

App Review Driven Collaborative Bug Finding

Xunzhu Tang · Haoye Tian ·
Pingfan Kong · Kui Liu · Xin Xia ·
Jacques Klein · Tegawendé F. Bissyandé

Received: date / Accepted: date

Abstract Software development teams generally welcome any effort to expose bugs in their code base. In this work, we build on the hypothesis that mobile apps from the same category (e.g., two web browser apps) may be affected by similar bugs in their evolution process. It is therefore possible to transfer the experience of one historical app to quickly find bugs in its new counterparts. This has been referred to as collaborative bug finding in the literature. Our novelty is that we guide the bug finding process by considering that existing bugs have been hinted within app reviews. Concretely, we design the BUGRMSYS approach to recommend bug reports for a target app by matching historical bug reports from apps in the same category with user app reviews of the target app. We experimentally show that this approach enables us to quickly expose and report dozens of bugs for targeted apps such as **Brave** (web browser app). BUGRMSYS’s implementation relies on DistilBERT to produce natural language text embeddings. Our pipeline considers similarities between bug reports and app reviews to identify relevant bugs. We then focus on the app review as well as potential reproduction steps in the historical bug report (from a same-category app) to reproduce the bugs. During evaluation, we consider 300 bug reports that we compare to thousands of user reviews from **Brave**: BUGRMSYS recommends 23 bug reports from **Firefox** that may be relevant for **Brave**. After reproduction, we indeed confirm 4 bugs in **Brave**. The whole process takes less than 1 hour. In contrast, until BUGRMSYS was

X. Tang, H. Tian, P. Kong, J. Klein, T. F. Bissyandé
SnT, University of Luxembourg
E-mail: {xunzhu.tang, haoye.tian, tagewende.bissyande, jacques.klein}@uni.lu

K. Liu
Hua Wei
E-mail: brucekuiliu@gmail.com

X. Xia
Hua Wei
E-mail: xin.xia@acm.org

applied, some buggy behavior mentioned in a user review of **Brave** was left unattended for nearly 3 months. At the time of submission, that bug was already fixed by the developers thanks to BUGRMSYS. Overall, after applying BUGRMSYS to six popular apps, we were able to identify, reproduce and report 20 new bugs: among these, 9 reports have been already triaged, 6 were confirmed, and 4 have been fixed by official development teams, respectively.

Keywords Bug finding · App review · Bug similarity

1 Introduction

Modern apps must evolve quickly to adapt to a fierce competition in app markets where users have varied choices among feature-rich apps [32]. Unfortunately, the fast iteration in app updates often results in defects being found by users after releases [3]. Various research efforts based on static analysis [14, 19, 48] and dynamic testing [13, 50, 6, 45, 22] have therefore been carried out to detect bugs before releasing apps. Bug-free apps remain however a myth and even popular apps, which are intensively used by large user communities, often display simple but annoying defects [6, 1, 47]. Through app reviews, users can provide feedback on buggy behaviour that sometimes go overlooked by app developers for various reasons: reviews can be redundant and uninformative (e.g. simple praise or dispraise repeating the star rating) [26]. App reviews are also time-consuming to exploit and can mislead the identification of fault locations [43]. In contrast, official bug reports filed in the issue tracker are the focus of developer communities since these reports tend to be more readily exploitable for bug resolution.

It is noteworthy that if recurring bugs are not swiftly addressed by developers, they will lead to negative app reviews with significant impact on app score in app markets and other severe consequences such as app fails [21]. The aforementioned situation calls for a more careful consideration of user reviews by developers. In particular, it would be appealing to translate app reviews into bug reports that can be used by developers as starting points in their fight against bugs. However, there exists a significant gap between the language of user reviews and the language of developer bug reports. The former is generally formal and technically-written while the latter is informal and colloquially-written. In a recent work, Haering et al. [10] proposed a deep learning approach to match app reviews and bug reports with the ambition of easily tracking whether an issue reported in app review was already filed as an official bug report, which should increase bug fixing priority. While we subscribe to the claim that user feedback often lacks information that is relevant for developers (such as steps to reproduce or affected versions) [31, 61], their approach (1) does not address the key problem of review deluge, and (2) misses the opportunity to reveal new bugs to the developers. Indeed, on the one hand, for a popular app, there can be thousands of new reviews every day, most of which are noisy for developers since they do not offer insights into bugs. On the other hand, some app reviews may actually mention important

and annoying bugs which can impact user experience for a large number of users without ever being reported formally in the issue tracker.

In another research direction, Tan et al. [49] have proposed Bugine [49], a collaborative bug recommendation system that aims at pairing similar issue reports across different apps. Thanks to Bugine, they have empirically shown that it is indeed possible to match similar issue reports across different apps. However, Bugine can only report issues across apps where the relevant UI design is of high visual similarity. Besides limited to only UI-related bugs, Bugine does not take target app’s review into consideration, which enable it no ability to pick up useful bugs as input.

Following up on the hypothesis of the work by Tan et al., we performed a preliminary study (cf. Section 2) to check whether apps in the same category (e.g., two web browser apps or two calendar apps) tend to face similar development issues. Eventually, we observed that apps in the same category share similar issues since these apps (1) are built by leveraging the similar development frameworks for similar functionalities (e.g., *Unity* for gaming apps), (2) have similar UI design logic, and (3) use the same storage/notification/hosting services (e.g., Firebase) [23,24]. It seems therefore promising to build on the experience of historical apps to improve new apps. Additionally, prior work [20, 2] have demonstrated that interactions among developers of different software can be effective to improve the quality of each software.

This paper. We hypothesize that if app **A** and app **B** belong to the same category (we consider the categories listed in the Wikipedia enumeration of popular free and open source Android apps¹, e.g., web browsers, Games, etc.), bug reports from one can be relevant for discovering bugs in the other. Unfortunately, there can be too many bug reports filed in some app categories. For example, in the web browser category, the issue tracker of Firefox alone has received over 20 thousand bug reports. It is therefore necessary to identify those issues that are more likely to be relevant for the app under-study (i.e., the target app for bug discovery). To that end, our novel strategy in this work is to explore app reviews written by users for the target app. Our idea is that app reviews, which may contain hints about buggy behaviour observed by users of the target app, can be matched to bug reports from other apps in the same category.

We propose BUGRMSys, a collaborative bug finding approach that is guided by user app reviews. BUGRMSys finds bugs by recommending a bug report of app **A** (e.g., the excerpted bug report of Signal in Figure 1(a)) as being relevant to the target app **B** (i.e., Wire) given the similarity of the bug report with app reviews from users of **B** (e.g., the excerpted user review of Wire in Figure 1(b)). With the app review in **B** matched with a similar bug report in **A**, we reproduce the bug in **B** by leveraging reproduction steps in the bug report and additional information details from the app review. If reproduction is successful, we can confirm having found a “new bug” that will be filed into

¹ https://en.wikipedia.org/wiki/List_of_free_and_open-source_Android_applications

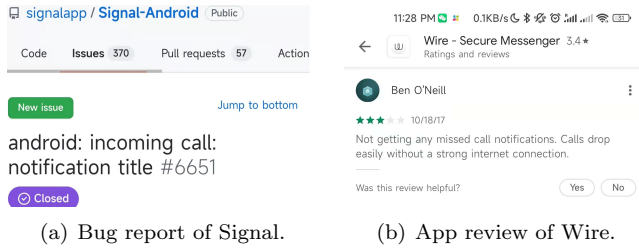


Fig. 1: Example of bug report and app review matched by BUGRMSys.

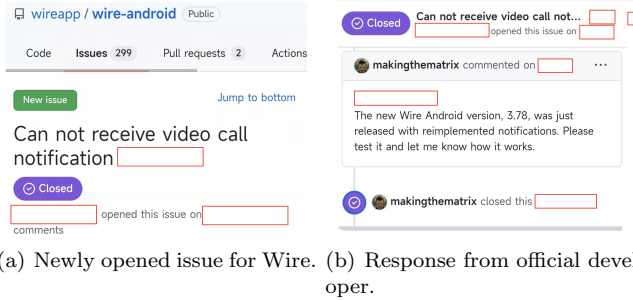


Fig. 2: Bug found by BUGRMSys and response from official developer (some texts have been hidden for anonymity).

the official issue tracker for app **B**. For example, the corresponding bug found by BUGRMSys in Wire was reported to its developers (shown in Figure 2(a)), and was finally got fixed by Wire’s official developer (Figure 2(b)) in one day. Surprisingly, the relevant user app review had been submitted since four years, but there is no any corresponding bug reported in the official issue tracker of the app yet. To the best of knowledge, this scheme of collaborative bug finding driven by app reviews, has not been previously explored in the literature.

The main contributions of our work are as follows:

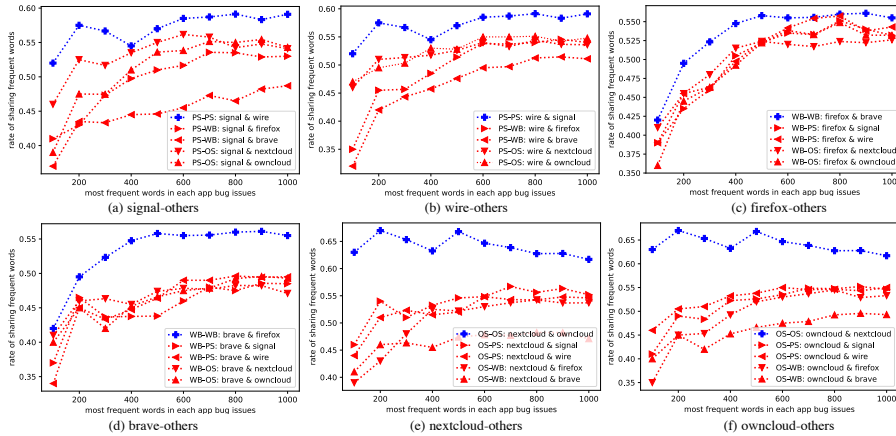
- We present insights from an empirical study about the similarity of issues reported across apps from the same category. These insights provide the intuitive basis for collaborative bug finding.
- We devise BUGRMSys, a bug recommendation system, which leverages similarities of user app reviews with bug reports to identify which bug reports from same category apps are good candidates to attempt a bug reproduction on a target app.
- We demonstrate experimentally that BUGRMSys is effective: applied to 6 apps from different categories, we find and reproduce 20 new bugs (i.e., bugs that are not yet reported in the issue trackers of these apps): 9, 6, and 4 reports have been already triaged, confirmed, fixed by official development teams, respectively.

Table 1: The Apps and their category.

| App Name | Repo Name | Category | # Downloads | #Bug Reports |
|-----------------|----------------|------------------|--------------|--------------|
| Signal | Signal-Android | privacy security | >50 million | 11,980 |
| Wire | wire-android | privacy security | >1 million | 3,677 |
| Firefox Android | fenix | Web Browser | >100 million | 24,087 |
| Brave | brave-browser | Web Browser | >50 million | 21,436 |
| NextCloud | nextcloud | Office Suite | >1 million | 9,890 |
| owncloud | owncloud | Office Suite | >0.1 million | 3,571 |

2 Preliminary Study

In this section, we conduct a preliminary study to evaluate the hypothesis of our work. Specifically, we seek to check that apps from the same category are more likely to share similar bug reports (and thus a bug report from app A could be relevant for app B if A and B are from the same category). We focus the comparison by estimating the overlap (i.e., the proportion of common words) between bug reports. To that end, we analyze the overlap rate of top_K frequent words of reports for apps from same or different categories. From a more qualitative perspective, we also analyze which types of words are frequently shared.

**Fig. 3:** Overlap rate of hot words of category apps' bug issues.

2.1 Empirical Setup

Apps: To conduct our preliminary study, we consider six popular apps listed in the Wikipedia page of free and open-source Android apps. Table 1 summarizes statistics about these apps. We consider *Signal* and *Wire* within the *Privacy-Security* category. Both apps have been downloaded more than one million times from the Google Play store. We also consider *Firefox* and *Brave*, two widely popular in the *Web Browser*, category. Finally, we consider *Nextcloud* and *Owncloud* among the apps in the *Office Suite* category.

Bug Reports: For each app, we have collected all bug reports that are present in their issue tracker system. The number of collected bug reports is

given in the last column of Table 1. Note that, for our experiment, we employ the Python NLTK [25] library and self-defined filters to pre-process the bug reports for natural language processing. We apply typical pre-processing tasks to remove stop words [54], punctuation, digits, etc.[9]. Meanwhile, to limit experimental bias, we set 10 different sizes K for the set of most frequent words, increasing step-wise until an order of magnitude: we consider Top_{100} , $\text{Top}_{200}, \dots, \text{Top}_{1000}$ frequent words. Concretely, to build each Top_K set for each app, we extract the K most frequent words in its bug reports. By analysing the Top_K frequent words, we can assess the differences in shared words between bug reports of same-category apps and different-category apps. Applied to all set ten Top_K sets, we can further check for potential trends, while empirically identifying the value of K under which the overlap (i.e., the proportion of shared words) is the highest.

Notation: In the rest of this paper, the category Privacy-Security, the category Web Browser, and the category Office-Suite are referred to as PS, WB, and OS, respectively. We also note the three pairwise combination of apps from the same category as follows: PS-PS: $\{ \text{Signal}, \text{Wire} \}$; WB-WB: $\{ \text{Firefox}, \text{Brave} \}$; OS-OS: $\{ \text{Nextcloud}, \text{Owncloud} \}$. Similarly, we consider 12 pairwise combinations of apps from different categories as follows: PS-WB: $\langle \text{Signal}, \text{Firefox} \rangle$, $\langle \text{Signal}, \text{Brave} \rangle$, $\langle \text{Wire}, \text{Firefox} \rangle$, and $\langle \text{Wire}, \text{Brave} \rangle$; PS-OS: $\langle \text{Signal}, \text{Nextcloud} \rangle$, $\langle \text{Signal}, \text{Owncloud} \rangle$, $\langle \text{Wire}, \text{Nextcloud} \rangle$, and $\langle \text{Wire}, \text{Owncloud} \rangle$; WB-OS: $\langle \text{Firefox}, \text{Nextcloud} \rangle$, $\langle \text{Firefox}, \text{Owncloud} \rangle$, $\langle \text{Brave}, \text{Nextcloud} \rangle$, and $\langle \text{Brave}, \text{Owncloud} \rangle$.

2.2 Hypothesis Analysis

In this section we check whether our hypothesis is valid from the quantitative and qualitative aspects of bug reports.

Quantitative Analysis: The results of the quantitative study are presented in Figure 3, where for each app X, we compute the percentage (or rate) of overlap, i.e., the percentage of shared words between the Top_K frequent words in the bug reports of X and the Top_K frequent words in the bug reports of another app Y. We note that, except for the $\{ \text{Owncloud}, \text{Nextcloud} \}$ pair, **both same-category and different-category pairs present that the percentage of shared words increases or reaches a plateau (often after $K=500$)**. In all six diagrams, and for all $K \in \{100, 200, \dots, 1000\}$, the percentage of shared Top_K frequent words is the highest when apps X and Y are from the same category (i.e., for Signal and Wire, Firefox and Brave, and, Owncloud and Nextcloud). It indicates that, **bug reports of same-category apps could present higher similarity than bug reports of different-category apps**.

Qualitative Analysis: We further investigate to what extent frequent words in bug reports of same-category and different-category apps share. To this end, we conduct the quantitative analysis by setting the threshold of frequency as 20 to select the most frequent words in bug reports (all words

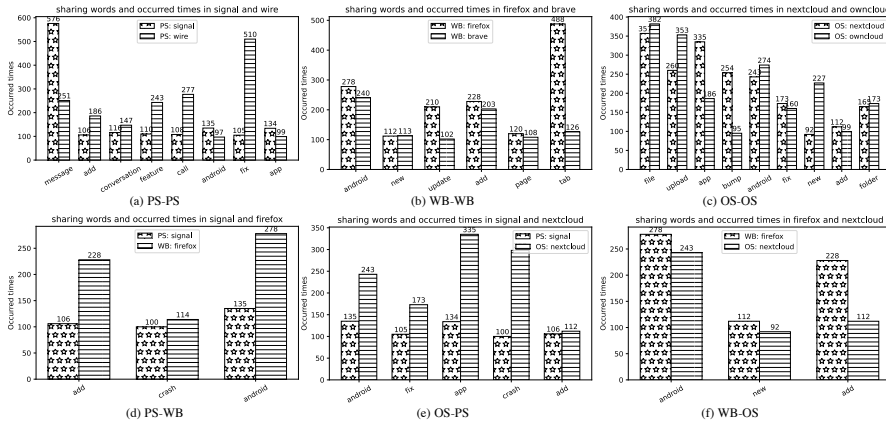


Fig. 4: Distribution of Top₂₀ shared frequent words in same-category app pairs and different-category app pairs.

are trimmed to keep their trunk, e.g., crashes→crash). With the top-20 most frequent words of each app, we assess to what extent different app pairs share these words, and how frequent the shared words occur in each app. Table 2 lists the top₂₀ frequent words of each app.

We have three pairs of same-category apps and twelve different-category apps. Considering the space limitation, we only show three pairs of different-category apps, and all data of twelve pairs of different-category apps are publicly available in our replication package². For the **PS** category, as shown in Figure 4(a), the top frequent words shared by Signal and Wire include message, conversation, call, etc. Apparently, these shared words represent the key features of these two apps (i.e., communication tool). They also share some general hot words about the Android Framework, such as android and app. For the **WB** category (cf. Figure 4(b)), the words shared by Firefox and Brave contain page and tab, of which both are highly relevant to the web browser. While for the category of **OS** (cf. Figure 4(c)), Nextcloud share “fix, file, folder, upload” words with Owncloud, which are related to the cloud storage system. Such results indicate that **the bug reports of same-category apps indeed share their field-specific information**. On the other hand, as shown in Figures 4(d), 4(e), and 4(f), the bug reports of different-category app pairs share less words than the same-category app pairs, and the shared words are mainly about the general information or problem of the Android framework itself (with words such as android, app, crash, add).

Regarding words that are not shared, by differentiating the top₂₀ frequent words and the shared words of the same/different-category pairs, we note that those non-shared frequent words of the same-category app pairs are about general features or problems (e.g., crash, notification, etc.), while those non-shared frequent words of the different-category app pairs are related to their main features of those apps. For example, given $\{Signal, Firefox\}_i$, a different-

² https://github.com/BugRMSys/BugRMSys/blob/main/12_different_category_apps

Table 2: Top₂₀ frequent words in bug reports for each app.

| App | | Top ₂₀ frequent words |
|-----|-----------|--|
| PS | Signal | signal, message, sms, group, android, app, contact, send, mms, conversation, notification, feature, call, add, fix, request, crash, textsecure, text |
| | Wire | fix, feature, add, conversation, avs, part, message, bump, wire, new, update, call, user, remove, app, android, video, version |
| WB | Firefox | bug, fnx, android, tab, add, tabs, update, search, issue, menu, components, fenix, button, page, crash, new, ui, strings, version |
| | Brave | brave, x, release, android, chromium, add, desktop, test, run, rewards, manual, ads, tab, browser, upgrade, wallet, new, settings, page, update |
| OS | Nextcloud | app, upload, bump, android, crash, file, nextcloud, fix, folder, auto, stable, add, error, rc, new, use, account, server |
| | Owncloud | android, upload, new, file, app, folder, fix, owncloud, feature, release, bug, request, update, share, add, view, arch, bump |

category app pair, the non-shared words in Firefox contain tabs, browser, pages, ads, which are associated with the traditional features of web browser: advertisement, web pages, etc.

To sum up, the bug reports of apps in the same category share more similar information than the apps in different categories, and the shared information is prone to the specific features related to the category with the frequently used words.

Hypothesis \Rightarrow “With the commonness/difference on the quantitative and qualitative analysis of bug reports in the same/different categories, we would like to intuitively hypothesize that it would be feasible to employ experience (e.g., bug reports) from same-category apps and the community (reviews) of target app *B* to discover potential bugs of *B*.”

3 BugRMSys

Figure 6 depicts the general workflow of the approach for app review driven collaborative bug finding.

Notation: we will refer to app A and app B as two apps that belong to the same category.

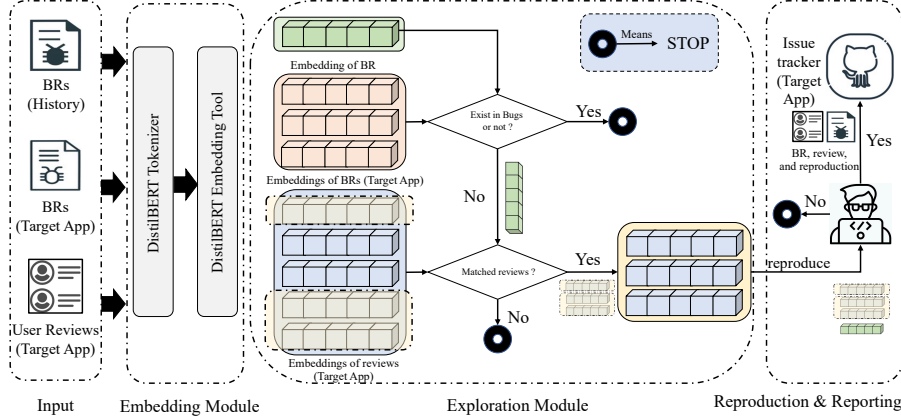


Fig. 5: Overview of BUGRMSys.

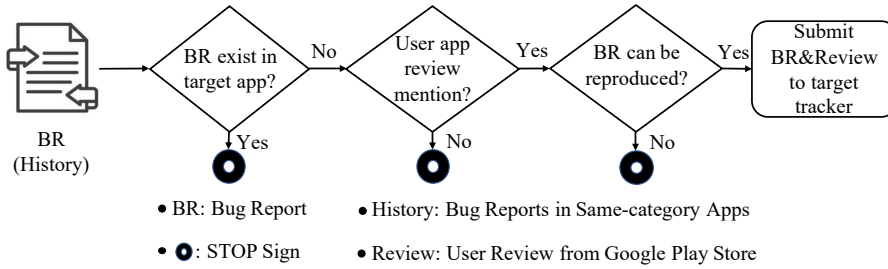


Fig. 6: General workflow of app review driven collaborative bug finding.

Our approach attempts to leverage the development experience of the historical app A to find bugs in the target app B. The workflow therefore starts with a representative bug report that has been handled in the development of app A. If a similar bug report exists in the issue tracker of the target app B under study, the bug finding process is halted and must be restarted with another bug report from app A. Otherwise, the workflow proceeds to check whether the bug report content is similar to some user reviews (that has therefore gone overlooked). If one or several relevant app reviews are found, we must attempt to reproduce the bug in app B based on reproduction steps in the bug report of app A as well as specific details in user review of B. In our evaluation, once the buggy behavior is confirmed through reproduction of the bug report, we further submit a new bug report in the issue tracker of app B. In the remainder of this section, we will present a real-world example before detailing the technical approach for bug recommendation.

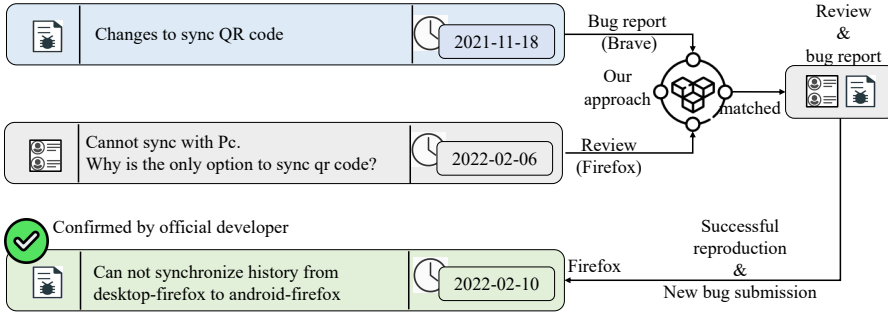


Fig. 7: Running example of BUGRMSys.

3.1 Running Example

Figure 7 illustrates the case where we leveraged BUGRMSys to discover a new bug in the web browser app Firefox. By iterating over bug reports from the active development repository of Brave, we identified a bug report which refers to synchronization with QR Code. A similar bug report was absent from the issue tracker of Firefox. After matching by BUGRMSys, a user review of Firefox had clearly stated a similar problem “*Cannot sync with Pc. Why is the only option to sync qr code?*”. With the user’s assessment, the bug report might indeed be relevant, we thus explore the steps enumerated in Brave’s bug report to reproduce the matched problem in Firefox. Surprisingly, the bug was successfully reproduced within a few minutes. We then submitted a bug report with two screenshots into the issue tracker of Firefox, which was eventually confirmed by the Firefox development team in 4 days.

Intuitively, our bug recommendation could have started with considering available user app reviews and try to correlate with historical bug reports from other apps in the same category of the target app. Unfortunately, in practice, most app reviews do not provide usable information. Therefore, we propose to initiate the search with the bug reports, which are in lesser numbers, and are more structured. Nevertheless, many bug reports in the same category as the target app are actually raising irrelevant issues. Therefore, it is important to further check if such potential issues have caught user attention and lead them to write reviews that mention them. This motivates the need to devise a reliable mechanism to precisely match useful app reviews with relevant bug reports.

3.2 Approach

Figure 5 illustrates the details of automatically finding bugs with BUGRMSys, which includes two main modules: an embedding module and an exploration module. The remainder of this section describes how the embedding module deals with bug report and app review text representation, and how the exploration module eventually identifies relevant bug reports for recommendation.

3.2.1 Embedding

To embed app reviews and bug reports into vector representation for natural language text semantic similarity computation, we resort to state-of-the-art pre-trained deep learning models. In practice, when reasoning about app reviews and bug reports, we face at least two major challenges:

1. App reviews and bug reports are organized differently, and their vocabularies differ from each other (technical vs non-technical terms). In addition, app reviews are more error-prone (e.g., spelling mistakes), may include emoticons, and may contain non-technical references (e.g., swearing). Reviews also often repeat, i.e., the same review can be submitted by the same user several times.
2. Bug reports are written based on a technical template that includes description, title, reproduction steps. They can also include code-specific technical terms such as test names (e.g., "RewardsBrowserTest"). A complete bug report is often very long, hence difficult to capture overall semantics.

Towards providing a precise and effective representation model for bug reports and app reviews, we address each challenge specifically in BUGRMSys:

For **challenge (1)**, we use common NLP pre-processing techniques to eliminate noisy words and/or reviews. We first lowercase all tokens and remove punctuation as well as stop words [35,51,44]. We then use a regularization function to remove digits and emoticons. In addition, we drop too short and too long reviews based on outlier thresholds: empirically, we determined that reviews containing fewer than 10 words or more than 200 words should be dropped as they are outliers. We adopt the rule-based methods in [28,52] to correct word repetitions and spelling mistakes.

For **challenge (2)**, we must carefully process bug reports to extract relevant information. Indeed, prior work [60,31] have established that the detailed bug report descriptions usually contain significant noise, including the report template itself. Previous study stated that the title of a bug report already includes the essential content of the bug report description [18,16,53]. Thus, in our approach, we consider the bug report titles as input to drive the recommendation system. Other details (e.g., reproduction steps) will be considered in the reproduction and submission phase.

To address the challenge for capturing semantic features from the app reviews and bug reports, we resort to an effective method of converting bug reports and app reviews into numerical representations. There are various techniques [58,29,36,33,15] in the NLP community for mapping texts into tokens and low-dimension representations. In recent years, pre-trained deep learning methods [5,38] have become popular across communities given their practicality and their effectiveness. In this work, we consider DistilBERT [38], a pre-trained lightweight network model built on the top of BERT [5], that offers similar performance with transformer-based models as BERT, but requires significantly less computer resource and training/tuning time than BERT. And DistilBERT employs a byte pair encoding method [40] as its word tokenizer

which empowers it to alleviate “out of vocabulary” issues and to recognize compound words, especially for variables and camel-case (and other composition-based) named elements. We reuse the same function defined in [5] and apply its default setting to get the embeddings of both bug reports and app reviews. The embeddings of bug reports and app reviews can be used for calculating the document’s cosine similarity.

3.2.2 Exploration

Once the embedding module has produced numeric representations for bug reports and app reviews, the exploration module attempts to identify those reviews that are relevant for a given bug report. If a user review is found, the bug report can be recommended (supposing that no similar bug report was already filed for the target app). To this end, following empirical findings from previous studies [8] dealing with text similarity metrics [55], we resort to employing the cosine similarity [42] between low-dimensional representations of bug reports and app reviews to measure their relatedness. BUGRMSYS then prioritizes app reviews as being potentially relevant or not by referencing a threshold value to dismiss app reviews that should not be listed. Therefore, if one app review of B is found to be similar with a bug report in A, the bug report will be recommended as describing a bug that is relevant for the target app B. Otherwise, the bug report will not be recommended. The threshold setting is presented in Section 5.1.

4 Experimental Design

4.1 Research Questions

Our experiments aim at answering the following research questions:

- **RQ-1: To what extent is the BugRMSys bug recommendation approach driven by app reviews effective?** Our approach is devised to recommend bugs with our hypothesis. In this RQ, we first investigate to what extent app reviews can be used to recommend the bug report for the target app. Before that, we build a “ground truth” set of bug reports pairwise combinations (from a target app and a same category app).
- **RQ-2: Can BugRMSys expose bugs in real-world apps based on existing app reviews?** With this research question, we explore the feasibility of exposing bugs in real-world apps by BUGRMSYS with 20 apps of 9 categories. We investigate both the Exploration module for bug recommendations, as well as the actual bugs that developers can confirm.
- **RQ-3: How does BugRMSys compare against previous bug recommendation systems?** With this research question, we compare the bug recommendation approach of BUGRMSYS against the bug recommendation systems (DeepMatcher [10] and Bugine [49]) in three dimensions (inputs, automation, effectiveness).

4.2 Dataset

We investigate the feasibility of our hypothesis with BUGRMSYS, we first curate a dataset that collects bug reports and app reviews from 20 free and open source apps of 9 categories. For each app, we collect bug reports from its GitHub issue tracking system. App reviews are selected by checking whether an app review is related to a concrete bug report. App reviews are further curated following the order given after a sort on “helpful”. A review is given a “helpful” score according to the number of users who agree with it. In a recent study study, Häring et al. [10] have used this score to measure the importance of a review. Note that, collected bug reports and app reviews are written in English. Table 3 summarizes the information of collected data.

Table 3: Dataset Summary.

| App Category | # Apps | # bug reports | # app reviews |
|-----------------------|--------|---------------|---------------|
| privacy security | 2 | 15,657 | 15,428 |
| web browser | 2 | 45,523 | 166,360 |
| office suites | 2 | 13461 | 2,260 |
| emulator | 3 | 25,822 | 43,709 |
| communication | 2 | 4,555 | 3,166 |
| game | 2 | 1,590 | 3,098 |
| multimedia | 2 | 5,898 | 53,746 |
| reading | 2 | 3,704 | 62,755 |
| science and education | 3 | 16,045 | 36,413 |
| Total | 20 | 132,255 | 386,935 |

4.3 Evaluation Metrics

Our study leverages a variety of metrics to validate the experiments. First, we rely on overlap rate analysis [17] to investigate our hypothesis about the similarity of bug reports across same-category apps.

Overlap Rate: Given two sets X and Y , Overlap rate of X with Y is computed as follows:

$$Overlap(X_Y) = \frac{size(|X \cap Y|)}{size(X)}, \quad (1)$$

where $size$ denotes the size function for sets. If both sets have the same size (e.g., in our case, we select top $_K$ frequent words in the different sets of bug reports), then $Overlap(X_Y) == Overlap(Y_X)$.

To evaluate the overall performance of BUGRMSYS, we use Acc@N and Mean Reciprocal Rank (MRR), which are widely used metrics for recommender systems [57, 59, 37].

Acc@N: Acc@N hit measures the retrieval precision over the top_N recommended issues or reviews in the ranked list:

$$Acc@N = \frac{\sum_i^N(pair(i))}{LENGTH}, \quad (2)$$

where *LENGTH* represents the length of tested ground truth; pair(i) means if the *i*-th issue *B* hit the *i*-th top_N reviews relevant to issue *A*, if yes, *pair(i)* = 1, else 0. Overall, *Acc@N* is an approach used in previous research describing how often the issue in target *B* is among the top_N Nearest Neighbours (by cosine) of an DistilBERT word space.

MRR: *MRR* is short for mean reciprocal rank and is a popular metric used to evaluate the efficiency of recommendation system [41, 4, 27]. The equation of accuracy is described as follows:

$$MRR = \frac{1}{N} * \sum_{i=1}^{|N|} \frac{1}{rank_i}, \quad (3)$$

where *N* is the length of ground truth; For the *i*-th issue (app A) in ground truth, rank_{*i*} represents the position of recommended review which also relevant to issue in app B.

4.4 Implementation & Availability

We first developed two crawlers for automatically collecting bug reports and app reviews based on python packages: PyGithub³ and google play scraper⁴, respectively. To filter most useless reviews, we only select reviews in the text size from 10 to 200. We have implemented a prototype version of BUGRMSYS using Python (and associated frameworks) with a well-known, light-weight, transformer-based, and contextual pre-trained model, DistilBERT, to extract vector representations for both bug reports and app reviews. The dataset and the replication package of BUGRMSYS are publicly available at:

<https://github.com/BugRMSys/BugRMSys>

5 Experimental Results

In this section, we conduct qualitative and quantitative analysis to evaluate BUGRMSYS. To this end, we evaluate the effectiveness of BUGRMSYS, we compare BUGRMSYS with state-of-the-art tools for bug recommendation, we study the characteristics of BUGRMSYS, and we explore the transferability of BUGRMSYS.

³ <https://pypi.org/project/PyGithub/>

⁴ <https://pypi.org/project/google-play-scraper/>

5.1 [RQ-1]: Effectiveness of BUGRMSYS

To answer RQ1, due to the huge manual effort for assessing similarity of app reviews with bug reports, we focus on a single same-category app pair (Firefox, Brave), where Firefox will be the input app and Brave the target app (i.e., we use bug reports from Firefox to find new bug reports in Brave by matching Brave user reviews). Note that in the running example (cf. Section 2), we had illustrated a bug case where Firefox was the target app and Brave was the input app. This is an additional argument that BUGRMSYS can explore the experience of any app and leverage it for any other same-category counterpart regardless of which app appears to have more historical data.

We start by building a ground truth dataset to assess the ability of BUGRMSYS to find relevant bug reports for recommendation. To that end, the idea is to first try to find existing Brave bug reports that are similar to the ones of Firefox to build a set of pairs of similar bug reports $\langle BR_{Firefox}, BR_{Brave} \rangle$. Then, for each pair $\langle BR_{Firefox}, BR_{Brave} \rangle$, we rely on BUGRMSYS to identify Brave app reviews that match $BR_{Firefox}$. However, we only consider Brave app reviews which precede the creation time of the corresponding BR_{Brave} . Finally, we manually check if the identified Brave app reviews match the corresponding Brave bug report BR_{Brave} . This would indicate that BUGRMSYS would have been useful to automatically recommend the Firefox bug report as relevant to Brave.

In practice, to build our set of similar bug reports, we randomly picked 3,000 bug reports from Firefox. By using a cosine similarity threshold of 0.91⁵, we were able to identify 81 bug report pairs $\langle BR_{Firefox}, BR_{Brave} \rangle$ where the Firefox and Brave bug reports are highly similar. Note that this ground truth construction may be too conservative: there are possibly other Brave bug reports that are also semantically similar to a given Firefox bug report.

Recall that the objective of BUGRMSYS is to match bug reports to app reviews. In this case, we assess whether, for each pair of the ground truth, we can match the Firefox bug report with app reviews that are relevant (based on human expertise⁶) to the associated Brave report. If so, we can conclude that BUGRMSYS would have been able to recommend the Brave bug report.

For our experiments, we applied BUGRMSYS on each of the 81 Firefox bug reports to match Brave app reviews. We retain only the top 3 matched reviews per bug report. Overall, among the 81 top₁ reviews recommended by BUGRMSYS, 21 could be confirmed as indeed semantically similar to the content in the Brave bug report associated in the ground truth pair.

We further confirm 32 and 38 (at Top 2 and Top 3 respectively) app review matches. Table 4 details our results by also providing the Accuracy and MRR scores. Overall, with Top 3 recommendations on a conservative ground truth,

⁵ We decided on 0.91 based on an empirical validation: the higher the similarity threshold is, the less pairs will be matched, and at the same time matched reviews are closer to bug reports. We decide on a high threshold to maximize high quality results.

⁶ This is the core challenge of BUGRMSYS: to find semantic similarity between a bug report and an app review

Table 4: Results of Acc@N and MRR@N. "App review Hits" represents the number of times BUGRMSYS matches the relevant app reviews associated to the ground truth bug report of Brave: this is a proxy for estimating that BUGRMSYS would have been able to recommend the Firefox bug report as relevant to Brave.

| 81 relevant bug pairs out of 3K bugs from Firefox | @1 | @2 | @3 |
|--|-------|-------|-------|
| App Review Hits | 21 | 32 | 38 |
| Value of Acc@N(%) | 25.93 | 39.51 | 46.91 |
| Value of MRR@N(%) | 25.93 | 26.50 | 35.19 |

we reach almost 50% hit ratio. Note that, while BUGRMSYS matches Firefox bug reports with Brave app reviews, our effectiveness evaluation is to check whether the matched app reviews are semantically relevant for Brave bug reports. This is the practical and ultimate concern of our bug recommendation scheme.

Table 7 presents an extract of our set of 81 similar bug reports pairs by focusing on three pairs, the corresponding app user reviews and the result of our manual analysis. The full results are detailed on our Github repository⁷.

↪ **Answer to RQ-1:** ▶ *The Acc@N and MRR@N values show that BUGRMSYS is reasonably effective in matching relevant app reviews of a target app to bug reports from same-category apps in order to drive bug recommendation.* ◀

5.2 [RQ-2]: Feasibility of BUGRMSYS

5.2.1 Bug recommendation in the wild

We conduct extensive execution of BUGRMSYS on data from 20 apps in 9 categories to recommend bugs. Due to space limitation, we report in Table 5 the statistics of bugs recommended by BUGRMSYS for the top₁₀ apps having the most recommended bugs. We note that, thanks to BUGRMSYS app-review driven approach, the collaborative bug finding allows to sift between a few hundreds to a few thousands bug reports from same-category apps in order to recommend⁸ only a few (1.63%) of bug reports as being relevant to the target apps.

Since each recommended bug is found by correlating information in its app reviews, we propose to estimate the potential time gain BUGRMSYS has brought by highlighting the buggy behaviour users complained about in unofficial channels. We compute the distribution of time elapsed since the app

⁷ https://github.com/BugRMSys/BugRMSys/blob/main/RQ1UserCase_firefox2brave.csv

⁸ We set a high similarity threshold at 0.9. This value can be fine-tuned following the practitioner’s objectives.

Table 5: Ranked apps based on the number of potential bugs.

| Target app | # of bug reports searched (from same category apps) | # app reviews (for the target app) | # recommended bugs |
|-------------|---|------------------------------------|--------------------|
| Brave | 10,000 | 10,000 | 208 |
| Nextcloud | 3,489 | 1,143 | 147 |
| Wire | 10,000 | 2,515 | 75 |
| VLC | 599 | 5,556 | 59 |
| Firefox | 10,000 | 10,000 | 52 |
| Dolphin | 2,437 | 2,498 | 44 |
| Wordpress | 4,139 | 2,714 | 44 |
| PPSSPP | 1,666 | 6,846 | 39 |
| Mupen64Plus | 2,437 | 1,401 | 39 |
| Fbreader | 726 | 3,914 | 33 |

review creation date and the BUGRMSYS bug recommendation date. On average, specially for Firefox and Brave in Table 5, the app reviews were created 22.2 and 33 days before we submit the bug reports, respectively.

✎ **Answer to RQ-2.a** ► BUGRMSYS can help triage bug reports from same-category apps to recommend a reasonable number of bugs in a target app. We also show that BUGRMSYS helps to highlight bugs that could have remained overlooked for a long time in app reviews, e.g. 22.2 days earlier for Firefox, and 33 days earlier for Brave. ◀

To assess whether the bugs recommended by BUGRMSYS are real bugs, we further reproduce each recommended bugs manually, and submit the successfully reproduced bugs in the issue tracker of the related app. Since the manual reproducing work requires extensive efforts, we will focus our investigation on four popular apps: Wire, Brave, FireFox and Nextcloud.

The numbers of recommended bugs, (successfully) reproduced bugs (these reproduced bugs are selected with their similarities, and the top-44 most similar ones are selected), and confirmed or fixed bugs are presented in Table 6. The details of the reported bugs can be found on our repository⁹. We remind that we can recommend a bug report from app A (the first column in Table 6) as relevant to the target app B when BUGRMSYS matches the bug report of A with user app reviews from the target app B. We then use the “steps to reproduce” present in the bug report of A, as well as the information present in the app reviews of app B to reproduce the bug in the target app B. Finally, for each bug that has been successfully reproduced in the target app B, we submit the bug in the issue tracker of the app. As shown in the the last column of Table 6, six bugs have been already confirmed or fixed by the developers before this submission.

⁹ https://github.com/BugRMSys/BugRMSys/blob/main/new_bugs.md

Table 6: Previously unknown bugs detected with BUGRMSYS.

| Input app | Target app | # of bug reports (input apps) | # app reviews (target app) | # recommended bugs | # of reproduction attempts | # successfully reproduced bugs | # replied, confirmed or fixed bugs |
|-----------|------------|-------------------------------|----------------------------|--------------------|----------------------------|--------------------------------|------------------------------------|
| Signal | Wire | 10,000 | 2,515 | 75 | 12 | 2 | 2, ✓ 2 |
| Firefox | Brave | 10,000 | 10,000 | 208 | 44 | 9 | (1), 2, ✓ 2 |
| Brave | Firefox | 10,000 | 10,000 | 52 | 24 | 4 | (1), 1 |
| Owncloud | Nextcloud | 3,489 | 1,143 | 147 | 12 | 5 | (1), 1 |
| | | | | 482 | 90 | 20 | (3), 6, ✓ 4 |

*“(#)” means the number of reported bugs replied by developers but not confirmed or fixed by them.

“✓ #” means the number of reported bugs confirmed and fixed by developers.

Table 7: Extract of our Ground Truth Dataset, Corresponding relevant app reviews, and Manual check result.

| Existing Bug Reports (input, target) | TOP 3 RELEVANT REVIEWS in Brave (creation time always prior to the corresponding Brave bug report) | MANUAL |
|--|---|-------------------------|
| App: Firefox Data: 2020-08-21 Reports: Download does not work on a custom tab (Slack) | App: Brave Date: 2021-03-20 Report: Download [Status Bar] Improvement Date: 2020-08-10 Review: Downloader is very bad...pls increase and more work on download manager Date: 2020-10-18 Review: There is no download option in this could u pls update on this issue Date: 2020-10-18 Review: ...it does not allow to manually add download tasks | TRUE TRUE TRUE |
| App: Firefox Data: 2020-06-17 Reports: Report clickbait sites, Protect user privacy | App: Brave Date: 2020-06-22 Report: Warn users about insecure Facebook and Google privacy settings Date: 2020-06-08 Review: Only browser that cheats about privacy. All claims about user privacy are bogus... Date: 2020-11-01 Review: Extremely convoluted privacy practices. They advocate for privacy but allow certain creepy sites ... Date: 2020-12-31 Review: Good privacy app. It doesn't prevent websites from annoying redirections | TRUE FALSE FALSE |
| App: Firefox Data: 2020-08-28 Reports: Invalid URLs can be bookmarked and they crash the browser | App: Brave Date: 2020-11-06 Reports: Clicking URLs outside of Brave opens a blank browser window with no URL Date: 2020-06-11 Review: "... has come under fire for automatically redirecting URLs typed into the browser's address bar ... Date: 2020-10-20 Review: One of the best browsers I'mo. Only wish I could set links on the brave homepage manually... Date: 2020-10-13 Review: Its a good browser sometimes it reload all tabs when i open newly | FALSE FALSE FALSE |

📌 **Answer to RQ-2.b:** ▶ Thanks to BUGRMSYS, we found 20 new bugs efficiently (i.e., within a short period of time), among which six have been confirmed or fixed by developers. These results demonstrate that BUGRMSYS is relevant for exploring news bugs using its app reviews driven collaborative bug finding scheme. Since there is no training involved in BUGRMSYS, it can be readily applied to various apps from various categories to help developers find bugs that have not yet been officially reported but which users may have witnessed already. ◀

5.3 [RQ-3]: BUGRMSYS vs Prior Works

Ideally, we should evaluate the performance of BUGRMSYS in comparison with prior works dealing with bug recommendations based on bug reports. There are two state of the art approaches, DeepMatcher [10] and Bugine [49], which are closely related.

DeepMatcher and BUGRMSYS both match app reviews with bug reports based on text embedding using pre-trained DistilBERT. Experimentally, we compare DeepMatcher against the BUGRMSYS by considering the ground truth data built for RQ-1 (cf. examples in Table 7): we propose to manually check whether the matched reviews with both approaches are relevant or not. While all reviews matched by BUGRMSYS are relevant to buggy behaviour, we observe that DeepMatcher only achieves a F1-score of 71% in filtering useful reviews. We postulate that BUGRMSYS performs better partly because it implements a focused collaborative bug finding approach where the matching is done on bug reports of same-category apps.

We also compare against Bugine, which also performs collaborative bug finding. We differ however as Bugine limits the matching to cases where apps

have the UI/components (while we consider apps from the same categories). We further consider app reviews to drive bug recommendation.

In the remainder, we further elaborate on the specific differences that prevent comparison between prior works and BUGRMSYS. These differences relate to three aspects: (1) Differences in input, output, and workflow; (2) Differences in automation level; (3) Performance in new bug finding.

5.3.1 Differences in input, output, and workflow

DeepMatcher employs user reviews for App B as input and recommends relevant bug reports for App B. The workflow is: App review \rightarrow Problem report \rightarrow Matched relevant bug reports. They only evaluate their tool on existing bugs instead of exploring new bugs. Furthermore, DeepMatcher does not leverage experience from other apps when investigating a target app. Their approach further suffers from the redundancy problem in app reviews.

Bugine employs issues in apps with same UI components as their database. They focus on building a automatic test generation from bug reproduction steps and run the test on target app with manual check. Bugine has been used to explore new bugs successfully. However, there is a great limitation in this approach: it only considers app issues with same-UI components into consideration, which can reduce the feasibility of learning from other apps.

5.3.2 Difference in automation level

Different from DeepMatcher, BUGRMSYS will not process a large number of bug reports: we focus on same-category apps to match relevant reviews of App B. After manual check, we have verified that when we feed a bug report into BUGRMSYS, the matched reviews are 100% related to some bugs. By building on same category apps (i.e., with similar functionality and usage steps) reproduction and localization of bugs is eased.

For Bugine, finding apps with same-UI components is a time-consuming task. In addition, using issues from same-UI apps makes it hard to transfer the learned expertise to other types of issues.

5.3.3 Performance in bug finding

The ability of DeepMatcher to find new bugs has not been evaluated. Bugine reported having found 34 new bugs in 5 evaluated apps. With BUGRMSYS, within a week, we were able to recommend, reproduce and identify 20 new bugs across 6 apps. 4 such bugs are already fixed by the app developers.

➤ **Answer to RQ-3:** ➤ BUGRMSYS, in comparison to Deepmatcher, is effective for filtering out relevant app reviews. Compared to Bugine, BUGRMSYS is scalable and can be applied to a larger range of bug types, while avoiding duplication of recommending bugs that were already reported in the target app. ◀

6 Discussion

6.1 Failures to reproduce recommended bugs

As illustrated previously, some of the reproduction attempts on the bug reports recommended by BUGRMSYS. lead to failures There are various reasons that explain such failures without suggesting that the recommended bug is not relevant. Prior studies have already largely elaborated on this difficulty to reproduce bugs: In Han’s work [11], an extensive classification of 8 categories of root causes for failed reproductions is provided: hardware dependency, operating system dependency, component dependency, unavailable source code, compilation error, installation error, missing step, and lack of symptom. Our failures causes span across these categories.

6.2 Threats to validity

Our design, implementation and evaluation of BUGRMSYS carries some threats to validity. First, when we are building the ground truth, we manually check whether the reviews are meaningful. Therefore, the ground truth may be biased by our own experience. Second, BUGRMSYS is not fully automated, i.e., manual effort is still needed when we reproduce from the recommended bug reports. Consequently, the success rate of reproduction could be dependent on the developing experience of individual developer.

7 Related Work

Collaborative Experience Sharing: Collaborative programming is common in the development of open source software. Consequently, similar bugs can emerge across different projects. Other researches attempted to leverage this fact to Recommend, Reproduce, and Repair inter-project bugs. For instance, the detection of duplicate bug reports have been studied in localizing fault of software [46,53]. Specifically, Yang *et al.* [56] combined information retrieval technique and word embedding technique to process the detailed information of bug reports to recommend similar bugs. On the other hand, Tan *et al.* [49] use three collaborative sources for bug finding: (1) bugs from the same programmer across different projects, (2) bugs from manually searching for bug reports in GitHub repositories, (3) bugs from a bug recommendation system. Based on these shared experience, they explored the concept of collaborative bug finding on improving the teaching of software testing course. In the experience-based collaborative learning of crowd-sourcing, Mao *et al.* [30]’s experimental results show that the generated replicable test scripts from crowd-based testing can improve the coverage attainment for automated mobile testing.

Recommendation of Bug Reports Based on App Reviews: The importance of app reviews in App vendors have been comprehensively demonstrated [34, 39]. Leaving the app reviews not to be addressed is harmful to the experience of the users and rating of the app, and further lead to uninstallation of the app [12]. To maintain the evolution of the app, researches started to leverage the app reviews. For instance, Gao *et al.* [7] developed a novel approach to automatically generate proper responses to the app reviews in Google Play. However, this approach mainly try to (1) soothe bad emotion of users, and (2) collect detailed user experience, but not to discover potential bugs in advance. Tan *et al.* [49] designed an approach to find bugs for Android apps. Their pipeline is to retrieve bugs in other similar apps that may also exist in the current app. This work validates the feasibility of searching for bugs in other projects to identify new bugs mentioned by app reviews. Afterwards, Marlo *et al.* [10] try to match bug reports with related app reviews to discover bugs by filling the gap of different languages between app reviews written by non-technical users and bug reports proposed by professional developers.

8 Conclusion and Future Work

In this paper, we introduce BUGRMSYS, a tool-supported approach for app reviews driven collaborative bug finding. Given a target app B, BUGRMSYS builds on the development experience of app A to identify bug reports in A that match app reviews of B. If such bug reports exist, they are considered as candidate for recommending bugs to the target app B. To that end, BUGRMSYS implements an embedding procedure to represent bug reports and app reviews text, and use cosine similarity to decide on matching similarity scores. Once bugs are recommend, we experimentally attempt reproduction to confirm the detection of new bugs in the target app. Our experimental results on free and open source apps in various categories show that BUGRMSYS is effective, scales to a variety of bug types, and does not yield too many irrelevant app review matches. Overall, with BUGRMSYS, we already successfully reproduced 20 new bugs in 6 apps across 3 categories. Several of these bugs have been acknowledge by the apps development communities and some have even already been fixed.

In future, we plan to address the question of automating the reproduction phase in order to scale the collaborative bug finding approach towards further increasing its practicality in real-world debugging scenarios.

References

1. Amalfitano, D., Riccio, V., Paiva, A.C., Fasolino, A.R.: Why does the orientation change mess up my android application? from gui failures to code faults. *Software Testing, Verification and Reliability* **28**(1), e1654 (2018)
2. Bevan, J., Werner, L., McDowell, C.: Guidelines for the use of pair programming in a freshman programming class. In: *Proceedings 15th Conference on Software Engineering Education and Training (CSEE&T 2002)*, pp. 100–107. IEEE (2002)

3. Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O’Hearn, P., Papakonstantinou, I., Purbrick, J., Rodriguez, D.: Moving fast with software verification. In: *NASA Formal Methods Symposium*, pp. 3–11. Springer (2015)
4. Cames, K., II-Grants, A.A.: Recommendation. City (2006)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
6. Fan, L., Su, T., Chen, S., Meng, G., Liu, Y., Xu, L., Pu, G., Su, Z.: Large-scale analysis of framework-specific exceptions in android apps. In: *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*, pp. 408–419. IEEE (2018)
7. Gao, C., Zeng, J., Xia, X., Lo, D., Lyu, M.R., King, I.: Automating app review response generation. In: *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pp. 163–175. IEEE (2019)
8. Ghosh, J., Strehl, A.: Similarity-based text clustering: A comparative study. In: *Grouping Multidimensional Data*, pp. 73–97. Springer (2006)
9. Haddi, E., Liu, X., Shi, Y.: The role of text pre-processing in sentiment analysis. *Procedia computer science* **17**, 26–32 (2013)
10. Haering, M., Stanik, C., Maalej, W.: Automatically matching bug reports with related app reviews. In: *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pp. 970–981. IEEE (2021)
11. Han, X., Carroll, D., Yu, T.: Reproducing performance bug reports in server applications: The researchers’ experiences. *Journal of Systems and Software* **156**, 268–282 (2019)
12. Hassan, S., Tantithamthavorn, C., Bezemer, C.P., Hassan, A.E.: Studying the dialogue between users and developers of free apps in the google play store. *Empirical Software Engineering* **23**(3), 1275–1312 (2018)
13. Hu, G., Yuan, X., Tang, Y., Yang, J.: Efficiently, effectively detecting mobile app bugs with appdoctor. In: *Proceedings of the Ninth European Conference on Computer Systems*, pp. 1–15 (2014)
14. Jiang, H., Yang, H., Qin, S., Su, Z., Zhang, J., Yan, J.: Detecting energy bugs in android apps using static analysis. In: *International Conference on Formal Engineering Methods*, pp. 192–208. Springer (2017)
15. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759* (2016)
16. Ko, A.J., Myers, B.A., Chau, D.H.: A linguistic analysis of how people describe software problems. In: *Visual Languages and Human-Centric Computing (VL/HCC’06)*, pp. 127–134. IEEE (2006)
17. Kowalski, G.: *Information retrieval architecture and algorithms*. Springer Science & Business Media (2010)
18. Lamkanfi, A., Demeyer, S., Giger, E., Goethals, B.: Predicting the severity of a reported bug. In: *2010 7th IEEE Working Conference on Mining Software Repositories (MSR 2010)*, pp. 1–10. IEEE (2010)
19. Lee, S., Dolby, J., Ryu, S.: Hybridroid: static analysis framework for android hybrid applications. In: *2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pp. 250–261. IEEE (2016)
20. Li, H., Fang, C., Wei, Z., Chen, Z.: Cocotest: collaborative crowdsourced testing for android applications. In: *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 390–393 (2019)
21. Li, H., Zhang, L., Zhang, L., Shen, J.: A user satisfaction analysis approach for software evolution. In: *2010 IEEE International Conference on Progress in Informatics and Computing*, vol. 2, pp. 1093–1097. IEEE (2010)
22. Liu, Z., Chen, C., Wang, J., Huang, Y., Hu, J., Wang, Q.: Guided bug crush: Assist manual gui testing of android apps via hint moves. *arXiv preprint arXiv:2201.12085* (2022)
23. Long, T., Yoon, I., Memon, A., Porter, A., Sussman, A.: Enabling collaborative testing across shared software components. In: *Proceedings of the 17th international ACM Sigsoft symposium on Component-based software engineering*, pp. 55–64 (2014)
24. Long, T., Yoon, I., Porter, A., Memon, A., Sussman, A.: Coordinated collaborative testing of shared software components. In: *2016 IEEE International Conference on Software Testing, Verification and Validation (ICST)*, pp. 364–374. IEEE (2016)

25. Loper, E., Bird, S.: Nltk: The natural language toolkit. arXiv preprint cs/0205028 (2002)
26. Maalej, W., Kurtanović, Z., Nabil, H., Stanik, C.: On the automatic classification of app reviews. *Requirements Engineering* **21**(3), 311–331 (2016)
27. Mahmood, T., Ricci, F., Venturini, A.: Improving recommendation effectiveness: Adapting a dialogue strategy in online travel planning. *Information Technology & Tourism* **11**(4), 285–302 (2009)
28. Man, Y., Gao, C., Lyu, M.R., Jiang, J.: Experience report: Understanding cross-platform app issues from user reviews. In: 2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE), pp. 138–149. IEEE (2016)
29. Manning, C., Schütze, H.: Foundations of statistical natural language processing. MIT press (1999)
30. Mao, K., Harman, M., Jia, Y.: Crowd intelligence enhances automated mobile testing. In: 2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 16–26. IEEE (2017)
31. Martens, D., Maalej, W.: Extracting and analyzing context information in user-support conversations on twitter. In: 2019 IEEE 27th International Requirements Engineering Conference (RE), pp. 131–141. IEEE (2019)
32. McIlroy, S., Ali, N., Hassan, A.E.: Fresh apps: an empirical study of frequently-updated mobile apps in the google play store. *Empirical Software Engineering* **21**(3), 1346–1370 (2016)
33. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
34. Oh, J., Kim, D., Lee, U., Lee, J.G., Song, J.: Facilitating developer-user interactions with mobile app review digests. In: CHI’13 Extended Abstracts on Human Factors in Computing Systems, pp. 1809–1814 (2013)
35. Panichella, S., Di Sorbo, A., Guzman, E., Visaggio, C.A., Canfora, G., Gall, H.C.: How can i improve my app? classifying user reviews for software maintenance and evolution. In: 2015 IEEE international conference on software maintenance and evolution (ICSME), pp. 281–290. IEEE (2015)
36. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pp. 1532–1543 (2014)
37. Poerner, N., Waltinger, U., Schütze, H.: E-BERT: Efficient-yet-effective entity embeddings for BERT. In: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 803–818. Association for Computational Linguistics, Online (2020). DOI 10.18653/v1/2020.findings-emnlp.71. URL <https://aclanthology.org/2020.findings-emnlp.71>
38. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108 (2019)
39. Sefferman, A.: Survey on user ratings and reviews. <https://www.apptentive.com/blog/2020/02/04/mobile-app-ratings-and-reviews/> (2015). [Online; accessed 2015]
40. Sennrich, R., Haddow, B., Birch, A.: Neural machine translation of rare words with subword units. arXiv preprint arXiv:1508.07909 (2015)
41. Shani, G., Gunawardana, A.: Evaluating recommendation systems. In: Recommender systems handbook, pp. 257–297. Springer (2011)
42. Sitikhu, P., Pahi, K., Thapa, P., Shakya, S.: A comparison of semantic similarity methods for maximum human interpretability. In: 2019 artificial intelligence for transforming business and society (AITB), vol. 1, pp. 1–4. IEEE (2019)
43. Stanik, C., Haering, M., Maalej, W.: Classifying multilingual user feedback using traditional machine learning and deep learning. In: 2019 IEEE 27th International Requirements Engineering Conference Workshops (REW), pp. 220–226. IEEE (2019)
44. Stanik, C., Montgomery, L., Martens, D., Fucci, D., Maalej, W.: A simple nlp-based approach to support onboarding and retention in open source communities. In: 2018 IEEE International Conference on Software Maintenance and Evolution (ICSME), pp. 172–182. IEEE (2018)

45. Su, T., Fan, L., Chen, S., Liu, Y., Xu, L., Pu, G., Su, Z.: Why my app crashes understanding and benchmarking framework-specific exceptions of android apps. *IEEE Transactions on Software Engineering* (2020)
46. Sun, C., Lo, D., Khoo, S.C., Jiang, J.: Towards more accurate retrieval of duplicate bug reports. In: 2011 26th IEEE/ACM International Conference on Automated Software Engineering (ASE 2011), pp. 253–262. IEEE (2011)
47. Sun, J., Su, T., Li, J., Dong, Z., Pu, G., Xie, T., Su, Z.: Understanding and finding system setting-related defects in android apps. In: Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis, pp. 204–215 (2021)
48. Talukder, M.A.I., Shahriar, H., Qian, K., Rahman, M., Ahamed, S., Wu, F., Agu, E.: Droidpatrol: a static analysis plugin for secure mobile software development. In: 2019 IEEE 43rd annual computer software and applications conference (COMPSAC), vol. 1, pp. 565–569. IEEE (2019)
49. Tan, S.H., Li, Z.: Collaborative bug finding for android apps. In: Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering, pp. 1335–1347 (2020)
50. Van Der Veen, V., Bos, H., Rossow, C.: Dynamic analysis of android malware. Internet & Web Technology Master thesis, VU University Amsterdam (2013)
51. Villarroel, L., Bavota, G., Russo, B., Oliveto, R., Di Penta, M.: Release planning of mobile apps based on user reviews. In: 2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE), pp. 14–24. IEEE (2016)
52. Vu, P.M., Nguyen, T.T., Pham, H.V., Nguyen, T.T.: Mining user opinions in mobile app reviews: A keyword-based approach. arXiv preprint arXiv:1505.04657 (2015)
53. Wang, X., Zhang, L., Xie, T., Anvik, J., Sun, J.: An approach to detecting duplicate bug reports using natural language and execution information. In: Proceedings of the 30th international conference on Software engineering, pp. 461–470 (2008)
54. Wilbur, W.J., Sirotkin, K.: The automatic identification of stop words. *Journal of information science* **18**(1), 45–55 (1992)
55. Wooditch, A., Johnson, N.J., Solymosi, R., Ariza, J.M., Langton, S.: Getting to know your data. In: A Beginner’s Guide to Statistics for Criminology and Criminal Justice Using R, pp. 21–38. Springer (2021)
56. Yang, X., Lo, D., Xia, X., Bao, L., Sun, J.: Combining word embedding with information retrieval to recommend similar bug reports. In: 2016 IEEE 27th international symposium on software reliability engineering (ISSRE), pp. 127–137. IEEE (2016)
57. Ye, X., Bunescu, R., Liu, C.: Learning to rank relevant files for bug reports using domain knowledge. In: Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, pp. 689–699 (2014)
58. Zhang, Y., Jin, R., Zhou, Z.H.: Understanding bag-of-words model: a statistical framework. *International Journal of Machine Learning and Cybernetics* **1**(1-4), 43–52 (2010)
59. Zhou, J., Zhang, H., Lo, D.: Where should the bugs be fixed? more accurate information retrieval-based bug localization based on bug reports. In: 2012 34th International Conference on Software Engineering (ICSE), pp. 14–24. IEEE (2012)
60. Zhou, Y., Tong, Y., Gu, R., Gall, H.: Combining text mining and data mining for bug report classification. *Journal of Software: Evolution and Process* **28**(3), 150–176 (2016)
61. Zimmermann, T., Premraj, R., Bettenburg, N., Just, S., Schroter, A., Weiss, C.: What makes a good bug report? *IEEE Transactions on Software Engineering* **36**(5), 618–643 (2010)