mattia Fazzini, *Member, IEEE*, Kevin Moran, *Member, IEEE*, Carlos Bernal-Cardenas, *Student Member, IEEE*, Tyler Wendland, *Student Member, IEEE*, Alessandro Orso, *Fellow, IEEE*, and Denys Poshyvanyk, *Member, IEEE*

Abstract—One of the primary mechanisms by which developers receive feedback about in-field failures of software from users is through bug reports. Unfortunately, the quality of manually written bug reports can vary widely due to the effort required to include essential pieces of information, such as detailed reproduction steps (S2Rs). Despite the difficulty faced by reporters, few existing bug reporting systems attempt to offer automated assistance to users in crafting easily readable, and conveniently reproducible bug reports. To address the need for proactive bug reporting systems that actively aid the user in capturing crucial information, we introduce a novel bug reporting approach called EBug. EBug assists reporters in writing S2Rs for mobile applications by analyzing natural language information entered by reporters in real-time, and linking this data to information extracted via a combination of static and dynamic program analyses. As reporters write S2Rs, EBug is capable of automatically suggesting potential future steps using predictive models trained on realistic app usages. To evaluate EBug, we performed two user studies based on 20 failures from 11 real-world apps. The empirical studies involved ten participants that submitted ten bug reports each and ten developers that reproduced the submitted bug reports. In the studies, we found that reporters were able to construct bug reports 31% faster with EBug as compared to the state-of-the-art bug reporting system used as a baseline. EBug's reports were also more reproducible with respect to the ones generated with the baseline. Furthermore, we compared EBug's prediction models to other predictive modeling approaches and found that, overall, the predictive models of our approach outperformed the baseline approaches. Our results are promising and demonstrate the feasibility and potential benefits provided by proactively assistive bug reporting systems.

Index Terms—Bug Reporting, Mobile Apps, Natural Language Processing, Language Modeling

♦

1 Introduction

Developers rely on bug reports to identify and fix problems in software. These bug resolution activities have been shown to constitute a major part of the software maintenance process, which in turn, typically accounts for a majority of the development effort [1]. There are generally two major types of reports that developers must manage: those produced by automated systems and those manually constructed by users. For in-field software failures that can be automatically recognized and captured, i.e., those that are revealed by a known oracle (e.g., a crash or a program assertion failure), researchers have developed automated techniques that capture fine-grained failure information and devised approaches to reproduce the failures based on the captured information. However, prior research that empirically analyzed issue trackers in open source projects found that software problems identified by known oracles

- M. Fazzini and T.Wendland are with the Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN, 55455.
 E-mail: [mfazzini,wendl155]@umn.edu
- K. Moran is with the Department of Computer Science, George Mason University, Fairfax, VA, 22030.
 E-mail: kpmoran@gmu.edu
- C. Bernal Cardenas, and D. Poshyvanyk are with the Department of Computer Science, William & Mary, Williamsburg, VA, 23185.
 E-mail: [cebernal.denvs]@cs.wm.edu
- A. Orso is with the School of Computer Science, College of Computing, Georgia Institute of Technology, Atlanta, GA, 30332.
 E-mail: orso@cc.gatech.edu

accounted for less than 50% of the studied issues [2]. This finding illustrates that a majority of software problems in the studied open source systems stemmed from functional issues that must be reported *manually*.

Despite the prevalence and importance of manually reported bugs, most existing issue tracking systems largely rely upon loosely structured, free-from text entry to capture information from reporters. The lack of structure generally results in reports whose quality is largely dependent upon the experience and thoroughness of the reporter. Past studies have illustrated that the most important information for developers in bug reports is reproduction steps, which is also the most difficult information for reporters to provide, resulting in reports that developers are unable to reproduce [3]. The difficulties often faced in manual bug reporting largely stem from the cognitive and lexical gap that exists between reporters and developers [4]. That is, there is typically a rift between the knowledge of the reporter of a bug and that of a developer, who has a wealth of domain specific knowledge about a given software system. In the current landscape of issue tracking systems, the task of bridging this gap falls almost entirely upon the stakeholders: either a reporter must spend extensive effort in crafting a report, or a developer must strive to "translate" the low-quality information provided in a poorly written report to map it to the code and ultimately reproduce and fix the issue.

To mitigate this problem, we investigate the potential of offloading part of the cognitive burden of bug reporting on

the reporting system itself. Specifically, we aim to *proactively assist* reporters in defining high-quality bug reproduction steps (S2Rs), so as to facilitate the reporting process and provide developers with more useful information, thus reducing the cognitive load on all the stakeholders. As past work showed, designing such a bug reporting system is a tenuous balancing act. On the one hand, in order to elicit more detailed information from reporters, a bug reporting system must impose some structure on the information being collected. On the other hand, doing so can make a system more difficult to use for reporters, which would ultimately limit its practical applicability and usefulness.

Prior work on the FUSION bug reporting interface has clearly illustrated this tension [4]. FUSION leveraged information gleaned from static and dynamic analysis to populate a series of dropdown menus having the objective to help reporters in providing S2Rs. That work demonstrated that an enhanced reporting system with populated, structured fields can increase the quality of reports as measured by their reproducibility. However, the higher quality of FU-SION's reports came at a cost, as it generally took reporters longer to write reports than with traditional free-form, textbased reporting systems. Longer report creation times may indicate a higher cognitive load for reporters, which is likely to hurt adoption in practice. In this paper, we further investigate these trade-offs in intelligent bug reporting systems, with the aim of designing a system that finds a sweet spot between ease of use and effectiveness in capturing high quality bug reports.

To this end, we introduce the EBUG reporting technique. At the core of EBUG there are (1) a real-time engine for understanding S2Rs written in natural language and (2) a predictive model capable of suggesting likely S2Rs, trained on real application usages. Supporting this core functionality are automated static and dynamic analyses capable of extracting a detailed GUI model from a given mobile app, including screenshots and GUI-related metadata. The primary means of interaction with EBUG is via a "smart" unstructured text field, wherein reporters are tasked with writing S2Rs in natural language. When a user begins writing, EBUG automatically recognizes different components of a given S2R (i.e., action, target GUI-component) and attempts to "auto-complete" missing information in a manner akin to Google's smart compose feature in Gmail [5]. EBUG's natural language understanding and predictive model are capable of recognizing and suggesting both entire S2Rs as well as individual components thereof. We developed EBUG's predictive modeling using a novel application of n-gram language modeling to sequences of actions from natural app usages collected by end-users. We provide a video demonstration of EBUG in its online appendix [6].

The main assumption for EBUG to work effectively is that reported S2Rs describe GUI actions in the app. The description of an S2R needs to be provided through text but does not need to follow a pre-imposed structure. Forthermore, EBUG does not make assumptions on the type of failure reported by the user as it focuses on the S2Rs associated with the failure.

We evaluated EBUG in two user studies that involved both bug reporting and bug reproduction tasks, based on 20 failures from 11 real-world apps. Specifically, the studies included ten participants submitting ten bug reports each and ten developers reproducing the submitted bug reports. The ten participants that submitted the bug reports were not developers but four of the participants have industry exposure through internships as software engineers. The participants that reproduced bug reports were all developers. These developers have experience in app development, but were not the developers of the apps considered in the study. Comparing bug creation and bug reproduction against the state-of-the-art FUSION approach, we found that participants were able to create high quality reports quicker, as compared to FUSION, and that the resulting reports were more reproducible that those created with FUSION.

This paper makes the following contributions:

- EBUG, a bug reporting system for mobile apps that leverages automated natural language understanding, static and dynamic analyses, and predictive modeling to facilitate the reporting process.
- An implementation of EBUG for Android apps that is publicly available, together with the artifacts and the infrastructure we used in our evaluation [6].
- An evaluation that provides initial evidence of the efficiency, effectiveness, and usefulness of EBUG.

2 TERMINOLOGY & MOTIVATING EXAMPLE

This section introduces some relevant terminology and presents an example that we use to motivate our approach.

2.1 Terminology

Given a bug report that describes a failure for an app, we informally use the terms relevant failure and relevant app to indicate the failure and the app. We use the term steps to reproduce (S2R) to indicate the textual description, in the bug report, of an operation that should be performed on the relevant app to reproduce the relevant failure (e.g., Enter "Transaction" in the "Description" text box). We use the terms GUI action (or simply action) and GUI interaction (or simply *interaction*) interchangeably to indicate the operation performed on relevant app's GUI. A GUI action is composed by its type (e.g., typing something in the GUI), the GUI element (or simply element) affected by the action (e.g., the text box in the GUI having the label "Description"), and, if present, the action's parameters (e.g., the "Transaction" text). We use the term target GUI element (or simply target) to indicate the GUI element affected by the action.

2.2 Motivating Example

Our motivating example, shown in Fig. 1, is a real bug report [7] for GNUCASH [8], a widely used real-world app to track expenses that has been installed over 100K times [9]. In the figure, the section labeled *Bug Report* contains the bug report submitted by user *21nds*, which contains the same information as the one in GNUCASH's issue tracking system [7]. The report describes a bug that manifested itself when the user changed the type of an account's transaction from Withdrawal to Deposit, and the app failed to store this change. The bug report is followed by a discussion (section labeled *Discussion* in Fig. 1), which was necessary for developers to understand and reproduce the bug. In the

1 Bug Report

Title: Deposit / Withdrawal change existing entry **Author:** 21nds

Text: An existing entry is created as a "Withdrawal". Try to change the existing entry from a "withdrawal" to "Deposit" button next to the amount, it does change the color of the amount to green and changes the button to say Deposit, but when you click the checkmark to save the entry it does not save the entry as a deposit. (also vice versa deposit ->Withdrawal) Also when entering it will auto-fill the amount from a previous transaction, and automatically mark it as withdrawal or deposit but will not save if you change it to the other.

2 Discussion

Message: #1

Author: codinguser

Text: Thanks for the feedback. Could you please provide some more information? What version of the app are you using? What Android version? Are you using double-entry accounting? How many splits were in the transaction? Or was it a simple transfer from one account to another.

Message: #6 Author: 21nds

Text: I have done a little more investigation. I have Deleted ALL accounts and created one account "Checking". I was using it for a simple checkbook register, so single accounting. It was working good about 2 releases ago in this fashion, but one of the updates in the last couple updates made it stop working for this. Thank You!

Message: #10 Author: fefe982 Text: Bug is confirmed.

Fig. 1. Bug report and reproduction discussion from GNUCASH app.

first message (message #1), an app developer (codinguser) asked the user to provide additional information on the bug. After a few messages among developers (message #6), the user reported that she deleted all accounts, created a single accounting account, and the transaction that she created in this account was the one leading to the bug described in the original report. After the user provided these additional details, an app developer (fefe982) could confirm the bug (message #10). The bug report took six days and ten discussion messages to be reproduced and confirmed.

In the bug report, the user provided a description containing some of the S2Rs but forgot to report some essential S2Rs that are necessary to reproduce the bug. These circumstances caused additional and unnecessary work for developers. Our bug reporting approach, EBUG, aim to decrease this extra work by guiding users in thoroughly and efficiently reporting S2Rs so that developers can more easily reproduce reported bugs.

3 THE EBUG APPROACH

In this section, we present EBUG, an approach for enhancing mobile app bug reporting. The basic idea behind EBUG is to help users write high-quality bug reports by processing

natural language S2Rs in real-time (i.e., as S2Rs are typed in the report) and guiding its users in adding S2Rs (or portions thereof) that make the bug reports complete and accurate.

Fig. 2 provides an overview of EBUG's workflow and its two main phases: *models generation* and *bug reporting*.

The models generation phase operates offline and generates three app-related models (depicted with solid, double borders in Fig. 2). First, in its *GUI model generation* step, this phase takes as input the *relevant app* and a *device* and combines static and dynamic analysis to compute the *GUI model* of the app. Then, in its *prediction models generation* step, this phase processes execution traces generated by the *app users* and creates two statistical models: the *GUI action prediction model* and the *GUI element prediction model*.

When a user is reporting a failure, the bug reporting phase processes the S2Rs in real-time using natural language processing, maps the content of the S2Rs to the GUI model of the relevant app, and uses the prediction models to suggest S2Rs (or portions thereof). The final output of this phase is a *bug report* containing S2Rs aimed at reproducing the failure experienced by the user. The rest of this section describes the two phases of EBUG in detail.

3.1 Models Generation Phase

3.1.1 GUI Model Generation

The GUI model (GM) captures which GUI actions are possible in the different screens of the relevant app and encodes GUI properties associated with actions. EBUG generates GM offline, as it might not be always practical to analyze S2Rs while the user is typing them in the bug report. For example, a solution that dynamically analyzes the GUI of the app while the user is typing the S2Rs might need to perform time-consuming operations for backtracking the exploration (e.g., restarting the app [10]) when the user edits/deletes previously typed S2Rs, which might significantly affect the performance of the approach.

EBUG builds the GM of the relevant app by (1) computing a static GUI model (GM_s) through static analysis, (2) extracting a dynamic GUI model (GM_d) through a dynamic analysis, and (3) combining GM_s and GM_d into GM. The use of both static and dynamic analysis allows EBUG to achieve high-coverage in modeling the GUI of the app and its corresponding GUI actions.

The GUI models computed by EBUG (GM_s , GM_d , and GM) are directed graphs G = (N, E), where N is the set of nodes in G and $E \subseteq N \times N$. EBUG creates two types of nodes. A node $s \in S$ corresponds to a screen in the app, whereas a node $v \in V$ corresponds to a GUI element. The set of nodes N contains the sets S and V, therefore having $N = S \cup V$. Both types of nodes have properties that are stored in the nodes as tuples. The properties of a screen sconsist of its name and screenshot ($\langle name, screenshot \rangle$). The properties of an element v consist of its text, identifier, type, and screenshot ($\langle text, id, type, screenshot \rangle$). (If one or more of these properties is missing or cannot be computed for a given element, EBUG stores an empty value for it.) EBUG uses the text and the identifier of an element to map S2Rs to GUI actions in the app during the bug reporting phase. We selected these properties as they are particularly suitable for

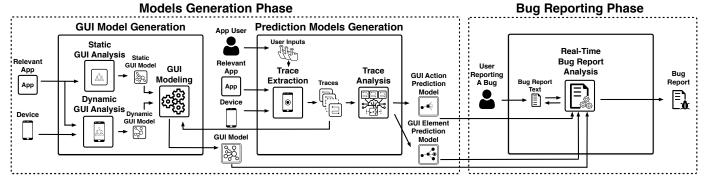


Fig. 2. High-level overview of the EBUG approach.

characterizing and identifying GUI elements from textual information in bug reports [11].

EBUG computes two types of edges $E = C \cup T$. A containment edge $c \in C \subseteq S \times V$ captures the containment relationship between a screen and a GUI element ($s \stackrel{c}{\rightarrow} v$), whereas a transition edge $t \in T \subseteq V \times S$ connects a GUI element to a screen $(v \xrightarrow{t} s)$ and models the transition to s after exercising v. Transition edges have two properties t ($\langle a-type, t-type \rangle$) that capture the (a-type) of action performed on v to trigger the transition to s, and the (*t-type*) of analysis—static or dynamic—used to extract the information from the app. Finally, EBUG models clicks, long clicks, scrolls, text typing as actions in the model. It also represents screen rotations in the model by associating to each screen s a screen rotation operation. Specifically, screen rotations are represented with a dummy GUI element v contained in s whose transition t reaches s. The rest of this section describes how EBUG computes GM.

To build the static GUI model GM_s , EBUG leverages a technique presented in related work [12] that captures (1) the creation and propagation of GUI elements in the code, (2) the relationship between screens and elements, (3) what actions can be performed on the elements, and (4) the effects of source code operations on the elements. To capture this information, the static analysis used by EBUG [12] builds and processes a constraint graph. The nodes in the constraint graph model statements creating and affecting the behavior of screens and GUI elements. Edges represent constraints on the flow of values between statements. The analysis identifies properties and relationships associated with screens and GUI elements by propagating values in the graph representing screens and GUI element objects. To propagate the values, the analysis uses a fixed-point algorithm. After running the fixed-point algorithm, the analysis extracts properties and relationships of screens and GUI elements by analyzing the values flowing to specific statements in the graph. Based on the results of the static analysis, EBUG first adds all the screens S_s to GM_s ; it then adds all the element nodes V_s , while creating a containment edge c_s between an element's screen and the element; it finally adds transition edges T_s between nodes in V_s and S_s (setting t-type = STATIC) based on the identified transitions.

EBUG builds the dynamic GUI model GM_d by dynamically exploring the relevant app through a depth-first traversal (DFT) that starts from the *initial/main* screen of the app and navigates through the app by interacting with the GUI

elements on each screen. Specifically, the traversal clicks on all the clickable components and types pre-defined text (i.e., "Test") on all editable elements that are reachable through the traversal. (The pre-defined value of text inputs could be extended in future work to use semantically relevant inputs as done in related work [13].) Before each step of the traversal is executed, the technique extracts relevant information about the current screen and its GUI elements. The technique then executes the action associated with each GUI component in a depth-first manner. During the exploration, the technique also captures the transitions between different screens by tracking the visited screens. In the traversal, if a GUI element is clicked and that operation would bring the traversal to a screen not belonging to the app (e.g., clicking a web link that would launch a browser app), the technique would execute a back command in order to continue the exploration in the relevant app. If the traversal exits the app and reaches the home screen of the device/emulator for any reason, the technique re-launches the app and continues the traversal using the depth-first strategy. During GUI exploration, there may be cases in which there are no GUI actions are available on the screen and EBUG must backtrack to make progress. When this occurs, our approach restarts the app (while resetting its data) and navigates to the screen containing the GUI element that should be exercised next. This step is possible because EBUG keeps track of visited screens and exercised GUI elements.

When EBUG visits a new screen s_d (i.e., $s_d \notin S_d$), it adds s_d to GM_d . It then adds all the elements $\{v_d^1,...,v_d^n\}$ in s_d to GM_d . Finally, it creates a containment edge c_d^i from s_d to v_d^i for each element v_d^i (with $i \in \{1, ..., n\}$) just added to GM_d . EBUG determines that two screens are the same if their GUI trees (composed of their GUI elements) are the same. In computing this information, EBUG does not consider the elements contained in "container elements" (e.g., list containers), as the number of contained elements can change dynamically. However, if a new element is present in such containers, this element is added to GM_d . When EBUG identifies that a screen s_d is different from the previously visited screen s_d^p , it adds a transition edge from v_d^p to s_d where v_d^p is is the element from s_d^p exercised by the approach (setting t-type = DYNAMIC). While adding a node or an edge to GM_d EBUG also collects the node's or edge's properties. For example, when EBUG adds a GUI element to the graph, the approach also saves the element's screenshot. Finally, for elements that do not have a *text* property, EBUG tries to extract this information from the screenshot of the element

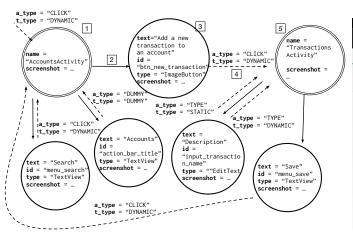


Fig. 3. Portion of the GUI model graph (GM) that includes parts of the AccountsActivity and TransactionsActivity screens.

using optical character recognition (based on the algorithm implemented in [14]).

EBUG creates GM by computing the union [15] of GM_s and GM_d that is, $GM = GM_s \cup GM_d = (N_s \cup N_d, E_s \cup E_d)$. While computing the union, EBUG identifies matching nodes in the two graphs. In this process, the approach does not consider the screenshot property. However, EBUG uses the screenshot values from the nodes in GM_d when the approach adds the nodes in GM. After taking the union of the two graphs, if a node $v \in V$ does not have any transition to any node in S, EBUG creates a "dummy" $s \in S \land s \xrightarrow{c} v \land c \in C$ and $t = \langle \text{DUMMY}, \text{DUMMY} \rangle$ to denote that after performing an action on v the screen does not change. During the bug reporting phase, EBUG uses GM to identify the GUI actions described by the S2Rs.

Fig. 3 provides a graphical representation for a portion of the *GM* computed from GNUCASH, which we presented in the motivating example (Sec. 2). This portion of the graph models the screens depicted in Fig. 4 and 5. Nodes with double borders are screens, while the remaining nodes are the elements of the screens. Edges with solid lines are containment edges, while the other edges are transitions.

3.1.2 Prediction Models Generation

In this step, EBUG builds two statistical models to capture how users interact with the relevant app and predict likely interactions given a certain context (i.e., a sequence of interactions). The two models are the GUI action and the GUI element prediction models (GAPM and GEPM in short, respectively). The bug reporting phase uses the GAPM to provide recommendations for not-yet-typed S2Rs. EBUG leverages the GEPM to, instead, help users complete partially written S2Rs.

EBUG computes the two models from user-generated execution traces. EBUG can collect these traces when developers distribute the relevant app for alpha or beta testing, the app runs in the field, or during in-house user testing. EBUG produces a trace by intercepting the actions performed on the app's GUI and logging them in the trace. The approach identifies clicks, long clicks, scrolling actions, text typing actions, and screen rotations. For each user execution, EBUG stores actions as a sequence of tuples in the trace. For each action, EBUG stores the tu-

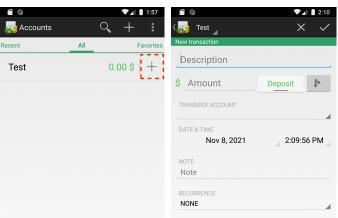


Fig. 4. Accounts Activity screen in the GNUCASH app.

Fig. 5. TransactionsActivity screen in the GNUCASH app.

ple: $\langle s\text{-}name, a\text{-}type, e\text{-}type, e\text{-}id \rangle$. The tuple contains the name of the screen (s-name) where the action was performed, the type (a-type) of the action, and the type (e-type) and the identifier (e-id) of the element exercised by the action. EBUG collects this information to uniquely identify an action during the app's execution. The approach uses the traces to also refine the GUI model. Specifically, EBUG translates the information contained in an action (containing screen, action type, and affected element) into corresponding nodes and edges in the GUI model.

Fig. 6 provides an example of a user trace. The figure reports a portion of a trace collected from the GNUCASH app. In this portion, the user created an account and added a transaction to the account. In Fig. 6, the fourth tuple represents a click on the "plus" button highlighted with a dashed red line in Fig. 4.

EBUG computes the two prediction models using two different sets of traces. EBUG derives these two sets of traces from the user traces. We call the two sets the *GUI action traces* (GATs) and the *GUI element traces* (GETs). EBUG uses the former set to create the GAPM and the latter set to generate the GEPM. EBUG creates a GAT and a GET for each user trace. The GAT and the GET are composed of a space-separated sequence of tokens. EBUG computes the tokens from the tuples in the corresponding user trace. Fig. 7 and 8 portray the GAT and the GET, respectively, derived from the user trace depicted in Fig. 6. (In both traces, the symbol \hookrightarrow indicates that the sequence included in the trace continues.)

EBUG uses *n*-gram-based language modeling [16] to build the GAPM and the GEPM from the GATs and the GETs, respectively. At a high level, language models assign probabilities to sequences of words and can be used to estimate the probability of a word given its preceding words (also known as the history of a word) [16]. EBUG uses this characteristic of language models during its bug reporting phase to predict the occurrence of an S2R (or part thereof) given a sequence of previously occurring S2Rs. Because the GUI action and element traces do not contain sequences of S2Rs but contain sequences of GUI-action-based tokens, in this step, EBUG creates token-based prediction models. EBUG maps tokens to S2Rs during its bug reporting phase. We chose to leverage n-gram based models over other more complex sequence based prediction models (RNNs and

```
...

\(\langle \text{AccountsActivity,CLICK,TextView,menu_add_account} \)

\(\langle \text{AccountsActivity,TYPE,EditText,edit_text_account_name} \)

\(\langle \text{AccountsActivity,CLICK,TextView,menu_save} \)

\(\langle \text{AccountsActivity,CLICK,ImageButton,btn_new_transaction} \)

\(\text{TransactionsActivity,TYPE,EditText,input_transaction_name} \)

\(\text{TransactionsActivity,TYPE,EditText,input_transaction_amount} \)

\(\text{TransactionsActivity,CLICK,TextView,menu_save} \)
```

Fig. 6. Portion of a user trace collected from GNUCASH app.

TransactionsActivity.CLICK.TextView.menu_save ...

Fig. 7. Portion of the GUI action trace computed from the user trace in Fig. 6.

Fig. 8. Portion of the GUI element trace computed from the user trace in Fig. 6.

other Neural Language models) as such models typically require large amounts of training data to be effective.

EBUG's *n*-gram models are able to provide suggestions for sequences that are shorter than the order of the *n*-gram models or for sequences that appear with a slightly different context (i.e., arrangement) by leveraging the Kneser-Ney smoothing method [17]. This method is one of the most commonly used and best performing *n*-gram smoothing methods [16]. Using the Kneser-Ney smoothing method, lower-order *n*-gram sequences are encoded in the models. Lower-order *n*-gram sequences are then used to make predictions for sequences shorter than the order of the models and to provide suggestions for sequences appearing with a slightly different context [16].

To compute EBUG's prediction models, EBUG treats trace tokens as words and uses the GATs and GETs as the training data to build the the *n*-gram-based models characterizing the GAPM and the GEPM. EBuG's n-gram models use closed vocabularies that are computed using the tokens of the training traces. Based on this characteristic, when EBUG translates the text of an S2R into a prediction model token in the bug reporting phase, and the token is not part of the model vocabulary, the technique does not make any suggestions to the reporter. Because EBUG considers ngram based models and the length of the token sequence used to make the prediction has a fixed length (i.e., n-1), a token that is not part of the vocabulary would impact (i.e., prevent) only *n*-1 predictions. This characteristic does not affect the ability of the models to identify suggestions for sequences that were not observed in the training traces.

To instantiate the language modeling framework described so far, EBUG needs to determine the order of the n-grams used by the models. EBUG determines the order using the concept of wasted effort. We define wasted effort (we) as the number of S2R-related suggestions that a user needs to process before finding the suggestion that completes the sequence of already entered S2Rs. (The concept of wasted effort is similarly used in the fault localization literature [18] to evaluate fault localization techniques.) To give an example, if a user needs to process two "unrelated" suggestions before finding the "useful" suggestion, then we=2. When considering multiple prediction tasks, we determine the performance of a prediction model using the wasted effort score metric:

$$wes = \sum_{i=1}^{k} we_i / \sum_{i=1}^{k} c_i \tag{1}$$

In Equation 1, we_i is the wasted effort in task i and c_i

is equal to 1 if the model provided the desired suggestion for task i, 0 otherwise. To give an example, in a model that provides three suggestions for every prediction task i, c_i will be equal to zero if none of the three suggestions is the desired suggestion, and the wasted effort we_i would be equal to three. If the model provided the desired suggestion in the third position of the suggestion list, c_i will be equal to one and we_i will be equal to two.

Because wes is affected by the length of the suggestions provided to the users, we search the best performing model across two dimensions: the *n*-gram order and the length of the suggestions. EBUG limits the search space by consider n-grams up to length 10 and suggestions up to length 10. (Our experiments detailed in Sec. 5 show that higher values for the order and the length is highly unlikely to lead to better results). EBUG searches the prediction model that minimizes wes and uses that model in the bug reporting phase. (The approach performs the search for each of the two prediction models used by EBUG.) EBUG leverages the GATs and the GETs to compute the wes of different models. If the relevant app has a database of traces extracted from existing bug reports, EBUG also uses these traces to compute the wes of different models. While computing wes, EBUG maps the concept of wasted effort to trace tokens and uses the token sequences from the traces to check if a token suggestion is a useful suggestion. A token suggestion is useful if the suggested token is the same as the next token in the sequence considered. EBUG computes wes for a model using leave-one-out (sequence) cross-validation. For each leave-one-out *test* sequence, the approach considers the first token in the sequence and uses the model to provide suggestions for the next token. (For the GEPM, EBUG makes predictions only for the element tokens.) EBUG repeats this process until it reaches the last token in the sequence. EBUG divides the sum of the wasted effort by number of "useful" suggestions across all folds to compute wes for the model. Finally, EBUG selects the model with the minimum westo find the model that minimizes the effort needed by the user to find a "useful" suggestion. In computing wes, we consider that the benefit of finding a "useful" suggestion has the same weight as the cost of processing an "unrelated" suggestion. We leave as future work the idea of investigating different weights in the computation of wes.

In the process of comparing prediction models, we do not need to use an intrinsic evaluation metric such as perplexity [16] as our extrinsic evaluation metric *wes* is fairly inexpensive to compute. Because EBUG can determine the *n*-gram order of the prediction models automatically, the

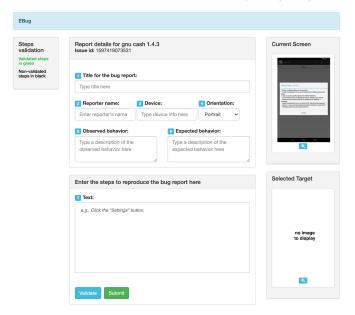


Fig. 9. EBug's bug reporting interface.

approach can update the models as new training traces or bug reports are collected from the users. Although EBUG provides a mechanism to automatically determine the *n*-gram order of the prediction models, developers can also specify this value as well as the desired length for the suggestions manually.

When EBUG reaches the end of this step, the GAPM and the GEPM are ready for use in the bug reporting phase.

3.2 Bug Reporting Phase

In this phase, EBUG guides its users in the process of reporting failures so that app developers can easily reproduce submitted bug reports. Specifically, when a user is reporting a failure, the approach analyzes the S2Rs contained in the report and provides S2R-related suggestions so that the user can include high-quality (i.e., no missing or ambiguous) S2Rs in the report. Additionally, EBUG also provides S2R-related suggestions to accelerate the bug reporting process.

Users can report bugs using the interface depicted in Fig. 9, where S2Rs are reported in the text box labeled with [7]. Algorithm 1 describes how EBUG guides users in the process of reporting S2Rs. The algorithm takes as inputs the GUI model (GM), the GUI action prediction model (GAPM), and the GUI element prediction model (GEPM). The algorithm's output is a bug report (bugReport) containing the S2Rs reported by the user. This output is also the final output of the approach. At a high level, the algorithm uses natural language processing to map the S2Rs to GUI actions in the application, analyzes the sequence of actions to provide suggestions for not-yet-typed S2Rs, and processes partial sentences to provide suggestions on how to complete partially written S2Rs.

Throughout the bug reporting phase, Algorithm 1 processes provided S2Rs and encodes them into a list of *S2R entities* (currS2REs in Algorithm 1). EBUG creates an entity for each S2R identifying a GUI action in the relevant app. An S2R entity is a tuple that contains: the text of the S2R

Algorithm 1: S2Rs analysis in the bug reporting phase.

```
Input: GM: GUI model for the relevant app
            GAPM: GUI action prediction mode
            GEPM: GUI element prediction model
   Output: bugReport: bug report describing the failure
  begin
     currS2REs = []
     currScreen = Get-Current-Screen(GM,currS2REs)
     Display-Screen(currScreen)
     while TRUE do
       upon event S2RS_TEXT_CHANGE do
        currS2RsText = Get-S2Rs-Text()
7
        textEdit = Get-Text-Edit()
8
        Check-Updated(currS2RsText,textEdit,currS2REs)
        prevS2REs = currS2REs
10
        textEditOp = textEdit.Get-Text-Operation()
11
        if textEditOp = INSERT then
12
          newText = textEdit.Get-New-Text()
13
           if textEdit.Get-New-Text() == ST then
14
            currS2REs = Compute-S2REs(GM, currS2RsText, prevS2REs)
15
16
            Display-Validated-S2Rs(currS2REs)
            Display-Screen(GM,currS2REs)
17
18
            S2RSuggestions = Get-GA-Suggestions(GAPM, currS2REs)
19
            {\tt Display-S2Rs-Suggestions}({\tt S2RSuggestions})
20
           else if newText == SPACE then
            currS2RText = Get-Current-S2R-Text(currS2RsText.textEdit.
21
             prevS2REs)
22
            suggestionType = Analyze-Partial-S2R-Text(currS2RText)
            if \ {\tt suggestionType} == {\tt TARGET} \ then
23
              currS2RAction = Get-Current-S2R-Action(currS2RText)
24
              precS2REs = Get-Prec-S2REs(Get-Prec-Text(currS2RText,
               currS2RsText),prevS2REs)
              GESuggestions = GEPM.Get-GE-Suggestions(precS2REs,
               currS2RAction)
              Display-GE-Suggestions(GESuggestions)
27
28
            else if suggestionType == PARAM then
              PSuggestion = Get-P-Suggestion(currS2RText)
29
30
              Display-P-Suggestion(PSuggestion)
            else if suggestionType == STRUCTURE then
31
              SSuggestion = Get-S-Suggestion(currS2RText)
32
33
             Display-S-Suggestion(SSuggestion)
         else if textEditOp = DELETE then
34
35
          currS2REs = Compute-S2REs(GM, currS2RsText, prevS2REs)
36
          Display-Validated-S2Rs(currS2REs)
         Display-Screen(GM, currS2REs)
37
       upon event SUBMIT_BUG_REPORT do
38
        currS2RsText = Get-S2Rs-Text()
39
        prevS2REs = currS2REs
40
        currS2REs = Compute-S2REs(GM, currS2RsText, prevS2REs)
41
        bugReport = Get-Bug-Report()
42
        bugReport.Set-PS2Rs(currS2REs)
43
        return bugReport
```

(*s2r-text*), the *abstract GUI action* (AGA) representing the S2R (*a-action*), the GUI action corresponding to the AGA (*action*), the GUI screen on which the action is performed (*b-screen*), and the screen displayed by the relevant app after performing the action (*a-screen*).

An AGA is composed of the elements of a sentence that describe a GUI action. More precisely, an AGA is characterized by the type of the action (a-type), the description of the element exercised by the action (e-desc), and the description of the parameters associated with the action (p-desc). All of the AGAs contain the type of the action, but not all of them contain the description of the element or the action, as this information might not apply. We call the action "abstract" because the action does not contain executable information. To provide an example, the AGA for the sentence Enter "Transaction" in the "Description" text box (that describes an operation in the app of the motivating example) is $\langle TYPE$, the "Description" text Description" text Description". We detail how EBUG computes AGAs in Sec. 3.2.1.

A GUI action contains the type of the action (*a-type*), the element affected by the action (*element*), and the parameters of the action (*params*). An element is composed of the GUI screen of the element (*e-screen*), the type of the element (*e-type*), the identifier of the element

TABLE 1
Some of EBug's rules to identify a GUI action and its details from a clause's dependency tree.

ID	Grammar Rules	Clause Example	Abstract GUI Action
A1	$(n_1 \in C) \xrightarrow{e_1 = dobj} (n_2)[tree(n_2)]_{e\text{-}desc}$	Click the "Transaction" element.	⟨CLICK, the "Transaction" element, ⟩
A2	$(n_1 \in C) \xrightarrow{e_1 = advmod} (n_2 = long) $ $\xrightarrow{e_2 = dobj} (n_3)[tree(n_3)]_{e-desc}$	Long click the "Transaction" element.	$\langle {\sf LONG_CLICK}, the "Transaction" element, \rangle$
A3	$(n_1 \in R)$	Rotate the screen.	$\langle ROTATE, , \rangle$
A4	$ \begin{array}{c} (n_1 \in S) \xrightarrow{e_1 = prt} (n_2 \in D)[n_2]_{p\text{-}desc} \\ \xrightarrow{e_2 = prep} (n_3) \xrightarrow{e_3 = pobj} (n_4)[tree(n_4)]_{e\text{-}desc} \end{array} $	Scroll up on the "Transactions" list.	$\langle SCROLL, the "Transaction" list, UP \rangle$
A 5	$(n_2)[tree(n_2) \setminus tree(n_2, prep)]_{p\text{-}desc} \xleftarrow{e_1 = dobj} (n_1 \in T)$ $\xrightarrow{e_2 = prep} (n_3) \xrightarrow{e_3 = pobj} (n_4)[tree(n_4)]_{e\text{-}desc}$	Enter "Transaction" in the "Description" text box.	\(\text{TYPE}\), the "Description" text box, "Transaction"\)

TABLE 2 Some of EBug's rules to identify suggestions for partial clauses.

ID	Grammar Rules	Clause Example	Suggestion
S1	$(n_1 \in C) \nrightarrow$	Click	PARTICLE
S2	$(n_1 \in T) \nrightarrow$	Туре	PARAM
S3	$ \begin{split} &(n_1 \in T) \nrightarrow \\ &(n_2) \xleftarrow{e_1 = aobj} (n_1 \in T) \\ &\xrightarrow{e_2 = prep} (n_3) \xrightarrow{e_3 = pobj} (n_4 \in DET) \nrightarrow \end{split} $	Type "Transaction" in the	TARGET

(e-id), and the text displayed by the element (e-text). To provide an example, the GUI action for the sentence Enter "Transaction" in the "Description" text box is $\langle TYPE, \langle TransactionsActivity, EditText, input_transaction_name, "Description" <math>\rangle$, ["Transaction"] \rangle . We describe how EBUG identifies GUI actions from AGAs in Sec. 3.2.2.

Algorithm 1 starts by initializing the current list of S2R entities to be empty (line 2). In this initialization phase, the algorithm also identifies the app's initial screen (line 3) and displays this information to the user while preemptively retrieving the screen information from GM. Although EBUG displays this information, the user can also report S2Rs starting from a different screen of the apps. After these initialization steps, the algorithm executes its main loop (lines 5-44), where EBUG helps the user in reporting S2Rs (lines 6-37). The loop terminates when the user decides to submit the bug report (lines 38-44). In this last step, the algorithm stores the content of the bug report together with the S2R entities associated with the S2R description.

We now detail EBuG's NLP analysis (Sec. 3.2.1) and then describe how EBuG leverages the analysis to helps users in reporting S2Rs (Sec. 3.2.2).

3.2.1 Natural Language Processing of S2Rs

EBUG processes each sentence in the text of the S2Rs using a two-step approach. First, EBUG preprocesses a sentence in the report to simplify its analysis. Specifically, the approach performs noise removal, lexicon normalization, and object standardization [19] on the sentence as described in related work [11]. For noise removal, EBUG discards content within parenthesis, which, in our experience, is generally unnecessary for mapping an S2R to the corresponding GUI action. For lexicon normalization, EBUG normalizes non-standard words to their canonical form (e.g., EBUG replaces the word "Tap" with "Click"). (We use the word replacements identified in related work [11] to perform this task.) For object standardization, the approach simplifies the text of a sentence by replacing sequences of words referring to a GUI element in the relevant app with a freshly-created textual identifier. EBUG performs this operation using the textual

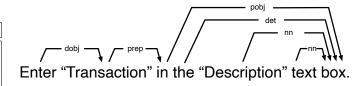


Fig. 10. Dependency tree for a sentence.

properties saved in the GUI model (GM) and leveraging the fact that the text displayed by an app usually follows a title or sentence-case convention [20]. This last operation helps in simplifying the analysis of the text of the S2Rs.

After its preprocessing step, EBUG analyzes the sentence. More precisely, the approach analyzes each clause that appears in a sentence, as each of the clauses could described a different action in the relevant app. To identify the clauses of a sentence, the approach leverages related work [21], which parses a dependency parse tree recursively and, at each step, predicts whether an edge should yield an independent clause. EBUG analyzes clauses using their dependency tree representation [16], [22]. A dependency tree is a directed graph that captures the syntactic structure of a clause and provides a representation of the grammatical relations between words in the clause. Words are nodes, and relations are edges in the tree. Fig. 10 provides an example of a dependency tree (extracted from the clause enter "Transaction" in the "Description" text box). (The dependency tree in the figure does not report relations between words and punctuation for readability.)

To identify whether a clause describes a GUI action, EBUG analyzes a clause's dependency tree using a rulebased approach as done in related work [11], [23]. Specifically, EBUG uses grammar rules to identify whether a tree represents a GUI action and to and extract the relevant information characterizing the action into an AGA. Table 1 reports some of EBUG's grammar rules to extract AGAs (For the complete list of grammar rules, please check our online appendix [6].) Column *Grammar Rule* reports the rule information, column Clause Example provides an example of a clause matched by the corresponding rule, and column Abstract GUI Action details the AGA extracted from the corresponding clause. For example, the rule $(n_1 \in C) \xrightarrow{e_1 = dobj}$ $(n_2)[tree(n_2)]_{e-desc}$ matches the clause Click the "Transaction" element., and (CLICK, the "Transaction" element,) is the AGA extracted from clause.

EBUG uses grammar rules also to analyze partial clauses. Table 2 reports some of the grammar rule used for this task. Given a partial clause, EBUG uses the grammar rules to

1a	S2Rs Description	[] Click on the "Save" element. Click the "new transaction" button.
1b	Abstract GUI Actions	[] {CLICK,the "Save" element,);{CLICK, the "new transaction" button,}
1c	GUI Actions	[] (CLICK, (AccountsActivity, TextView, menu_save),); (CLICK, (AccountsActivity, ImageButton, btn_new_transaction, "Add transaction to an account"),)
1d	S2R Entities	[_] (Click on the "Save" element.,(CLICK,the "Save" element,),(CLICK,(AccountsActivity,TextView,menu_save),), AccountsActivity, AccountsActivity) (Click the "new transaction" button.,CLICK, the "new transaction" button,), (CLICK,(AccountsActivity,ImageButton,btn_new_transaction,"Add transaction to an account"),), AccountsActivity, TransactionsActivity)
2a	GUI Action Prediction Model Suggestion	[] AccountsActivity.CLICK.TextView.menu_save AccountsActivity.CLICK.ImageButton.btn_new_transaction TransactionsActivity.TYPE.EditText.input_transaction_name
За	GUI Action for Prediction	(TYPE,(TransactionsActivity,EditText,input_transaction_name,"Description"),[])
3b	S2Rs Suggestion	[] Click on the "Save" element. Click the "new transaction" button. Type "value" in the "Description" text box
Зс	EBug's S2Rs Interfatce	Enter the steps to reproduce the bug report here 7 Text: Click on the "Dismiss" button. Click on the "Cancel" button. Click on the "Create Account" button. Type "Checking" in the "Account name" text box. Click on the "Save" element. Click the "new transaction" button. Type "value" in the "Description" text box

Fig. 11. Example of the information used by EBug to analyze new text in the S2Rs description.

identify which suggestion to make. The approach makes three types of suggestions: PARTICLE, PARAM, and TARGET. Suggestions of type PARTICLE (e.g., "the") have the objective to help users write well-formed clauses, which help EBUG's analysis. Suggestions of type PARAM remind the users that they also need to specify a parameter for certain types of action (e.g., when users describe that they typed some text in the GUI, they should also specify the text they typed). Suggestions of type TARGET provide users with a list of likely targets for the partial S2R, and this list is computed using the GUI element prediction model. For example, given the partial clause *Type "Transaction" in the*, EBUG uses the rule \$3\$ to identify that this clause requires an action target to define a complete S2R.

3.2.2 S2Rs Reporting

This part of EBUG is summarized in Algorithm 1 at lines 5-44. EBUG's analysis performs different operations based on whether the user enters (lines 12-33) or removes (lines 34-37) text to/from the S2Rs description. At a high level, when the user enters new text (Case 1), the algorithm helps the user by validating the user-provided S2Rs (i.e., informs the user whether the S2Rs can be mapped to GUI actions in the app) and by providing S2R-related suggestions on how to complete the report. When the user removes text from the S2Rs description, EBUG re-validates user-provided S2Rs (Case 2). We now describe Case 1 and Case 2 in detail.

Case 1: Handling Additions to the S2Rs Description

EBUG processes new text in the S2Rs description only when the user either enters a sentence terminator (ST) or a space character (SPACE). When the user enters a sentence terminator (Case 1a), EBUG assumes that the user finished describing a certain S2R and provides suggestions for possibly following S2Rs. When the users enters a space character (Case 1b), the algorithm considers the description of the certain S2R to be partial, and helps the user in completing the S2R by providing suggestions on its content and structure.

Case 1a: Processing Complete S2Rs

When the user enters a sentence terminator (lines 14-19), the algorithm (i) maps the S2Rs to GUI actions in

the app while encoding the GUI actions into S2R entities (Compute-S2REs), (ii) informs the user about validated S2Rs and, if necessary, updates EBUG's interface to display the app's screen associated with the sequence of GUI actions extracted from the user-provided S2Rs (Display-Screen), (iii) computes suggestions for not-yet-typed S2Rs (Get-GA-Suggestions), and (iv) provides the suggestions to the user (Display-S2Rs-Suggestions). Fig. 11 provides an example of some of the information computed by EBUG in this process. We now describe the four steps of EBUG that perform these operations.

Step 1: Mapping S2Rs to S2R entities. After the user completes a specific S2R, EBUG uses the NLP analysis described in Sec. 3.2.1 to extract a list of abstract GUI actions (AGAs) from the S2Rs description. Given a list of AGAs (e.g., the list of tuples labeled with 1b in Fig. 11), EBUG uses the GUI model (GM in Algorithm 1) to map the AGAs into a sequence of GUI actions and encode the actions into S2R entities. At a high level, EBUG identifies GUI actions based on the description contained in the AGAs. The mapping process operates as follows. EBUG uses the list of S2R entities from the previous mapping (prevS2REs) to identify the oldest S2R that the approach did not process yet. To do so, EBUG processes the newly computed AGAs and finds the longest sequence matching the AGAs in prevS2REs. The approach starts processing AGAs from the first one that does not have a match, and also processes all subsequent AGAs. EBUG looks at past matches to avoid processing AGAs that did not change. If prevS2REs is empty, EBUG starts the mapping task from the first AGA in the list.

For each AGA, EBUG identifies the corresponding action using four pieces of information: (i) the AGA (aga), (ii) GM, (iii) the screen (rs) "reached" by the GUI action matched with the previously considered AGA, and (iv) the screen (lrs) "reached" by the last GUI action that was successfully matched to an AGA. For example, based on the GUI model represented in Fig. 3, if action $\langle \text{CLICK}, \langle \text{AccountsActivity}, \text{ImageButton}, \text{btn_new_transaction}, "Add transaction to an account"} \rangle$, $||\rangle$ was the action matched to the previous AGA, then rs would correspond to TransactionsActivity, and lrs would correspond to the same screen. For the first AGA in the

$$R(lrs, e\text{-}desc, v) = \begin{cases} \alpha * Max(S(e\text{-}desc, v.e\text{-}text), S(e\text{-}desc, v.e\text{-}id)) + (1-\alpha) * \frac{1}{1+D(lrs, v)} & \text{if } Max(S(e\text{-}desc, v.e\text{-}text), S(e\text{-}desc, v.e\text{-}id)) > \beta \\ 0 & \text{otherwise} \end{cases}$$

Fig. 12. Equation to rank relevant GUI elements when mapping abstract GUI actions into GUI actions.

list, EBUG uses the initial screen of the relevant app as the value for rs and lrs. When the approach is not able to find a matching GUI action while analyzing the preceding AGA, EBUG assigns the special value (*) to rs, which indicates that the currently reached screen could be any screen in GM.

To find the GUI action corresponding to an AGA, the approach performs four operations. First, EBUG identifies all the relevant elements v in rs. (If rs=*, EBUG selects all the elements in GM.) Second, EBUG refines the list of elements by selecting only the ones that have an outgoing transition edge $(\stackrel{\tau}{\rightarrow})$ such that the action type of the edge is equal to the action type of the AGA. Third EBUG ranks the elements based on how relevant they are to the element description (e-desc) contained in the AGA and how close they are to lrs. The approach uses the equation in Fig. 12 to determine how relevant (function R(lrs, e-desc, v) or R in short) an element is. Finally, if the top-ranked element has R > 0, the approach selects the element as the target of the GUI action and creates the corresponding S2R entity. For example, considering the $\langle CLICK, \text{the "new transaction" button}, \rangle$, represented GUI model in Fig. rs=lrs=AccountsActivity, EBUG would find that the top-ranked element is \(\text{AccountsActivity}, \text{ImageButton}, \) btn_new_transaction, "Add transaction to an account" \>. When the top-ranked element has R=0, EBUG creates an S2R entity without a GUI action, indicating that the approach could not find a matching GUI action for the S2R. We now describe the details of the equation in Fig. 12.

The equation is based on three parameters (lrs,e-desc, and v). The first part of the equation has two components, $\alpha * Max(S(e\text{-}desc, v.e\text{-}text), S(e\text{-}desc, v.e\text{-}id))$ and $(1 - \alpha) *$ $\frac{1}{1+D(lrs,v)}$. EBUG uses the first component to compute the similarity between e-desc and v. The approach computes the similarity by taking the maximum (function Max) of two values. The first value corresponds to the semantic similarity between e-desc and the text of the GUI element (e-text). The second value represents the semantic similarity between e-desc and the identifier of the GUI element (e-id). EBUG uses semantic similarity values, and not stringdistance-based metrics (e.g., the Levenshtein distance [24]), to account for the fact that users might use, in their bug reports, words that have a different representation from the ones in the developer-defined labels, but these words could be either synonyms or similar in meaning (e.g., "create" and "add"). EBUG computes semantic similarity values using vectors based on word embeddings extracted from a FastText model [25]. Our approach uses FastText as it is has been shown to outperform related work [26], [27], [28], [29], [30] in different contexts [25].

The approach represents each of the three textual properties (*e-desc*, *e-text*, and *e-id*) in Fig. 12 as a single vector computed by removing stopwords (as they introduce unnecessary noise) and by averaging the vectors of the remaining words. By taking the average, EBUG is able to incorporate

the meaning of every word in the vector representation. For the element identifier, EBUG also performs an additional preprocessing step in which the approach splits the value of the property into words. Specifically, EBUG replaces underscores with spaces and splits apart composite words that follow a camel case convention, both of which are common occurrences in apps [31]. To determine the similarity of two properties (e.g., e-desc and e-text), the approach computes the cosine similarity between the corresponding vectors. The cosine similarity ranges between [0,1], where 1 corresponds to the highest similarity. If both S(e-desc, v.e-text)and S(e-desc, v.e-id) are below a threshold value β , EBUG does not consider the element as being relevant to the AGA and R returns 0. The approach uses 0.5 as the default value for β . EBUG uses this value to disregard unrelated matches and the value is based on an empirical analysis done in related work [11].

EBUG uses the second component $((1-\alpha)*\frac{1}{1+D(lrs,v)})$ to prioritize matches with elements belonging to screens that are in the proximity of the screen containing the previously matched element. Specifically, D(lrs,v) is the distance between lrs and the screen vs containing v. The approach measures the distance by finding the shortest path between lrs and vs in GM and counting the number of transition edges in the path. If EBUG does not find a path connecting lrs and vs, the approach sets D(lrs,v)=0. $\frac{1}{1+D(lrs,v)}$ ranges between [0,1], where 1 corresponds to the shortest distance (i.e., the screens are the same). When lrs=rs and $rs\neq *$, the value of $\frac{1}{1+D(lrs,v)}$ is always 1 as EBUG selects elements from rs.

Finally, the equation in Fig. 12 uses the constant α to define the weight of the two equation's components. EBUG uses 0.5 as the default value for α to assign the same weight to the two components.

Step 2: Displaying Validated S2Rs and Current Screen. In this step, EBUG provides feedback to the user based on whether the approach was able to map the S2Rs into GUI actions. The goal of this step is to inform the user that certain S2Rs might require further refinement. Specifically, for the S2R entities that EBUG could map to GM (i.e., entities containing the corresponding GUI action), the approach displays and highlights (in green) the text of the S2Rs in its graphical interface. For S2R entities not mapped to GM, the approach reports their S2Rs in its graphical interface but does not highlight them, as we wanted to avoid reporting false negatives (i.e., valid actions not mapped to GM due to the possibly partial information encoded in the model). EBUG displays the results of the validation process in the section of Fig. 9 labeled with Steps validation. During this step, the approach also updates its interface to display the reached screen rs associated with the last reported S2R entity. If rs=*, EBUG displays a screen informing the user that the last S2R was not matched. The screen is displayed in the part of Fig. 9 labeled with Current Screen.

Step 3: Computing S2R Suggestions. EBUG uses the GUI action prediction model (GAPM in Algorithm 1) to provide

suggestions for not-yet-typed S2Rs. To identify relevant suggestions, the approach first encodes the S2R entities into a sequence of tokens and then uses GAPM to identify the next token (S2R) for the sequence. Specifically, EBUG translates the last n-1 S2R entities into tokens. (The tokens are in the same order as they appear in the S2Rs description.) EBUG uses the GUI actions associated with the S2R entities to compute the sequence tokens. The structure and the content of the tokens processed by GAPM are detailed in Sec. 3.1.2. In Fig. 11, the parts labeled with $\boxed{1d}$ and $\boxed{2a}$ show how EBUG translates S2R entities into tokens suitable for GAPM. Once the approach retrieves the token suggestions from GAPM, EBUG uses them to display suggestions to the user.

Step 4: Providing S2R Suggestions to the User. Given a list of suggestions, EBUG displays them to the user in the same order as provided by GAPM. First, for each token, the approach retrieves additional information associated with the token (e.g., its text) from GM Then, the approach displays the text of the S2R corresponding to the token, and, when available, a screenshot of the GUI element affected by the action associated with the token. (EBUG does not always have a screenshot for a GUI element, as the approach uses static analysis to identify some of the elements in the GUI model). To translate a token into an S2R, EBUG uses textual templates. EBUG has templates for each GUI action supported by the approach. The approach fills the template with the information associated with the action represented by a token. In Fig. 11, the parts labeled with 2a | 3a | 3b | and 3c show how EBUG displays suggestions to the user.

Case 1b: Processing Partial S2Rs

When the user enters a space character (lines 20-33), EBUG (i) analyzes the text of the S2R to identify whether the S2R needs further refinement (Analyze-Partial-S2R-Text), (ii) computes suggestions on how to complete the S2R, and (iii) provides the suggestions to the user.

Step 1: Analyzing Partial S2Rs. EBUG uses the NLP analysis described in Sec. 3.2.1 to analyze the current S2R and identify how to help the user in completing the S2R. EBUG can (i) suggest likely targets (i.e., GUI elements) for the type of GUI action described in the S2R (lines 23-27), (ii) inform the user that the S2R should also specify a parameter (i.e., the text to be typed in a text box) for the GUI action described in the S2R (lines 28-30), and (iii) add grammatical particles to the S2R so that the S2R is easier to analyze (lines 31-33).

Step 2: Computing Partial S2R Suggestions. The approach leverages the GUI element prediction model (GEPM in Algorithm 1) to suggest likely targets for a certain S2R. To identify relevant suggestions, the approach extracts the action type associated with the S2R, encodes the action type into a model token, transforms preceding S2R entities (precS2REs) into model tokens, and uses the sequence of tokens to identify predictions based on GEPM. Similarly to what EBUG does in the case of a complete S2R suggestion, if an S2R entity does not have a corresponding GUI action, EBUG does not compute partial S2R suggestions for sequences that include the entity. Once the approach retrieves the token suggestions from GEPM, EBUG stores them in the same order as they are provided by the model and uses them to display suggestions to the user. Additionally, the technique is also

able to provide suggestions based on partially typed GUI elements. In this case, the suggestions are not based on the GEPM but are based on text matching. When EBUG identifies that the S2R needs a parameter or a suggestion for improving the structure of the S2R, the approach provides the corresponding textual suggestion.

Step 3: Providing Partial S2R Suggestions to the User. EBUG offers suggestions for GUI elements similarly to how it provides suggestions for complete S2Rs. Suggestions for parameters and the structure of certain S2Rs are automatically added to the bug report and the user can accept them by pressing the tab character.

Case 2: Handling Removals from the S2Rs Description

When the user removes text from the S2Rs description (lines 34-37), the algorithm analyzes whether the sequence of GUI actions associated with the S2Rs has changed, and updates EBUG's interface to display the app's screen associated with the action sequence. Specifically, EBUG (i) maps S2Rs into S2R entities, (ii) validates the S2Rs, and (iii) displays the screen reached by the last S2R entity in the sequence. The approach performs these steps following the methodology described in *Case 1a-Step 1* and *Case 1a-Step 2*.

4 IMPLEMENTATION

We implemented EBUG in a system that supports bug reporting for Android apps. The system consists of three main modules. The GUI models generation module is written in Java and leverages Gator [12] to perform EBUG's static analysis, builds on UiAutomator [32] to complete the approach's dynamic analysis, and stores the GUI model using the Neo4J [33] graph database. The prediction models generation module leverages the Getevent [34] and UIAutomator utilities to collect user traces and builds on the MITLM toolkit [35] and the MonkeyLab infrastructure [36] to generate the n-gram based language models characterizing the GUI-action and GUI-element prediction models. Finally, the bug reporting module is implemented as a web application. This web application leverages the Google Cloud Natural Language infrastructure [37] to perform EBUG's natural language processing analysis and builds on the Quill library [38] to implement the interactive S2R reporting approach. The module uses the FastText model trained on the Common Crawl dataset [39] to identify the semantic similarity between S2R descriptions and GUI element properties. The web application also provides access to bug reports saved in a MySQL database.

5 EMPIRICAL EVALUATION

To determine the effectiveness and efficiency of EBUG, we used the implementation of our approach and performed two user studies. The studies' main focus is to assess the extent to which EBUG is able to facilitate the creation of highly reproducible bug reports. In the studies, we also compared the efficiency and effectiveness of EBUG with those of FUSION [4]. We selected FUSION as a baseline because the technique also aims to improve bug reporting, handles Android apps, and has been shown to create bug reports that are easier to reproduce compared to the ones

generated using traditional issue tracking systems [4]. The effectiveness of this reporting system [4] makes it a strong baseline against which to compare EBUG.

In the evaluation, we also compare the effectiveness of EBUG's GUI-action and GUI-element prediction models with those of prediction models based on AKOM [40] and CPT+ [41]. We selected AKOM as a baseline because the technique generates a model to perform sequence predictions, the technique functions effectively even without a large amount of training data, and the approach has been used to perform sequence predictions in numerous applications and various domains [40]. We selected CPT+ as it has similar characteristics as AKOM (e.g., the technique functions effectively even without a large amount of training data) and, in some domains, the technique was shown to provide better accuracy than AKOM [41].

In the evaluation, we targeted the following research questions (RQs):

- **RQ1**: Do developers using EBUG's reports reproduce more failures compared to when using FUSION's reports?
- RQ2: Is bug reporting with EBUG more efficient than
- RQ3: Do developers using EBUG's reports reproduce failures more efficiently compared to when using FUSION's reports?
- **RQ4**: Is bug reporting with EBUG leading to a better user experience than the one with FUSION?
- **RQ5**: Do developers using EBUG's reports have a better user experience compared to FUSION?
- RQ6: Do EBuG's GUI-action and GUI-element prediction models provide a better wasted effort score than prediction models based on AKOM and CPT+?

5.1 Experimental Context

In this section, we present the benchmarks used in the evaluation, describe the process used to collect user traces, and detail the user studies' characteristics.

5.1.1 Benchmarks

Our empirical evaluation is based on 20 unique, real-world failures from 11 apps. We identified the failures by randomly selecting bug reports from a collection of bug reports used in related work. These bug reports were used to evaluate different aspects of mobile app bug reporting [4], [42]. Before including a bug report in our benchmark dataset, we first confirmed that the failure described in the report was still reproducible. To reproduce the failure, we ran the app on a Nexus 7 emulator. The versions of the app and the Android system running on the emulator were the same as those used in related work. (We used a Nexus 7 emulator as FUSION's empirical evaluation also used this emulator.) The first 20 bug reports that we randomly selected were reproducible. Nineteen failures were reproducible on the emulator running Android 4.4. One failure was reproducible on the emulator running Android 5.0. Nine bug reports were used in FUSION's empirical evaluation [4], whereas the remaining eleven bug reports were used in recent work by Chaparro et al. [42]. All the 20 bug reports belong to free-text-based issue tracking systems. Among the failures described in the 20 bug reports, four failures manifest as crashes and sixteen failures provide an incorrect output to

TABLE 3

Apps, bug reports, and failures used in the empirical evaluation. *OIB* = observable incorrect behavior of the reports (O = output, C = crash).

App ID	App Name	Category	Version	Report ID	Failure ID	Min GUI Actions	OIB	
A01	AARD	Books	1.6.10	104	F01	4	О	
A02	ACV	Comics	1.4.1.4	<u>11</u>	F02	7	С	
A03	CAR REPORT	Auto	3.4.1	<u>43</u>	F03	10	О	
A04	Doc Viewer	Prod.	2.2	48	F04	9	О	
A05	DROIDWEIGHT	Hoath	1.5.4	<u>38</u>	F05	5	О	
A03	DROIDWEIGHT	1 leaut	1.5.4	<u>25</u>	F06	11	О	
Δ062	GNUCASH	Finance	1.4.3	247	F07	16	О	
71000	GNUCASII	1 marice	1.4.5	<u>256</u>	F08	20	О	
				<u>615</u>	F09	9	О	
A06b	GNUCASH	Finance	2.1.1	<u>620</u>	F10	11	О	
				<u>663</u>	F11	10	С	
				<u>49</u>	F12	15	О	
A07	MILEAGE	Finance	3.0.8	<u>53</u>	F13	12	0	
				<u>64</u>	F14	11	О	
A08	Notepad	Tools	1.06	<u>23</u>	F15	6	C	
A09	Olam	Books	1	2	F16	2	C	
A10	SCHEDULE	Books	1.32.2	<u>154</u>	F17	7	О	
				24	F18	11	О	
A11	TIME TRACKER	Prod.	0.17	<u>25</u>	F19	10	О	
				<u>46</u>	F20	8	О	

the user. Examples of output errors include displaying an erroneously computed value, incorrectly displaying a GUI element, and bringing the user to the wrong screen.

Comparing our dataset with the ones we used to create it, we can report the following. The dataset from the baseline technique we considered [4] includes 15 failures from 14 apps. That dataset has three failures that manifest as crashes and 12 failures that manifest as output failures. The second dataset [42] contains 24 failures from six apps. Four failures manifest as crashes and 12 manifest as output failures. We believe that the size and diversity between our dataset and these two other datasets are comparable.

Table 3 details some of the characteristics of the apps, bug reports, and failures. Specifically, column APP ID provides a unique identifier for the app, App Name reports the name of the app, Category provides the category of the app on the Google Play store [43], Version details the app's version used to reproduce the failure, Report ID reports the identifier of the bug report describing the failure, Failure ID provides a unique identifier for the failure, Min GUI Actions details the minimum number of GUI actions necessary to reproduce the failure, and OIB labels the type of observable incorrect behavior (i.e., crash or output failure). For the GNUCASH app, we use two different identifiers (A06a and A06b) because we reproduced the failures on two different versions of the app. The average and the median number of GUI actions necessary to reproduce the failures are 9.7 and 10, respectively. These numbers highlight that the failures we considered require multiple steps to reproduce the issue.

5.1.2 GUI Models and User Traces

EBUG uses pre-computed GUI models in its bug reporting phase. In the evaluation, we computed one GUI model for each app version considered. The generation of each GUI model is based on static and dynamic app analysis. The static analysis took, on average, 16 seconds to complete. For the dynamic analysis, we did not perform the analysis

exhaustively but, instead, used a timeout of 60 minutes, as this timeout was shown to be effective in related work [4].

In our experiments, we used user traces to generate the prediction models and gathered the traces by designing a data collection activity. In this activity, we invited nine participants to generate the traces. The participants' demographics include three participants without any computer science (CS) background, three undergraduate students in CS, and three graduate students in CS. In the activity, we assigned four apps to each participant, asked the participants to first get familiar with each app's functionality for five minutes, and then invited them to collect at least three traces for each app while exercising the app's functionality as they would do in a normal use of the app. In total, the participants collected 152 traces. The overall number of GUI actions in the traces is 3, 580. We asked participants to collect traces for both A06a and A06b as the apps' GUI differed significantly. Finally, the data collection activity took about two weeks to complete.

Based on this set of traces, we generated the prediction models following the methodology described in Section 3.1.2. Table 7 reports the n-gram order (column *Order* under EBUG headers) and the number of suggestions provided by each prediction (column *SN* under EBUG headers) for both the GUI-action and GUI-element prediction models.

5.1.3 Bug Reporting Study

To determine the extent to which EBUG is able to facilitate the creation of highly reproducible bug reports and how efficiently users can do so, we ran a user study involving ten participants. The participants' demographics include two undergraduate and eight graduate students in CS, and none of the participants took part in the data collection activity. Additionally, the participants had been studying CS for an average of 7.65 years (min=2, max=14), had 7.8 years (min=3, max=14) of programming experience, and had an average of self-determined experience level in mobile app programming of 4.7 (min=2, max=8) based on 1-10 assessment scale. None of the participants in this study is a professional developer, but four of the participants have industry exposure through internships as software engineers.

In the study, we asked the participants to report the failures from Table 3, using both FUSION and EBUG. Specifically, each participant reported five failures using FUSION and five failures using EBUG. We assigned failures to participants making sure that participants would not report the same failure with both tools and would not report two failures belonging to the same version of the app while using the same tool. We took these two measures to minimize bias in the reporting task. In the study, the subjects used anonymized variants of the tools, and we did not inform them about which tool was ours and which one was the baseline. It should be noted that we ensured that the order in which each participant was exposed to a tool was set such that half of the participants used EBUG first and half used FUSION first to mitigate potential effects of learning bias.

Before starting the study, we asked participants to watch two tutorial videos (about five-minutes long) explaining how to use FUSION and EBUG. (The participant could skip watching the tutorial videos if they chose to do so.) At the beginning of the study, we provided each participant with a document containing the failures the participant should report and the tool the participant needed to use to report a certain failure. We provided information about each failure in the form of a video. The video contained the sequence of GUI actions necessary to reproduce the failure and how the failure manifested itself. We recorded these videos by performing the GUI actions reported in Table 3 on a Nexus 7 emulator. In the document, we asked participants to report failures based on what they saw in the videos. We chose to expose bugs through video because we are interested in participants experience reporting bugs assuming that an end-user has a good understanding of the failure that they are reporting. As such, while exposing users to annotated videos of bugs may not represent a real-world scenario, it provided us with an avenue to clearly illustrate app failures in a manner in which we could be reasonably certain that the participant understood the bug. The document listed failures so that half of the participants would start reporting failures using FUSION, and the remaining half would start with EBUG. We randomized the order of the remaining failure-tool pairs in each document. The document also instructed participants to report the failures in the order they appeared in the document. During the study, we asked the participants to record the time it took them to report each failure and record their screen as they were reporting the failures. (We used the screen recordings to check whether the information entered by participants was the same as the one saved in the two tools' database. After watching the screen recordings, we could confirm that the information was the same.) Finally, during the study, the participants did not have access to the original bug reports of the failures.

At the end of the study, we asked the participants to complete a survey. In the survey, participants entered demographics information, recorded the time it took them to report each failure, and answered ten user experience (UX) questions for both FUSION and EBUG. Among the ten questions, six were usability questions, and four were user preference questions. Table 4 reports the usability (R-UN questions) and user-preference questions (R-PN questions) we asked to the participants. We formulated the usability questions based on the system usability scale (SUS) by John Brooke [44], as this scale is often used to compare the usability between systems [44]. The answers to these usability questions are structured using a five-level Likert scale that ranges from "Strongly disagree" to "Strongly agree". We formulated the user preference questions based on the user experience honeycomb developed by Peter Morville [45] and asked the participants to answer the questions using free-form text fields. Finally, the user experience questions were also used in the original evaluation of FUSION [4].

We ran the study over a three-week time frame and obtained a dataset of 80 bug reports. The dataset includes 40 bug reports generated by each tool and has two bug reports for each failure-tool pair.

5.1.4 Bug Reproduction Study

To determine whether the bug reporting study participants created reproducible bug reports, we ran a user study involving ten professional developers with experience in mobile app development from ten different companies. The developers had different levels of experience in mobile app

TABLE 4

User experience questions asked to participants in the bug reporting study. When asking questions about FUSION, we replaced <sys> with System B.

QID	Question
R-U1	I think that I would like to use <sys> frequently.</sys>
R-U2	I found <sys> very cumbersome to use.</sys>
R-U3	I found the various functions in <sys> were well integrated.</sys>
R-U4	I thought <sys> was easy to use.</sys>
R-U5	I found <sys> unnecessarily complex.</sys>
R-U6	I thought that <sys> was very useful for reporting a bug.</sys>
R-P1	What functionality did you find useful when reporting bugs with <sys>?</sys>
R-P2	What information (if any) were you not able to report with <sys>?</sys>
R-P3	What elements do you like the most from <sys></sys>
R-P4	What elements do you like the least from <sys></sys>

programming and bug reporting management. Specifically, the developers had an average programming experience of 9.4 years (min 2, max=12), had an average of selfdetermined experience level in mobile app programming of 5.4 (min=2, max=9) based on 1-10 assessment scale, and had an average of self-determined experience level in bug reporting and management of 6.8 (min=1, max=9) based on 1-10 assessment scale. Although the professional developers involved in the study have experience in app development, these developers were not the developers of the apps considered in the study. We opted not to involve the developers of the benchmark apps as we wanted to avoid including developers that could have been already aware of the reported bugs, which could have impacted the reproduction part of the study. In the study, we asked the developers to reproduce the bug reports gathered in the bug reporting study. In this study, we also asked the developers to reproduce the original bug reports (i.e., the bug reports listed and linked in Table 3). We assigned ten bug reports to each developer. Four bug reports were generated using FUSION, four were created using EBUG, and two were from the set of original bug reports. We assigned bug reports to developers making sure that they would not try to reproduce the same report. Similar to the bug reporting study, it should be noted that we ensured that the order in which each participant was exposed to each type of report was set such that one-third of the participants used EBUG first and one-third used FUSION first, and one-third used the original reports first to mitigate potential effects of learning bias.

During the study, the developers could access the bug reports using the anonymized variants of the tools. Like in the bug reporting study, we did not inform the subjects about which tool was ours and which one was the baseline. For the original bug reports, developers had access to a screenshot of the original bug reports. We took this measure so that the developers could not access any follow-up discussion in the bug tracking system of the original bug report. In the study instructions, we asked developers to reproduce the bug reports using a set of Nexus 7 emulators. These emulators contained the apps associated with the reports that the developers needed to reproduce. We also informed the developers that they needed to record the time it took them to reproduce each of the bug reports. We asked the developers to collect evidence (in the form of a screenshot or files stored in the emulator) that they reproduced a bug report (without including the time it took to perform this operation in the time they logged to reproduce a bug report).

TABLE 5

User experience questions asked in the bug reproduction study. QID = an identifier for the question, Question = user experience question.

QID	Question
D-U1	I think that I would like to use this type of bug report frequently.
D-U2	I found this type of bug report unnecessarily complex.
D-U3	I thought this bug report was easy to read/understand.
D-U4	I found this type of bug report very cumbersome to read.
D-U5	I thought that this type of report was very useful for reproducing the bug.
D-P1	What feature(s) did you find useful when reproducing the bugs with using this type of bug report?
D-P2	What additional information (if any) would you like to see in this type of bug report?
D-P3	What elements do you like the most from this type of bug report?
D-P4	What elements do you like the least from this type of bug report?

We used the evidence to verify that the subjects succefully reproduced the bug reports. Finally, we informed the developers that they should stop trying to reproduce a report if the report took more than 10 minutes to reproduce. To select the 10 minutes value, we relied on our experience from related work [4] and used the same value as in that work. We used the same value to provide a consistent evaluation of the techniques, and in related work, we also observed that the value had mitigated participant fatigue. Although setting the limit could provide a partial view of the total number of bug reports that could be reproduced, developers generally need to decide how much time they would like to invest in reproducing a report, and our results could also be interpreted as that time being set to 10 minutes.

Like in the bug reporting study, we asked participants to complete a survey at the end of the study. In the survey, developers entered demographics information, recorded the time it took them to reproduce each bug report, and answered nine user experience (UX) questions about FUSION, EBUG, and the original bug reports. Among the nine questions, five were usability questions, and four were user preference questions. Table 5 reports the usability (D-UN questions) and user-preference questions (D-PN questions). We formulated the questions following the same principles adopted in the bug reporting study.

5.2 Results

5.2.1 RQ1: Effectiveness

To answer RQ1, we identified how many failures developers could reproduce using the reports generated in the bug reporting study. Table 6 provides this information. Specifically, column *Repro* under the EBUG header reports which failures developers could reproduce using the reports generated with EBUG. Column *Repo* under the FUSION header reports the results for FUSION. Table 6 also reports the identifier of the developer that tried to reproduce a certain report (columns *Dev ID*).

Overall, developers using EBUG's reports could successfully reproduce the report's failure in 37 out of 40 cases (92.5% success rate). Developers using FUSION's reports could do so in 32 out of 40 cases (80% success rate). After inspecting the bug reports, we explain this result with the fact that bug report participants created more complete bug reports while using EBUG. This characteristic can be seen in Table 6 by comparing the *S2Rs* columns in the EBUG and FUSION sections of the table. Specifically, participants reported a total of 364 S2Rs using EBUG and 339 using

TABLE 6

Bug reporting and bug reproduction studies results. $App\ ID$ = identifier of the benchmark app, $Fail.\ ID$ = identifier for the failure in the benchmark app, $Rep\ ID$ = participant identifier in the bug reporting study, S2Rs = number of steps to reproduce in the bug report submitted by the participant, T_{report} = time taken by the participant to submit the bug report, $Dev\ ID$ = developer identifier in the bug reproduction study, T_{repro} = time taken by the participant to submit the bug report, Repro = \checkmark if a developer reproduced the bug report (\checkmark otherwise), $Val\ S2Rs$ = number of S2Rs validated by EBug, $CAS\ A$ = number of GUI action suggestions (GAS), $CAS\ A$ = number of GUI action suggestions accepted by the participants of the bug reporting study, $CAS\ A \rightarrow E$ = number of GUI

Action suggestions accepted and then edited by the participants of the bug reporting study, $GAS A \rightarrow D$ = number of GUI Action suggestions accepted and then deleted by the participants of the bug reporting study, GUI Element Sugg. = GUI element suggestions (GES), GESP = number of GUI Element suggestions provided to the participants of the bug reporting study, GESA = number of GUI Element suggestions accepted by the participants of the bug reporting study, $GESA \rightarrow E$ = number of GUI Element suggestions accepted and then edited by the participants of the bug reporting study, $GESA \rightarrow D$ = number of GUI Element suggestions accepted and then deleted by the participants of the bug reporting study, $GESA \rightarrow D$ = number of particle suggestions accepted by the participants of the bug reporting study.

				Fusi	EBUG																		
App ID	Fail.		Repor			eproducing	\exists	Reporting Reproducing															
	ID.	Rep			Dev		5	Rep		Val	Gl	II Ac	tion Si			I Elei	ment S	ugg.	D.C.		Dev		
		ΙĎ	S2Rs	T_{report}	ID	T_{repro}	3	ΙĎ	S2Rs	S2Rs	P	Α		A→D	Р	Α		$A \rightarrow D$	PS_A	T_{report}	ID	T_{repro}	Rep
_	H	R03	4	9m10s	D10	1m56s	=	R02	5	3	5	1	1	1	1	1	0	0	0	4m58s	D02	2m36s	-
A01	F01	R08	4	9m14s	D03			R10	4	3	3	0	1	0	3	1	0	0	0	2m30s		1m13s	1
		R04	4	7m18s	D09	4m30s		R02	6	4	1	0	0	0	5	4	1	0	0	5m17s		1m45s	-
A02	F02	R08	6	6m20s	D04	2m37s	.	R09	6	4	7	1	0	5	15	3	0	11	0	9m42s		59s	1
		R05	8	7m48s	D09	4m10s	_	R02	10	10	9	0	3	0	7	7	0	0	2	4m55s		1m04s	
A03	F03	R10	10	5m40s	D01	9m55s	ا ۔ ا	R09	16	16	14	3	5	0	6	5	0	0	0	6m42s		1m47s	
101	F0.4	R05	7	6m34s	D07	10m00s		R07	8	4	0	0	0	0	12	6	1	3	1	7m29s		5m40s	
A04	F04	R06	4	11m57s	D02	3m33s	- 1	R01	10	7	6	1	0	0	4	2	0	2	0	6m27s		2m30s	1
	E05	R03	5	7m36s	D07	1m40s	-	R07	5	5	1	0	0	0	5	4	0	1	1	3m42s		45s	
4.05	F05	R08	3	4m36s	D01	10m00s	x	R01	5	4	2	0	0	0	3	2	0	0	0	4m07s	D05	1m13s	
A05	For	R09	7	10m38s	D02	4m18s		R06	5	5	2	0	0	0	5	5	0	0	1	8m02s		10m00s	
	F06	R02	5	2m50s	D06	l I	x	R04	5	5	0	0	0	0	3	3	0	0	1	3m15s		1m20s	
	FOF	R03	13	16m51s	D09	3m30s	7	R02	12	11	15	5	2	1	6	4	0	0	0	10m30s		2m26s	_
100	F07	R07	12	21m29s	D04	3m10s	/	R09	9	9	31	6	3	11	5	1	0	1	0	20m11s		1m07s	1
A06a	F08	R05	13	19m01s	D07	10m00s	x	R04	15	11	14	2	0	4	20	6	0	10	4	18m40s	D04	7m36s	1
		R10	19	17m20s	D02	6m19s		R08	13	12	5	0	2	0	13	10	0	0	0	8m11s	D09	10m00s	X
	F09	R03	9	10m33s	D06	10m00s	X	R10	9	8	12	4	0	2	6	4	0	0	0	3m35s	D08	3m18s	1
	1.09	R09	9	14m12s	D03	3m30s	/	R01	7	6	7	2	0	1	4	1	0	3	0	4m07s	D01	3m00s	1
A06b	E10	R02	8	12m29s	D08	8m10s	7	R05	9	9	2	0	0	0	13	8	0	3	0	8m14s	D05	3m35s	1
AUUU	F10	R06	10	21m15s	D01	3m15s	/	R09	7	7	12	1	0	3	12	4	0	5	0	10m04s	D10	2m55s	1
	F11	R04	6	11m52s	D10	1m08s	7	R08	10	9	3	0	0	0	9	8	0	1	0	5m03s	D06	5m10s	1
	1.11	R07	15	12m36s	D05	2m33s	/	R03	11	7	9	4	0	2	2	1	0	0	0	9m49s	D03	3m30s	1
	F12	R08	13	10m03s	D01	10m00s	X	R07	22	20	5	0	0	0	16	14	0	0	7	7m32s	D06	3m00s	1
	1.17	R05	11	4m50s	D09	8m50s	/	R01	21	12	2	0	0	0	3	3	0	0	0	6m23s	D03	4m00s	1
A07	F13	R02	16	7m55s	D10	1m58s	7	R03	17	16	15	7	6	0	2	2	0	0	0	6m22s	D02	2m18s	1
1107	113	R06	6	8m32s	D05	3m01s		R08	10	9	0	0	0	0	7	7	0	0	0	5m01s	D07	1m22s	1
	F14	R04	11	5m37s	D08	3m26s	7	R05	12	10	1	0	0	0	12	11	0	1	1	5m13s		2m50s	
	111	R09	16	10m08s	D03	7m00s		R10	10	10	8	2	4	1	5	4	0	0	0	3m33s	D09	3m15s	_
A08	F15	R10	6	4m03s	D05	1m14s	/	R04	5	5	3	1	0	0	3	3	0	0	1	2m48s	D02	1m22s	1
7100	1 13	R01	7	5m00s	D08	1m01s		R06	6	6	7	4	1	0	3	1	0	1	0	4m58s		25s	
A09	F16	R07	3	2m40s	D05	3m27s	- 1	R03	2	2	3	1	0	0	1	1	0	0	0	2m03s		1m00s	
1107	110	R01	2	1m49s	D06	16s	_	R06	2	2	1	0	0	0	2	2	0	0	1	2m47s		29s	_
A10	F17	R06	7	9m11s	D04		×	R07	7	7	0	0	0	0	7	5	0	0	0	5m38s		35s	
1110	1	R01	7	6m19s	D06	2m11s	_	R05	7	7	3	0	0	0	10	7	0	1	0	3m04s		5m00s	_
	F18	R01	4	2m57s	D07	5m51s	- 1	R10	10	9	0	0	0	0	9	8	0	1	0	2m41s		4m10s	1
		R07	9	7m41s	D03		X	R04	7	6	0	0	0	0	2	0	0	2	0	5m28s		10m00s	
A11	F19	R02	8	13m45s	D08	2m00s		R05	7	5	0	0	0	0	11	5	0	0	0	5m45s		2m34s	1
		R10	11	5m55s	D04	7m22s	_	R06	8	6	0	0	0	0	10	6	0	0	1	7m40s		48s	
	F20	R04	10	4m27s	D10			R08	12	12	0	0	0	0	9	9	0	0	0	3m31s		45s	
		R09	11	7m02s	D02	3m42s		R03	12	12	1	1	0	0	7	7	0	0	0	5m50s		2m58s	
		-	339	365m13s	-	201m13s 3	32	-	364	315	209	46	28	31	278	185	2	46	21	251m47s	-	120m20s	37

FUSION. Additionally, the result can also be related to the fact that 315 out of the 364 S2Rs reported using EBUG, were also validated by the technique. As another explanation of the result, the clarity of the the S2Rs provided by EBUG also appeared as a comment in the survey answers of seven out the 10 developers. Finally, developers using EBUG's bug reports could reproduce all of the 20 failures we selected for the empirical evaluation.

We manually analyzed the three reports submitted using EBUG that could not be reproduced by developers and

identified that all three reports had missing S2Rs. The first report (F11 submitted by R04) had two missing S2Rs, the second report (F06 submitted by R06) had six missing S2Rs, and the third report (F08 submitted by R08) had seven missing S2Rs. Because of the missing S2Rs, two of the reports (F06 submitted by R06 and F08 submitted by R08) had some S2Rs that could not be validated by EBUG. In the third report (F11 submitted by R04), all S2Rs could be validated by EBUG as the missing S2Rs corresponded to

GUI actions in the screens as already validated S2Rs, and therefore EBUG did not identify issues with the sequence of S2Rs. Although the developers did not provide the reasons for which they could not reproduce these reports in their exit surveys, we believe that these missing S2Rs might have affected the bug reproduction outcome. We believe that automated bug reproduction techniques could be combined with EBUG to provide feedback on whether the reports could be reproduced upon submission, and this feedback might incentivize users to refine their reports.

In the bug reproduction study, we also asked developers to reproduce the failures using the original bug reports. Developers were able to do so in 15 out of 20 cases (75% success rate). These results show that EBUG is more effective than FUSION and the original bug tracking systems in providing reproducible bug reports to developers.

In summary, we believe that the results show that EBUG can be effective in providing reproducible bug reports and that EBUG is more effective than FUSION.

5.2.2 RQ2: Bug Reporting Efficiency

To answer RQ2, we compare the amount of time the participants took to submit a bug report using FUSION and EBUG. For each bug reporting task, Table 6 reports the time (columns T_{report}) participants took to report a specific failure.

In total, when using EBUG, the participants submitted the 40 bug reports in 251m47s (about 6m18s per report). When using FUSION instead, the participants took 365m13s to report the same failures (about 9m8s per report). These results show a decrease of 31.06% in the bug reporting time when using EBUG. If we only considered successfully reproduced bug reports, the average bug reporting time for EBUG is 6m14s and for FUSION is 9m13s. These numbers indicate that the decrease in bug reporting time not only applies to all bug reports but, in particular, also applies to successfully reproduced ones.

The bug reporting time of EBUG is characterized by a number of suggestions provided by EBUG and accepted by the participants of the study. For both GUI action and GUI element suggestions (GUI Action Sugg. and GUI Element Sugg. headers), Table 6 reports provided (P), accepted (A), accepted and then edited $(A \rightarrow E)$, and accepted and then deleted $(A \rightarrow D)$ suggestions. In total, EBUG provided 209 GUI Action suggestions, and participants accepted 46 suggestions (without any modification). The participants also accepted 28 additional suggestions and then edited the suggestions. In the modifications, the participants refined the content of the suggestion by providing additional information in the S2R. The participants also accepted 31 suggestions that they then removed from the bug report. One participant (R09) rewrote part of two bug reports multiple times before submitting the results, and this contributed to 12 suggestions being deleted. Although deleted suggestion could represent suggestions that were not ultimetly helpful for participants, deleted suggestions might be leveraged by automated bug reproduction techniques [11], [23] to guide the exploration performed by the techniques to reproduce bug reports. The number of GUI element suggestions is 278. The participants accepted 185 suggestions, accepted and then edited two suggestions, and accepted and then

deleted 46 suggestions. The GUI element suggestions reported in Table 6 only include suggestions provided by the GUI element prediction model. As a side note, the participants also accepted 42 suggestions created for partially written GUI elements. (These suggestions are based on text matching and not the GUI element prediction models). The participants also accepted 21 particle suggestions (PS_A). The small number of accepted particle suggestions is due to the significant number of accepted GUI action and element suggestions. The usefulness of the suggestions provided by EBUG appeared as a comment in the survey answers of nine out the ten participants and the comments highlighted that the suggestions made reporting faster.

Analyzing failures individually, we identified two cases (F02 and F16) in which the average time to report the failure (the study had two submitted bug reports per failure and tool) was lower in FUSION as compared to EBUG. The first failure (F02) is an output failure and the second one (F16) is a crash. In the first case, the average reporting time with EBUG was higher because one of the two users (R09) had a significantly higher reporting time. Although we do not know the exact reason behind this higher reporting time, we observed that the user rewrote four S2Rs three times because the subsequent S2R could not be validated by EBUG. This result shows that improving the accuracy of EBUG's GUI model, which is used to validate S2Rs, can further improve EBUG's efficiency. The second failure (F16) is the failure caused by the shortest number of GUI actions (two). In this case, the average reporting time with EBUG (2m25s) is slightly slower than FUSION (2m15s). We believe that EBUG does not perform better than FUSION in this case because users cannot take full advantage of EBUG's suggestions.

Give the results of our bug reporting study, we can conclude that EBUG allows users to submit bug reports more efficiently than FUSION.

5.2.3 RQ3: Bug Reproduction Efficiency

To answer RQ3, we compare the time developers took to reproduce FUSION's and EBUG's bug reports. Table 6 reports the time (columns T_{repro}) developers took to reproduce the failure described in the bug reports.

Overall, developers took 120m20s to reproduce EBUG's reports and 201m13s to reproduce FUSION's reports. These numbers identify a 40.2% decrease in the time to reproduce bug reports when EBUG's reports. (Note that this number is affected by the 10m time limit set for the bug reproduction task.) As in RQ1, we explain this improvement due to the fact that EBUG provides more complete bug reports with respect to FUSION. If we only consider bug reports successfully reproduced by developers, developers could reproduce EBUG's bug reports in 2m26s and FUSION's bug reports in 3m46s on average. This result highlights that developers can reproduce EBUG's bug reports faster even without considering the cases where the bug reproduction task is not successful. Finally, developers could also reproduce EBUG's bug reports faster than the original bug reports, which took 3m5s on average in the successful cases.

Considering failures for which developers could reproduce both participant-submitted bug reports (the study had two submitted bug reports per failure and tool), developers using bug reports submitted with EBUG were slower in only

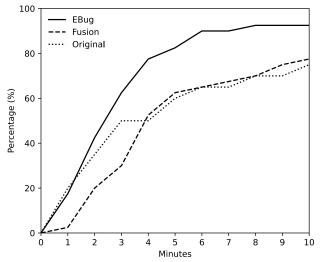


Fig. 13. Percentage of bug reports reproduced by developers over time.

one case (F11). To report this failure, users had to provide a value from a dropdown list that contained a large number of values (225). Because EBUG also allows users to report this information just by using text instead of specifically selecting a value from a suggestion, reporters using EBUG provided this information using a textual description without going through the technique's suggestions. The description provided by both users was not complete, and we believe that this aspect might have caused the higher reproduction time with EBUG, as developers had to identify the right value among a large number of values. This result highlights a tradeoff between flexibility in the bug reporting process and efficiency in bug reproduction time. We believe that this tradeoff is worth investigating in future work through additional user studies.

Fig. 13 reports the percentage of bug reports reproduced within n minutes (where n ranges from 0 to 10). The graph confirms that developers could reproduce more bug reports and faster using EBUG's reports as compared to FUSION's reports and the original reports. The only case where developers were not the fastest using EBUG's reports is within the first minute, where 20% of the original reports were reproduced as opposed to 17.5% of EBUG's reports. We believe that this last result can be explained by the fact that some of the original reports were slightly quicker to process as they contained less information, but in the following minutes, when relevant S2Rs were necessary to reproduce the failures, EBUG outperformed both FUSION and the original reports.

Because reports for the same failure were submitted by different participants and reproduced by different developers, it is not possible for us to compare results individually and identify exactly why it took longer for a certain failure to be reproduced using reports from EBUG. However, analyzing the content of submitted reports, we believe that missing S2Rs and the developers' experience could be factors that affected the failure reproduction time.

In summary, we can conclude that EBUG allows developers to reproduce bug reports more efficiently than FUSION and the issue tracking systems of the original bug reports.

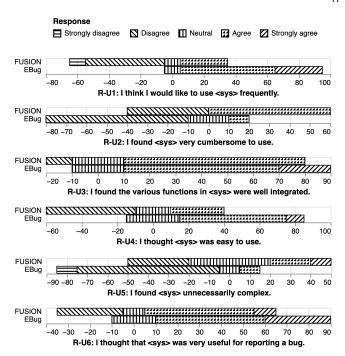


Fig. 14. Usability answers in the survey of the bug reporting study.

5.2.4 RQ4: Bug Reporting User Experience

To answer RQ4, we collected and analyzed the participants' answers to the exit survey questions of the bug reporting study. We report the answer to the usability questions (R-UN questions in Table 4) in Figure 14 using centered stacked bar charts. For questions R-U1, R-U3, R-U4, and R-U6 agreement is better, while for questions R-U2 and R-U5 disagreement is better.

Overall, the participants provided very positive answers to the usability questions related to EBUG. Nine of the ten participants either agreed (6) or strongly agreed (3) with the statement "I think I would like to use EBUG frequently" (R-U1). Participants also provided positive answers (5 agree and 3 strongly agree) in relation to the statement "I thought that EBUG was very useful for reporting a bug" (R-U6). Among all answers provided by all of the participants, only one was negative. Specifically, one participant agreed with the statement "I found EBUG very cumbersome to use" (R-U2). The same participant had a neutral answer in relation to the statement "I thought EBUG was easy to use" (R-U4). We believe that this answer was related to the fact that the participant experienced networking issues during the study, and our system relies on a network connection to operate.

In relation to user experience with EBUG, we also want to report a few excerpts from the free-text answers we received from the participants when they answered the user preference questions (R-PN questions in Table 4). Some answers indicate that the suggestions provided by EBUG make the bug reporting process more efficient: "[...] Autocomplete really speeds up the process [...]" (participant R02), "[...] The context-aware list of components helps finding the one I'm thinking of [...] Templates of step description reduce a lot of typing" (R03), and "[...] I like the action and GUI components' suggestions. That functionality is very useful for saving time [...]" (R6). Finally, some answers indicate that EBUG is useful and a natural way to submit bug reports: "I liked that I was basically typing a paragraph of instructions that happened

to automatically link with images." (R01), "It feels natural while reporting bugs with [EBUG] [...]" (R05), and "I felt that this tool was very useful and would prove very easy to use, even if you had no coding experience [...]" (R07).

The free-text answers also provided useful feedback for improving EBUG and bug reporting through future work. A comment (from participant R03) mentioned that "If the app contains same-looking components on multiple screens, the components list shows me multiple options, when there should be only one valid option if the tool considers the current state of the app". This situation can appear when EBUG is not able to validate previously typed S2R. In this case, EBUG provides a suggestion that contains all available GUI actions in the app (ordered by their distance in the GUI model with respect to the last validated step). This feedback highlights the opportunity to improve EBUG in two ways. First, future research could define techniques to create an even more accurate GUI model. Second, future research could focus on user studies in which users receive a limited number of options in the suggestions based on a pre-defined GUI model distance value (and not include all available GUI actions as the options). A second comment (from R05) highlighted that "[...] external controls such as power button, or emulator controls for changing orientation were not accessible.. Power button support is currently not provided by EBUG, but it could be added by representing the operation in the GUI model. For both types of actions, a user might not receive such suggestions also because the prediction model did not contain such actions. Having a more complete prediction model could be achieved by collecting additional user traces, but there is always the possibility of not having some actions in the model. Considering the results obtained while assessing EBUG's bug reporting efficiency, we believe that the technique has good results even if some actions are not present in the prediction model. Another comment (from R06) reported that "The selected target screen is unnecessary. At least in my case, I never used it at all.". The selected target screen is the one in the lower-right part of Fig. 9. Future user studies related to EBUG's usage could focus on assessing the usefulness of each of the technique's interface components.

In summary, these preliminary results demonstrate that the participants had a positive user experience when submitting reports with EBUG, and this experience was better than the one with FUSION.

5.2.5 RQ5: Bug Reproduction User Experience

We answer RQ5 by analyzing the developers' answers to the exit survey of the bug reproduction study. We report the answer to the usability questions (D-UN questions in Table 5) in Figure 15 using centered stacked bar charts. For questions D-U1, D-U3, and D-U5 agreement is better, while for questions D-U2 and D-U4 disagreement is better.

Also for RQ5, the subjects provided very positive answers to the usability questions related to EBUG. All developers either agreed (four) or strongly agreed (six) with the statement "I think that I would like to use this type of bug report frequently" (D-U1). All developers also agreed (three) or strongly agreed (seven) with the statement "I thought this bug report was easy to read/understand" (D-U3). Additionally, developers provided positive answers (two agree and seven strongly agree) in relation to the statement "I thought that

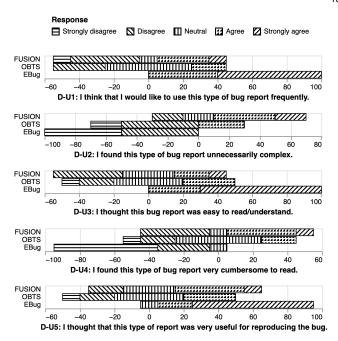


Fig. 15. Usability answers in the survey of the bug reproduction study. *OBTS* is the original bug tracking system.

this type of report was very useful for reproducing the bug" (D-U5). If we consider the participants' answers to all the questions, we found that all the answers were either neutral (two answers) or positive (48 answers). Developers did not provide as positive answers for FUSION's usability as the ones provided for EBUG. In fact, some participants disagreed with statement D-U1 (four disagree and one strongly disagree) and D-U5 (two disagree). Figure 15 also shows that the answers about the original bug tracking systems are not as positives as the ones provided for EBUG.

The developers' answers to the user preference questions (D-PN questions in Table 5) suggest that EBUG's reports are actionable and easy to understand: "Simple step details, ordered and simple" (developer D03), "list of operation is very easy to understand" (D04), "I found both the summary of the bug on the left side, with expected vs observed behavior to be useful to provide context. [...] I think the list of steps with screenshot on the right is what makes this type of bug report system particularly effective" (D06), "I liked that it was clear and with the right amount of information to reproduce the steps. The list of actions/screenshots is the most useful." (D07), and "The screenshots were very useful in finding the bug [...] (D10).

The free-text answers from the bug reproduction study also provided useful insights for improving EBUG and bug reproduction in future work. A comment (from developer D01) suggested that an "[i]mage of the final (buggy) state of the app" would be a good addition to EBUG. Currently, EBUG does not focus on improving the bug reporting process of the observed behavior. However, we believe that automated, real-time, and iterative bug reproduction techniques could help users report a screenshot of the observed behavior and believe that this is an interesting direction for future work. A second comment (from D02) mentioned that "I think the description field doesn't have relevant information if it's a copy of the steps, but could be useful to hold additional information". The description field mentioned by the developer is the description typed by the reporter. The developer found this description redundant as compared to the list of S2Rs

TABLE 7

Prediction models comparison. *App ID* = app identifier, *Traces* = number of traces , *Predictions* = number of predictions, *Order* = order of AKOM's and EBUG's models, *SN* = number of suggestions provided by the models, *wes* = wasted effort score, *PL* = sequence prefix length for CPT+.

		Predictions			G	Actio	n Mod	els		GUI Element Models										
App ID	Traces		AKOM			CPT+			EBug			A		CP'	Γ+	EBUG				
			Order	SN	wes	PL	SN	wes	Order	SN	wes	Order	SN	wes	PL	SN	wes	Order	SN	wes
A01	12	285	4	1	1.159	4	1	1.568	3	1	1.021	4	1	0.727	3	1	2.314	5	1	0.727
A02	15	283	7	1	0.85	3	1	1.55	5	1	0.952	6	1	0.626	3	1	2.291	9	1	0.675
A03	14	458	4	1	0.877	2	1	1.018	4	1	0.885	6	1	0.69	3	1	3.362	8	1	0.703
A04	15	284	2	1	2.595	2	1	3.239	3	1	3.239	4	1	1.427	2	1	8.793	5	1	1.513
A05	12	211	1	1	1.542	3	1	1.931	3	1	1.638	2	1	1.542	3	1	4.024	4	1	1.269
A06a	15	464	2	1	1.549	2	1	1.698	7	1	1.275	4	1	1.468	3	1	8.098	9	1	1.275
A06b	14	398	1	1	1.355	2	1	1.745	6	1	1.151	4	1	0.98	3	1	4.169	9	1	0.961
A07	13	285	2	1	0.717	5	1	0.759	3	1	0.686	4	1	0.647	3	1	2.132	5	1	0.566
A08	12	250	3	1	0.799	2	1	0.88	4	1	0.603	8	1	0.563	3	1	0.908	8	1	0.429
A09	9	126	2	1	0.636	2	1	0.909	2	1	0.575	4	1	0.313	3	1	1.1	5	1	0.355
A10	11	265	2	1	2.193	2	1	2.581	2	1	2.118	3	1	1.325	2	1	4.096	4	1	1.284
A11	10	271	3	1	1.221	2	1	1.398	5	1	1.117	6	1	0.936	2	1	4.42	5	1	0.726

validated by EBUG. The approach displays the information to the developer in case S2Rs were not suitably validated, and, in this way, the developer could always refer to the description originally provided by the user. We believe that additional user studies could further improve EBUG's interface, and these studies could possibly display the user description on an on-demand basis. Finally, another comment (from D04) highlighted that images associated with S2Rs are not always present, and this increases the difficulty in understanding the reported bug: "[n]ot always image[s] are present, this increase[s] [the] difficult[y] to understand". This situation can appear when EBUG is not able to validate an S2R. Devising techniques to improve EBUG's GUI model would mitigate this type of situation.

The answers provided in the study offer some preliminary results on the user experience with EBUG and reveal that participants had a positive user experience when reproducing bug reports with EBUG. Additionally, the results provide initial evidence that this experience was better than the one with FUSION's reports or the original bug reports.

5.2.6 RQ6: Prediction Models' Effectiveness

To answer RQ6, we leveraged the user traces collected for EBUG's evaluation and compared EBUG's prediction models with prediction models generated using AKOM [40] and CPT+ [41]. Specifically, we compared the prediction models of EBUG used in the bug reporting study with the best performing models generated using AKOM and CPT+. To compute AKOM and CPT+ prediction models, we follow a similar methodology as the one described in Section 3.1.2. For AKOM models, we identified the best performing model for each app by varying the model order and doing leaveone-out (sequence) cross-validation. We varied the sequence prefix length to identify the best performed CPT+ models. We use the wasted effort score (wes) metric to find the best performing models for a specific technique and compare models across approaches. The wasted effort score is an extrinsic evaluation metric and we used this metric as extrinsic evaluations are the most suitable way to evaluate models in sequence prediction tasks [16].

Table 7 reports the performance of the models. For each app, the table reports the best performing models for predicting GUI actions and GUI elements using AKOM, CPT+, and EBUG. Table 7 is divided into three sections. The

first section reports the identifier (column App ID) of the benchmark apps considered in the evaluation, the number of traces (Traces) to evaluate the models, and the number of predictions used to evaluate the models (Predictions). The GUI action and GUI element models were evaluated using the same number of predictions because when we evaluated GUI element models, we only assessed predictions on the tokens representing GUI elements (see Section 3.2 for more details). The second section reports the best performing models for making GUI Action suggestions and the third section for making GUI element suggestions. Columns Order provide the order of the models, columns SN provide the number of suggestions provided by the models for each prediction task, columns wes provide the wasted effort score for the models, and columns *PL* provide the sequence prefix length for the CPT+ models. Table 7 reports in bold the wasted effort score of the best performing model across all models for a certain app.

Overall, EBUG provides the best performing models in 16 out of the 24 cases considered (one of the cases is a tie with a model based on AKOM). AKOM offers the best performing model in nine out of the 24 (including the tie). CPT+ never provides the best performing model. If we considered the wasted effort score of the GUI action prediction models across all apps, EBUG's wes is 1.127 and AKOM's wes is 1.198. For GUI element prediction models, EBUG's wes is 0.846 and AKOM's wes is 0.902.

Although the overall difference between AKOM models and EBUG's n-gram models is not big, the n-gram models performed better than the AKOM models both in terms of wasted effort and correctly made predictions. Across all the GUI action prediction tasks, EBuG's n-gram models produced 1,683 correct predictions and a wasted effort of 1,897. AKOM models led to 1,629 correct predictions and a wasted effort of 1,951. For GUI element prediction tasks, EBUG's *n*-gram models produced 1,939 correct predictions and a wasted effort of 1,641. AKOM models led to 1,882 correct predictions and a wasted effort of 1,698. Even if these differences are not big, when considering that the same failure might be reported by multiple users, the differences can become more significant. Considering these results, EBUG's n-gram models can help reporters by providing more useful predictions while lowering the wasted effort. However, additional studies are also needed to confirm and

extend these results and we suggest those as possible future work.

Based on these results, we can conclude that EBUG provides the best performing prediction models in the majority of the cases, and when we consider all models of all apps, EBUG has the best wasted effort score as compared with baseline approaches.

5.3 Threats to Validity

As it is the case for most empirical evaluations, there are both external and construct threats to validity associated with our results. In the bug reporting study, the participants were CS students. These participants might have a different background compared to other bug reporters. To mitigate this threat to validity, the study included participants from three institutions and at a different level of their studies. Additionally, past work has illustrated that performance of CS students and professional developers do not differ dramatically in some software engineering tasks [46]. The results of the bug reporting study might also be affected by the user traces collected in the data collection activity. To mitigate this threat, we invited participants with different backgrounds (some of which did not have a CS background) and asked them to use the benchmark apps as they would do in a normal use of the apps. The results of the bug reproduction study might also not generalize to other developers. To mitigate this threat to validity, the study included developers from ten different companies at different career stages. In the bug reproduction study, we used 10 minutes as a time limit for developers to reproduce bug reports. This value could provide a partial view of the total number of reproduced bug reports as more bug reports could be reproduced after the time limit. We used this value to be consistent with related work [4]. We also decide to use the value, as, in related work, we also observed that the value had mitigated participant fatigue.

Another threat to external validity is that our results might not generalize to other failures or apps. In particular, we only considered 20 failures in our empirical evaluation. This limitation is an artifact of the complexity of our evaluation. To mitigate this threat, we used randomly selected real-world failures of varying types and complexity. These failures belong to real-world apps from different categories and functions. These apps include apps such as GNUCASH, MILEAGE, and DOC VIEWER, which have complex functionality, hundreds of widgets, and hundreds of thousands of users. We believe that, given the complexity of the failures and apps considered in the studies, EBUG should also apply to other types of failures and apps.

Finally, in terms of construct validity, there might be errors in the implementation of our approach or our experimental infrastructure. To mitigate this threat, we extensively inspected the results of the evaluation manually.

6 Discussion

Through our empirical evaluation, we found that EBUG is more efficient and effective than the baseline. Based on the evaluation results and the feedback received from the participants of the bug reporting study, we believe that the

results can be attributed to the S2R suggestions and the validation operations performed by EBUG. For one of the failures considered, bug reporting with EBUG was not more efficient. In this case, one of the bug reporters rewrote four S2Rs three times. We believe that this situation appeared because EBUG could not validate a subsequent S2R in the report. This result reveals that improving the accuracy of the GUI models generated by EBUG, which are used to validate the S2Rs, can further improve EBUG's efficiency. Developers using EBUG's bug reports were generally more efficient in reproducing bug reports as compared to developers using the baseline's bug reports. Developers using bug reports submitted with EBUG were slower only in the case of one specific failure. To report this failure, the reporters had to provide a value from a dropdown list that contained a large number of values, and the reporters decided to use free text instead of inspecting a list of options provided by EBUG. The description provided by both users was not complete, and we believe that this aspect might have caused the higher reproduction time with EBUG. This result shows a tradeoff between flexibility in the bug reporting process and efficiency in bug reproduction time. We believe that this tradeoff is an interesting aspect related to bug reporting and should be the target of future user studies.

During our evaluation, we focused on examining how useful EBUG is for reporting bugs, however, it is also important to comment on the general scalability and applicability of our technique. Given that a majority of the "heavy" processing required by EBUG is performed a-priori, that is before the reporting process commences, and the "online" portions of EBUG generally scale well, we expect EBUG to scale reasonably well for a variety of complex applications. The one potential limitation with scaling may be with EBUG's GUI-component matching procedure. That is, if there is a very large number of components present within an app, searching for the corresponding component may introduce a short delay in the reporting interface. However, while we cannot confirm EBUG's performance outside the context of our study, we speculate that, based on the authors' experience of the general size of most Android applications, EBUG would generally scale to a majority of apps available today.

Another important point of investigation to understand the reasons for EBUG general performance in our empirical evaluation is the extent to which static and dynamic GUI models contribute to the accuracy of the suggestion engine. To identify the contributions of the static and dynamic GUI models (which together form the GUI model used by EBUG), we analyzed the S2Rs that were entered by the participants of the bug reporting study and that were also validated by EBUG. As reported in Table 6, the technique validated 315 S2Rs across all apps. One of those S2Rs had the corresponding GUI action that was present only in the static GUI model, 140 S2Rs had corresponding GUI actions appearing in both the static and the dynamic GUI model, and 174 S2Rs had corresponding GUI actions present only in the dynamic GUI model. This result highlights that the dynamic GUI is of primary importance for EBUG. If we look at all the GUI actions (1,255) encoded in the GUI models of all the apps considered, 155 GUI actions are present only in the static GUI models, 274 appear in both the static and the dynamic

GUI models, and 826 are present only in the dynamic GUI models. Although a large number of GUI actions are part of the dynamic GUI models, the number of GUI actions provided by the static GUI model is not negligible, making them potentially useful for validating S2Rs.

The predictions accepted during bug reporting are another important aspect of EBUG. In our study we found that participants accepted a large number of both GUI action and GUI element suggestions. The wasted effort score during bug reporting was higher for GUI action predictions and lower for GUI element predictions as compared to when models were built (i.e., when evaluating the models on user traces). Although the wasted effort scores associated with GUI action prediction models are higher during bug reporting, the number of accepted suggestions, the study participant's feedback, and the overall lower bug reporting time of EBUG suggests that the predictions can be helpful during bug reporting. We believe that performing additional studies on wasted effort during bug reporting is an interesting venue for future work.

Another important aspect of EBUG's practical applicability is concerned with the collection of the app usage models. One of the key facets of EBUG is the utilization of "realworld" usage traces of a given application to help prioritize potential predicted S2Rs. There are several ways in which such traces could be collected. For example, as developers release beta versions of an application to crowd-testers, traces could be collected from crowd-testers [47]. Additionally, traces could be collected from developer written test cases, or even usages collected by developers themselves during the testing process of each version of their app. EBUG uses traces to gather an understanding of commonly performed GUI action sequences and provides suggestions when those GUI action sequences are part of a sequence that reveals a bug. For this reason, EBUG does not need to collect traces that expose bugs to provide suggestions.

There are many potential future directions of work for improving EBUG. For instance, currently our execution traces are tied to a specific application, meaning that data collection must be done on a per-app basis. However, recent work from the Ubicomp community has illustrated the benefits of modeling user behaviors across applications [48]. This suggests that future work could look to learn a generalized abstract usage model across many applications that is supplemented with a smaller amount of app-specific data from a diverse set of users. This may help to improve the need of collecting usage data for EBUG. We also believe future work could examine different types of reporting interfaces, such as on-device reporting or chatbot-based systems, given the success of prior programming tools that use these interfaces [49]. Finally, given the additional information that is captured via EBUG, we believe future work could examine potential benefits to downstream bug reporting tasks such as bug triaging or localization.

7 LIMITATIONS

EBUG computes a GUI model of the relevant app in its models generation phase. GUI model generation is based on static and dynamic analysis. In the evaluation, we used a timeout of 60 minutes for the dynamic analysis. Although

the dynamic analysis might take longer if the analysis is run exhaustively (or a higher timeout is used), our evaluation indicates that even with a timeout of 60 minutes EBUG can achieve good results. Furthermore, the analysis needs to be run only once for every app release. For this reason, we believe that the cost associated with this part of EBUG should have a limited impact on EBUG. The analyses performed by EBUG might not fully capture all the transitions and screens in the relevant app. EBUG's evaluation results give us confidence that generated models are effective, but this aspect of the technique could be improved in future work by leveraging ongoing research on GUI model abstraction and representation (e.g., [50]).

Our current prototype tool, in its static analysis part of the GUI model generation phase, does not currently create different screens for the different fragments in an activity of a certain app. This characteristic can lead the prototype tool to create an overapproximated GUI model of the app. The GUI model is overapproximated as it might contain transitions that do not actually exist in the app due to overapproximations (e.g., context-insensitive analysis) made in the static analysis used by EBUG. Although this aspect might lead to false positives in the S2Rs validation process, we did not encounter significant challenges related to this aspect in the studies performed to evaluate EBUG.

The main assumption for EBUG to work effectively is that reported S2Rs describe GUI actions in the app. The description of a S2R needs to be provided through text but does not need to follow an imposed structure. EBUG is able to process S2R descriptions having different structures by leveraging the dependency trees associated with the descriptions [16], [22]. To take full advantage of EBUG, the text of a S2R should describe the components of the corresponding GUI action, that is, the type of action, the GUI element affected by the action, and the properties of the action (when applicable). EBUG is also able to operate even if one or more S2R descriptions do not provide the desired information. Although the technique will not be able to map those GUI descriptions into GUI actions, it can map subsequent descriptions leveraging the GUI model and feedback from the user (i.e., the user can correct the provided mapping). This part of EBUG also allows the technique to support incomplete or missing S2R descriptions.

Our approach does not currently offer support for actions based on multi-touch gestures (e.g., pinch in). If an app uses that type of action, EBUG would not be able to validate or provide suggestions for that type of action but would still be able to support other actions in the app. We are currently investigating ways to support that type of action, and, based on our experience, that type is not the predominant one in a large number of apps.

Our approach binds app interactions to GUI elements. Certain Android apps, however, rely on bitmapped (rather than GUI) elements. Hence, the approach cannot currently handle such apps. Luckily, the vast majority of these apps are games, whereas other types of apps tend to rely on standard GUI elements. We also believe that those apps could be supported by identifying interactable elements using approaches based on computer vision techniques (e.g., edge detection).

Finally, our predictive models require training data.

Although it might not be possible to collect such data from users in some cases, we believe that developers can mitigate this problem by performing this activity in-house as described also in Sec. 6.

8 RELATED WORK

8.1 Automation for Bug Reporting Systems

FUSION is a bug reporting system developed by Moran et al. [4] that has the goal of facilitating the reporting process by providing pre-populated structured fields that reporters can use to compose the reproduction steps of their reports. While this approach led to the construction of bug reports that were more reproducible than a baseline issue tracking system, it came at the cost of longer reporting times. Our work in developing the EBUG approach differs in two important aspects. First, we are the first to explore how predictive modeling can enhance the bug reporting process by proactively providing suggestions to users about yetto-be-completed information. Second, our interface blends a predictive free-form text field in conjunction with visual information (e.g., screenshots), whereas FUSION used more structured menus. We compared EBUG against FUSION in our comprehensive evaluation and found that reporters could create bug reports faster with EBUG and these reports were more reproducible in far less time, illustrating the advantages of EBUG.

Other approaches from related work focus on the analysis of bug reports after they are created. Specifically these techniques analyze completed bug reports and (i) detect missing information [51], (ii) assess the quality of the reproduction steps [42], and (iii) automatically reproduce bug reports [11], [23], [52]. While such research is related to ours, we view our proposed approach as complimentary, as it aims to increase bug report quality at the time of reporting which in turn, can help to improve the performance of the techniques in this area of research by providing higher quality bug information to analyze. Additionally, automated reproduction techniques of bug reports [11], [23] and app reviews [52] could be integrated with bug reporting systems like EBUG so that the bug reporting process would have an iterative feedback loop in which the results of automated bug reproduction are provided to the users as they are reporting bugs to help them refine the reports.

Finally, a wealth of work has been conducted on automating downstream SE tasks that are tied to bug reporting, such as (i) triaging, (ii) fault localization, and (iii) duplicate detectio. We now provide a summary of the work in this area of research but refer the reader to surveys on this line of research for a more thorough discussion [18], [53], [54]. Related work on triaging improved assignment of bug reports to developers [55], [56], [57], [58], reduced bug report tossing [59], and prioritized bug reports [60]. Work on bug reports and fault localization used information retrieval to identify files that likely need fixing [61], leveraged multi-step recommendation models based on information retrieval [62], used version control history information to identify relevant buggy files [63], utilized stack traces to identify files leading to crashing faults [64], [65]. Related work on duplicate bug report detection used information retrieval [66], leveraged natural language processing and

execution information [67], and adopted a learning to rank approach [68]. We view EBUG as complimentary to the techniques in this line of research, as its main goal is to improve the quality of bug reports *before* they are submitted, which can ultimately benefit all of the downstream approaches by providing reports of higher quality.

8.2 Bug Reporting Studies

A large body of work focused on understanding the various aspects of the bug reporting processes, and we draw inspiration from several of the findings of this research area. Bettenburg et al. [3] performed a study with developers of open source software systems and found that (i) steps to reproduce, (ii) stack traces, and (iii) test cases were the most important pieces of information to be included within a bug report. However, these pieces of information were also the most difficult for reporters to provide. This helps to illustrate the importance of EBUG's real-time understanding and auto-completion of S2Rs as these features can help reporters in quickly writing effective S2Rs. Bettenburg et al. [69] also examined developer's perceptions of duplicate bug reports, and found that many developers did not necessarily view them as harmful, but rather, as providing important information. While EBUG does not focus on duplicate reports, we hope that the accurate reporting mechanism provided by our approach could aid future automated analysis tools that aim to find related or duplicate bug reports for developers. The same authors also proposed a tool that is capable of extracting structural information from bug reports [70]. Similarly, Song & Chaparro studied and proposed techniques for automatically identifying different types of lexical information from bug reports [51], [71]. EBUG is complementary to this line of research as the approach introduces a new form of real-time analysis of S2Rs that can identify various components in the S2Rs.

8.3 Field Failure Reproduction

The research area of field failure reproduction shares similar goals with our work, and with bug reporting systems in general, in that they attempt to capture and relay fault information to developers. Related work on field failure reproduction proposed techniques to reproduce failures using genetic programming [72], [73], crowdsourcing [74], symbolic execution [75], [76], [77], [78], static analysis combined with symbolic execution [79], model checking [80], dynamic exploration [81], sequential pattern minin on user traces [82], and test case mutaiton [83].

However, we view EBUG, and other user-facing bug reporting systems as largely complimentary to these field failure reproduction approaches. This is mainly due to the fact that in-field techniques require a known oracle, such as a crash, to identify that a failure has occurred. Conversely, user-facing bug reporting systems can be used for failures discovered by end-users or other developers. These bugs have been shown to account for more than half of the reported bugs in open source systems [2].

8.4 User Support Systems and Interactive Debugging

Previous work in the SE research community has focused on creating tools that emulate a chatbot-like interface for debugging, question answering, and user support. For instance, the work by Ko et al. focused on helping programmers ask questions about program behavior [49]. Parmit et al. introduced the LemonAid system that provided contextual assistance on websites through crowdsourcing techniques. There are also a number of commercial tools such as BugClipper [84] or BugSee [85] provide features for screen recording, and bug report annotation. In contrast to these tools, EBUG offers proactive suggestions to help speed up and improve the quality of information included in the bug reporting process. While EBUG does not make use of a chatbot interface, we believe this is an interesting direction of future work to explore for the domain of bug reporting.

8.5 Automated GUI Exploration for Mobile Apps

The modeling techniques employed by EBUG share certain similarities with automated GUI exploration techniques for mobile apps. Automated input generation approaches for mobile apps perform *random-based* input generation [86], [87], [88], [89], do *systematic* input generation [90], [91], [92], [93], [94], employ *model-based* input generation [10], [36], [92], [95], [96], [97], [98], use a *search-based* approach [99], [100], and perform *symbolic* input generation [101], [102]. Related work, also investigated how to integrate user feedback to improve mobile app testing [103], [104].

The two most highly related techniques to those proposed in EBUG are the MonkeyLab approach proposed by Linares-Vasquez et al. [36], and the Polariz approach proposed by Mao et al. [105]. Both of these approaches attempt to incorporate app usages gathered from crowd workers in order to generate a more diverse set of actions to guide automated input generation. EBUG's input generation approach and prediction modeling function in a similar manner, but with a markedly different end goal: aid users in predicting reproduction steps during the bug reporting process. As such, EBUG's approach uses an n-gram model to predict likely reproduction steps and target GUI elements for reporters, rather than using this information to guide an input generation technique.

9 Conclusion

When users experience a software failure, they have the option of submitting a bug report to provide information about the failure. Developers rely on submitted bug reports to identify and remedy the faults in their software. Unfortunately, the quality of manually constructed bug reports can vary widely due to the effort required to include essential information, such as a detailed description of the S2Rs. This issue affects and complicates the developers' task of reproducing the failures experienced by the users. To help users in writing bug reports that are easier to reproduce, we propose the EBUG bug reporting approach. EBUG understands S2Rs written in natural language and, as the user reports them, uses a novel predictive model to suggest additional S2Rs for helping the user complete the bug report. We implemented EBUG to support bug reporting in the context of Android apps and empirically evaluated it in two user studies. The studies involved ten participants who submitted ten bug reports each, and ten developers who reproduced the submitted bug reports. Our results show that reporters were able to write bug reports 31% faster with EBUG as compared to a state-of-the-art bug reporting system used as a baseline.

The results also show that developers could reproduce more failures using EBUG's reports than using those generated using the baseline.

We foresee a number of venues for future work. First, we will investigate ways to help users in defining observed and expected behavior. Second, we will study how to automatically extract more domain knowledge from the relevant app so that EBUG can understand and provide suggestions for macro-S2Rs (i.e., S2Rs that should be interpreted as sequences of multiple actions on the GUI). Third, we will evaluate the application of the approach in the context of web and GUI-based desktop software. Fourth, we also think that EBUG could be extended by evaluating it in additional studies. Specifically, EBUG could be evaluated in studies where developers use EBUG to report bugs and studies where EBUG is used by the developers of the relevant apps upon receiving newly identified bugs. Fifth, future studies could also analyze the tradeoffs between closed and open vocabularies in EBUG's prediction models. Finally, and more on the engineering side, we will extend EBUG to model fragments as independent screens in the relevant app and add support for multi-touch gestures, which will allow us to handle reports involving these kinds of interactions.

ACKNOWLEDGMENT

W&M and GMU co-authors have been supported in part by the NSF CCF-1955853 and CCF-2007246 grants. Any opinions, findings, and conclusions expressed herein are the authors' and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] G. Tassey, "The economic impacts of inadequate infrastructure for software testing," National Institute of Standards and Technology, Tech. Rep., 2002.
- [2] L. Tan, C. Liu, Z. Li, X. Wang, Y. Zhou, and C. Zhai, "Bug characteristics in open source software," *Empirical Software Engineering*, vol. 19, no. 6, pp. 1665–1705, 2014.
- [3] N. Bettenburg, S. Just, A. Schröter, C. Weiss, R. Premraj, and T. Zimmermann, "What makes a good bug report?" in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. New York, NY, USA: ACM, 2008.
- [4] K. Moran, M. Linares-Vásquez, C. Bernal-Cárdenas, and D. Poshyvanyk, "Auto-completing bug reports for android applications," in *Proceedings of the 2015 10th Joint Meeting on Foun*dations of Software Engineering. New York, NY, USA: ACM, 2015.
- [5] (2021) Google smart compose. [Online]. Available: https://www.blog.google/products/gmail/subject-write-emails-faster-smart-compose-gmail
- [6] M. Fazzini, K. Moran, C. Bernal-Cardenas, T. Wendland, A. Orso, and D. Poshyvanyk. (2021) EBUG's online appendix. [Online]. Available: https://www-users.cs.umn.edu/ ~mfazzini/ebug.html
- [7] (2021) Deposit/withdrawal change existing entry. [Online].Available: https://github.com/codinguser/gnucash-android/issues/247
- [8] (2021) Gnucash github. [Online]. Available: https://github.com/ codinguser/gnucash-android
- [9] (2021) Gnucash. [Online]. Available: https://play.google.com/ store/apps/details?id=org.gnucash.android
- [10] W. Choi, G. Necula, and K. Sen, "Guided gui testing of android apps with minimal restart and approximate learning," in Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications. New York, NY, USA: ACM, 2013, pp. 623–640.
- [11] M. Fazzini, M. Prammer, M. d'Amorim, and A. Orso, "Automatically translating bug reports into test cases for mobile apps," in Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis. New York, NY, USA: ACM, 2018, pp. 141–152.

- [12] S. Yang, H. Zhang, H. Wu, Y. Wang, D. Yan, and A. Rountev, "Static window transition graphs for android (t)," in *Proceedings of the 2015 30th IEEE/ACM International Conference on Automated Software Engineering*. Piscataway, NJ, USA: IEEE Press, 2015.
- [13] T. Wanwarang, N. P. Borges, L. Bettscheider, and A. Zeller, "Testing apps with real-world inputs," in *Proceedings of the IEEE/ACM 1st International Conference on Automation of Software Test*, ser. AST '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–10.
- [14] (2021) Tesseract ocr. [Online]. Available: https://github.com/ tesseract-ocr/tesseract
- [15] F. Harary, Graph Theory. USA: Avalon Publishing, 1969.
- [16] D. Jurafsky and J. H. Martin, Speech and Language Processing (2nd Edition). USA: Prentice-Hall, Inc., 2009.
- [17] S. F. Chen and J. Goodman, "An empirical study of smoothing techniques for language modeling," Computer Science Group, Harvard University, Tech. Rep. TR-10-98, 1998.
- [18] W. E. Wong, R. Gao, Y. Li, R. Abreu, and F. Wotawa, "A survey on software fault localization," *IEEE Transactions on Software Engineering*, vol. 42, no. 8, pp. 707–740, 2016.
- [19] G. Miner, J. Elder, T. Hill, R. Nisbet, D. Delen, and A. Fast, Practical Text Mining and Statistical Analysis for Non-structured Text Data Applications. Orlando, FL, USA: Academic Press, 2012.
- [20] J. Saito. (2016, Jul.) Making a case for letter case. [Online]. Available: https://medium.com/@jsaito/ making-a-case-for-letter-case-19d09f653c98
- [21] G. Angeli, M. J. J. Premkumar, and C. D. Manning, "Leveraging linguistic structure for open domain information extraction," in Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing. Stroudsburg, PA, USA: The Association for Computer Linguistics, 2015, pp. 344–354.
- [22] M. de Marneffe, B. MacCartney, and C. D. Manning, "Generating typed dependency parses from phrase structure parses," in *Proceedings of the Fifth International Conference on Language Resources and Evaluation*. Bern, Switzerland: European Language Resources Association (ELRA), 2006, pp. 449–454.
- [23] Y. Zhao, T. Yu, T. Su, Y. Liu, W. Zheng, J. Zhang, and W. G. J. Halfond, "Recdroid: Automatically reproducing android application crashes from bug reports," in *Proceedings of the 41st International Conference on Software Engineering*. Piscataway, NJ, USA: IEEE Press, 2019, p. 128–139.
- [24] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," in Soviet physics doklady, 1966.
- [25] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Transactions of the Association for Computational Linguistics*, pp. 135–146, 2017.
- [26] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," 2013.
- [27] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proceedings of the 26th International Confer*ence on Neural Information Processing Systems. Red Hook, NY, USA: Curran Associates Inc., 2013, p. 3111–3119.
- [28] M.-T. Luong, R. Socher, and C. D. Manning, "Better word representations with recursive neural networks for morphology," in Proceedings of the Seventeenth Conference on Computational Natural Language Learning, 2013, pp. 104–113.
- [29] S. Qiu, Q. Cui, J. Bian, B. Gao, and T.-Y. Liu, "Co-learning of word representations and morpheme representations," in *Proceedings* of the 25th International Conference on Computational Linguistics: Technical Papers, 2014, pp. 141–150.
- [30] R. Soricut and F. J. Och, "Unsupervised morphology induction using word embeddings," in Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2015, pp. 1627–1637.
- [31] J. Ostrander, Android UI Fundamentals: Develop and Design. Berkeley, CA, USA: Peachpit Press, 2012.
- [32] (2021) UI automator. [Online]. Available: https://developer.android.com/training/testing/ui-automator
- [33] (2021) Neo4j. [Online]. Available: https://neo4j.com
- [34] (2021) Getevent. [Online]. Available: https://source.android. com/devices/input/getevent
- [35] B.-J. Hsu and J. Glass, "Iterative language model estimation: efficient data structure & algorithms," in Ninth Annual Conference of the International Speech Communication Association, 2008.

- [36] M. Linares-Vásquez, M. White, C. Bernal-Cárdenas, K. Moran, and D. Poshyvanyk, "Mining android app usages for generating actionable gui-based execution scenarios," in *Proceedings of the 12th Working Conference on Mining Software Repositories*. Piscataway, NJ, USA: IEEE Press, 2015, p. 111–122.
- 37] (2021) Google natural language. [Online]. Available: https://cloud.google.com/natural-language
- [38] (2021) Quill. [Online]. Available: https://quilljs.com
- [39] (2021) English word vectors. [Online]. Available: https://fasttext.cc/docs/en/english-vectors.html
- [40] J. Pitkow and P. Pirolli, "Mining longest repeating subsequences to predict world wide web surfing," in Proceedings of the 2nd Conference on USENIX Symposium on Internet Technologies and Systems. USA: USENIX Association, 1999, p. 13.
- [41] T. Gueniche, P. Fournier-Viger, R. Raman, and V. S. Tseng, "Cpt+: Decreasing the time/space complexity of the compact prediction tree," in *Proceedings of the 19th Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Berlin, Germany: Springer, 2015.
- [42] O. Chaparro, C. Bernal-Cárdenas, J. Lu, K. Moran, A. Marcus, M. Di Penta, D. Poshyvanyk, and V. Ng, "Assessing the quality of the steps to reproduce in bug reports," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. New York, NY, USA: ACM, 2019, p. 86–96.
- [43] (2021) Gnucash. [Online]. Available: https://play.google.com
- [44] J. Brooke, "SUS: A quick and dirty usability scale," in *Usability evaluation in industry*. London: Taylor and Francis, 1996.
- [45] P. Morville. (2021) User experience design. [Online]. Available: http://semanticstudios.com/user_experience_design
- [46] I. Salman, A. T. Misirli, and N. Juristo, "Are students representatives of professionals in software engineering experiments?" in Proceedings of the 37th International Conference on Software Engineering Volume 1, ser. ICSE '15. IEEE Press, 2015, p. 666–676.
- [47] K. Mao, M. Harman, and Y. Jia, "Crowd intelligence enhances automated mobile testing," in *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE 2017. IEEE Press, 2017, p. 16–26.
- [48] X. Chen, Y. Wang, J. He, S. Pan, Y. Li, and P. Zhang, "Cap: Context-aware app usage prediction with heterogeneous graph embedding," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 3, no. 1, mar 2019. [Online]. Available: https://doi.org/10.1145/3314391
- [49] A. J. Ko and B. A. Myers, "Debugging reinvented: Asking and answering why and why not questions about program behavior," in *Proceedings of the 30th International Conference on Software Engineering (ICSE'08)*, 2008, p. 301–310.
- [50] D. Lai and J. Rubin, "Goal-driven exploration for android applications," in 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE), 2019, pp. 115–127.
- [51] O. Chaparro, J. Lu, F. Zampetti, L. Moreno, M. Di Penta, A. Marcus, G. Bavota, and V. Ng, "Detecting missing information in bug descriptions," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. New York, NY, USA: ACM, 2017, p. 396–407.
- [52] S. Li, J. Guo, M. Fan, J.-G. Lou, Q. Zheng, and T. Liu, "Automated bug reproduction from user reviews for android applications," in 2020 IEEE/ACM 42nd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP). USA: IEEE, 2020, pp. 51–60.
- [53] J. Zhang, X. Wang, D. Hao, B. Xie, L. Zhang, and H. Mei, "A survey on bug-report analysis," *Science China Information Sciences*, vol. 58, no. 2, pp. 1–24, 2015.
- [54] J. Uddin, R. Ghazali, M. M. Deris, R. Naseem, and H. Shah, "A survey on bug prioritization," Artificial Intelligence Review, vol. 47, no. 2, pp. 145–180, 2017.
- [55] R. Shokripour, J. Anvik, Z. M. Kasirun, and S. Zamani, "Why so complicated? simple term filtering and weighting for locationbased bug report assignment recommendation," in *Proceedings* of the 10th Working Conference on Mining Software Repositories. Piscataway, NJ, USA: IEEE Press, 2013, pp. 2–11.
- [56] H. Naguib, N. Narayan, B. Brügge, and D. Helal, "Bug report assignee recommendation using activity profiles," in *Proceedings* of the 10th Working Conference on Mining Software Repositories. Piscataway, NJ, USA: IEEE Press, 2013, pp. 22–30.
- [57] J. woo Park, M.-W. Lee, J. Kim, S. won Hwang, and S. Kim, "Costriage: A cost-aware triage algorithm for bug reporting systems," 2011

- [58] M. Linares-Vasquez, K. Hossen, H. Dang, H. Kagdi, M. Gethers, and D. Poshyvanyk, "Triaging incoming change requests: Bug or commit history, or code authorship?" in Proceedings of the 28th IEEE International Conference on Software Maintenance. Piscataway, NJ, USA: IEEE Press, 2012, pp. 451-460.
- G. Jeong, S. Kim, and T. Zimmermann, "Improving bug triage with bug tossing graphs," in Proceedings of the the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering. New York, NY, USA: ACM, 2009, pp. 111-120.
- [60] H. Shen, J. Fang, and J. Zhao, "Efindbugs: Effective error ranking for findbugs," in Proceedings of the 2011 IEEE Fourth International Conference on Software Testing, Verification and Validation. Piscataway, NJ, USA: IEEE Press, 2011, pp. 299-308.
- J. Zhou, H. Zhang, and D. Lo, "Where should the bugs be fixed? more accurate information retrieval-based bug localization based on bug reports," in Proceedings of the 34th International Conference on Software Engineering. Piscataway, NJ, USA: IEEE Press, 2012.
- [62] D. Kim, Y. Tao, S. Kim, and A. Zeller, "Where should we fix this bug? a two-phase recommendation model," IEEE Transactions on Software Engineering, vol. 39, no. 11, pp. 1597–1610, 2013.
- S. Wang and D. Lo, "Version history, similar report, and structure: Putting them together for improved bug localization," in Proceedings of the 22Nd International Conference on Program Comprehension. New York, NY, USA: ACM, 2014, pp. 53-63.
- R. Wu, H. Zhang, S.-C. Cheung, and S. Kim, "Crashlocator: Locating crashing faults based on crash stacks," in Proceedings of the 2014 International Symposium on Software Testing and Analysis. New York, NY, USA: ACM, 2014, pp. 204-214.
- L. Moreno, J. J. Treadway, A. Marcus, and W. Shen, "On the use of stack traces to improve text retrieval-based bug localization,' in 2014 IEEE International Conference on Software Maintenance and Evolution. USA: IEEE, 2014, pp. 151–160.
- A. T. Nguyen, T. T. Nguyen, T. N. Nguyen, D. Lo, and C. Sun, "Duplicate bug report detection with a combination of information retrieval and topic modeling," in Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering. New York, NY, USA: ACM, pp. 70-79.
- X. Wang, L. Zhang, T. Xie, J. Anvik, and J. Sun, "An approach to detecting duplicate bug reports using natural language and execution information," in Proceedings of the 30th International Conference on Software Engineering. New York, NY, USA: ACM, 2008, pp. 461-470.
- [68] J. Zhou and H. Zhang, "Learning to rank duplicate bug reports," in Proceedings of the 21st ACM International Conference on Information and Knowledge Management. New York, NY, USA: ACM, 2012, pp. 852-861.
- N. Bettenburg, R. Premraj, T. Zimmermann, and S. Kim, "Duplicate bug reports considered harmful... really?" in Proceedings of the 2008 IEEE International Conference on Software Maintenance. Piscataway, NJ, USA: IEEE Press, 2008, pp. 337-345.
- -, "Extracting structural information from bug reports," in Proceedings of the 2008 International Working Conference on Mining Software Repositories. New York, NY, USA: ACM, 2008, pp. 27–30.
- Y. Song and O. Chaparro, BEE: A Tool for Structuring and Analyzing Bug Reports. New York, NY, USA: ACM, 2020, p. 1551–1555.
- M. Soltani, A. Panichella, and A. van Deursen, "A guided genetic algorithm for automated crash reproduction," in Proceedings of the 39th International Conference on Software Engineering. Piscataway, NJ, USA: IEEE Press, 2017, pp. 209-220.
- F. M. Kifetew, W. Jin, R. Tiella, A. Orso, and P. Tonella, "Reproducing field failures for programs with complex grammar-based " in Proceedings of the 2014 IEEE Seventh International Conference on Software Testing, Verification and Validation. Piscataway, NJ, USA: IEEE Press, 2014, pp. 163–172.
- M. Gómez, R. Rouvoy, B. Adams, and L. Seinturier, "Reproducing context-sensitive crashes of mobile apps using crowdsourced monitoring," in Proceedings of the International Conference on Mobile Software Engineering and Systems. New York, NY, USA: ACM, 2016, pp. 88-99.
- N. Chen and S. Kim, "Star: Stack trace based automatic crash reproduction via symbolic execution," IEEE Transactions on Software Engineering, vol. 41, no. 2, pp. 198-220, 2015.
- W. Jin and A. Orso, "Automated support for reproducing and debugging field failures," ACM Trans. Softw. Eng. Methodol., vol. 24, no. 4, 2015.

- ----, "Bugredux: Reproducing field failures for in-house debugging," in Proceedings of the 34th International Conference on Software Engineering. Piscataway, NJ, USA: IEEE Press, 2012, pp. 474–484.
- Y. Cao, H. Zhang, and S. Ding, "Symcrash: Selective recording for reproducing crashes," in Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering. New York, NY, USA: ACM, 2014, pp. 791-802.
- C. Zamfir and G. Candea, "Execution synthesis: A technique for automated software debugging," in *Proceedings of the 5th European Conference on Computer Systems*. New York, NY, USA: ACM, 2010, pp. 321-334.
- M. Nayrolles, A. Hamou-Lhadj, S. Tahar, and A. Larsson, "A bug reproduction approach based on directed model checking and crash traces," Journal of Software: Evolution and Process, vol. 29, no. 3, 2017.
- M. White, M. Linares-Vásquez, P. Johnson, C. Bernal-Cárdenas, and D. Poshyvanyk, "Generating reproducible and replayable bug reports from android application crashes," in Proceedings of the 2015 IEEE 23rd International Conference on Program Comprehension. Piscataway, NJ, USA: IEEE Press, 2015, p. 48-59.
- T. Roehm, S. Nosovic, and B. Bruegge, "Automated extraction of failure reproduction steps from user interaction traces," in Proceedings of the 2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering. Piscataway, NJ, USA: IEEE Press, 2015, pp. 121-130.
- J. Xuan, X. Xie, and M. Monperrus, "Crash reproduction via test case mutation: Let existing test cases help," in Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering. New York, NY, USA: ACM, 2015, pp. 910-913.
- "Bugclipper http://bugclipper.com," 2021.
- "https://www.bugsee.com," 2021.
- [86] (2021) Android ui/application exerciser monkey. [Online]. Available: http://developer.android.com/tools/help/monkey.html
- A. Machiry, R. Tahiliani, and M. Naik, "Dynodroid: An input generation system for android apps," in Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering. York, NY, USA: ACM, 2013, pp. 224-234.
- R. Sasnauskas and J. Regehr, "Intent fuzzer: Crafting intents of death," in Proceedings of the 2014 Joint International Workshop on Dynamic Analysis and Software and System Performance Testing, Debugging, and Analytics. New York, NY, USA: ACM, 2014.
- L. Ravindranath, S. Nath, J. Padhye, and H. Balakrishnan, "Automatic and scalable fault detection for mobile applications, Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services. New York, NY, USA: ACM, 2014, pp. 190-203.
- [90] D. Amalfitano, A. R. Fasolino, P. Tramontana, S. De Carmine, and A. M. Memon, "Using gui ripping for automated testing of android applications," in Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering. New York, NY, USA: ACM, 2012, pp. 258-261.
- S. Anand, M. Naik, M. J. Harrold, and H. Yang, "Automated concolic testing of smartphone apps," in Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering. New York, NY, USA: ACM, 2012, pp. 59-69.
- T. Azim and I. Neamtiu, "Targeted and depth-first exploration for systematic testing of android apps," in Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications. New York, NY, USA: ACM, 2013, pp. 641-660.
- K. Moran, M. Linares-Vásquez, C. Bernal-Cárdenas, C. Vendome, and D. Poshyvanyk, "Automatically discovering, reporting and reproducing android application crashes," in Proceedings of the IEEE International Conference on Software Testing, Verification and Validation. Piscataway, NJ, USA: IEEE Press, 2016, pp. 33-44.
- (2021) Google firebase test lab robo test. [Online]. Available: https://firebase.google.com/docs/test-lab/robo-ux-test
- D. Amalfitano, A. R. Fasolino, P. Tramontana, B. D. Ta, and A. M. Memon, "Mobiguitar: Automated model-based testing of mobile apps," IEEE Software, vol. 32, no. 5, pp. 53-59, 2015.
- R. N. Zaeem, M. R. Prasad, and S. Khurshid, "Automated generation of oracles for testing user-interaction features of mobile apps," in Proceedings of the 2014 IEEE International Conference on Software Testing, Verification, and Validation. Washington, DC, USA: IEEE Computer Society, 2014, pp. 183–192. W. Yang, M. R. Prasad, and T. Xie, "A grey-box approach
- for automated gui-model generation of mobile applications,"

in *Proceedings of the 16th International Conference on Fundamental Approaches to Software Engineering.* Berlin, Heidelberg: Springer-Verlag, 2013, pp. 250–265.

[98] H. Zhang and A. Rountev, "Analysis and testing of notifications in android wear applications," in *Proceedings of the 2017 IEEE/ACM 39th International Conference on Software Engineering*. Piscataway, NJ, USA: IEEE Press, 2017, pp. 347–357.

Piscataway, NJ, USA: IEEE Press, 2017, pp. 347–357.

[99] R. Mahmood, N. Mirzaei, and S. Malek, "Evodroid: Segmented evolutionary testing of android apps," in *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. New York, NY, USA: ACM, 2014, pp. 599–609.

[100] K. Mao, M. Harman, and Y. Jia, "Sapienz: Multi-objective automated testing for android applications," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*. New York, NY, USA: ACM, 2016, pp. 94–105.

[101] C. S. Jensen, M. R. Prasad, and A. Møller, "Automated testing with targeted event sequence generation," in *Proceedings of the* 2013 International Symposium on Software Testing and Analysis. New York, NY, USA: ACM, 2013, p. 67–77.

[102] N. Mirzaei, H. Bagheri, R. Mahmood, and S. Malek, "Sig-droid: Automated system input generation for android applications," in *Proceedings of the 2015 IEEE 26th International Symposium on Software Reliability Engineering*. USA: IEEE Computer Society, 2015, p. 461–471.

[103] G. Grano, A. Ciurumelea, S. Panichella, F. Palomba, and H. C. Gall, "Exploring the integration of user feedback in automated testing of android applications," in 2018 IEEE 25Th international conference on software analysis, evolution and reengineering (SANER). USA: IEEE, 2018, pp. 72–83.

[104] L. Pelloni, G. Grano, A. Ciurumelea, S. Panichella, F. Palomba, and H. C. Gall, "Becloma: Augmenting stack traces with user review information," in 2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER). USA: IEEE, 2018, pp. 522–526.

[105] K. Mao, M. Harman, and Y. Jia, "Crowd intelligence enhances automated mobile testing," in *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*. Piscataway, NJ, USA: IEEE Press, 2017, p. 16–26.



Mattia Fazzini is an Assistant Professor in the Department of Computer Science & Engineering at the University of Minnesota. He completed his Ph.D. in Computer Science from the Georgia Institute of Technology in 2019. His research interests lie primarily in the area of software engineering, with emphasis on techniques for improving software quality. The central theme of his research is the development of approaches for testing and maintenance of mobile apps. He has published in several top peer-reviewed

software engineering venues including: ASE, ICSE, ICST, and ISSTA. He received the Facebook Testing and Verification award in 2019 for his research on mobile app testing. More information is available at https://www-users.cs.umn.edu/~mfazzini/.



Kevin Moran is currently an Assistant Professor at George Mason University where he directs the SAGE research group. He received a Ph.D. degree from William & Mary in August 2018. His main research interest involves facilitating the processes of software engineering, maintenance, and evolution with a focus on mobile platforms. He has published in several top peer-reviewed software engineering venues including: ICSE, ESEC/FSE, TSE, USENIX, ICST, ICSME, and MSR. His research has been recognized

with ACM SIGSOFT distinguished paper awards at ESEC/FSE 2019 and ICSE 2020. He is a member of the ACM and IEEE. More information is available at http://www.kpmoran.com



Carlos Bernal-Cárdenas is currently Ph.D. candidate in Computer Science at the College of William & Mary as a member of the SEMERU research group advised by Dr Denys Poshyvanyk. His research interests include software engineering, software evolution and maintenance, information retrieval, software reuse, mining software repositories, mobile applications development, and user experience. He has published in several top peer-reviewed software engineering venues including: ICSE, ESEC/FSE, ICST, and

MSR. He has also received the ACM SigSoft Distinguished paper award at ESEC/FSE'15 & '19 and ICSE'20. More information is available at http://www.cs.wm.edu/~cebernal/.



Tyler Wendland is a graduate student at the University of Minnesota. He graduated with a B.S. in Computer Science from the University of Minnesota in 2020. His research focuses on devising automated software engineering techniques to improve bug reporting processes.



Alessandro Orso is a Professor in the School of Computer Science and an Associate Dean in the College of Computing at the Georgia Institute of Technology. He received his M.S. degree in Electrical Engineering (1995) and his Ph.D. in Computer Science (1999) from Politecnico di Milano, Italy. From March 2000, he has been at Georgia Tech. His area of research is software engineering, with emphasis on software testing and program analysis. His interests include the development of techniques and tools for improv-

ing software reliability, security, and trustworthiness, and the validation of such techniques on real-world systems. Dr. Orso has received funding for his research from both government agencies, such as DARPA, ONR, and NSF, and industry, such as Facebook, Fujitsu Labs, Google, IBM, and Microsoft. He served on the editorial boards of ACM TOSEM and on the Advisory Board of Reflective Corp, and he is serving on the editorial boards of IEEE TSE. Ha also served as program chair for ACM SIGSOFT ISSTA 2010, program co-chair for IEEE ICST 2013, ACM-SIGSOFT FSE 2014, and ACM-SIGSOFT/IEEE ICSE 2017, and as a technical consultant to DARPA. Dr. Orso has received multiple Best Paper and ACM Distinguished Paper Awards and three Impact and Most Influential Paper Awards. Dr. Orso is a Distinguished Member of the ACM and an IEEE Fellow. More information is available at http://www.cc.gatech.edu/~orso/.



Denys Poshyvanyk is a Professor of Computer Science at William and Mary. He received the MS and MA degrees in Computer Science from the National University of Kyiv-Mohyla Academy, Ukraine, and Wayne State University in 2003 and 2006, respectively. He received the PhD degree in Computer Science from Wayne State University in 2008. He served as a program co-chair for ASE'21, MobileSoft'19, ICSME'16, ICPC'13, WCRE'12 and WCRE'11. He currently serves on the editorial board of IEEE Transac-

tions on Software Engineering (TSE), Empirical Software Engineering Journal (EMSE, Springer), Journal of Software: Evolution and Process (JSEP, Wiley) and Science of Computer Programming. His research interests include software engineering, software maintenance and evolution, program comprehension, reverse engineering and software repository mining. His research papers received several Best Paper Awards at ICPC'06, ICPC'07, ICSM'10, SCAM'10, ICSM'13, CODAPSY'19 and ACM SIGSOFT Distinguished Paper Awards at ASE'13, ICSE'15, EC/FSE'15, ICPC'16, ASE'17, ESEC/FSE'19 and ICSE'20. He also received the Most Influential Paper Awards at ICSME'16, ICPC'17 and ICPC'20. He is a recipient of the NSF CAREER award (2013). He is a member of the IEEE and ACM. More information is available at: http://www.cs.wm.edu/~denys/