

LadyBug: A GitHub Bot for UI-Enhanced Bug Localization in Mobile Apps

Junayed Mahmud*, James Chen*, Terry Achille*, Camilo Alvarez-Velez*, Darren Dean Bansil*, Patrick Ijeh*, Samar Karanch*, Nadeeshan De Silva†, Oscar Chaparro†, Andrian Marcus‡, Kevin Moran*

*University of Central Florida (USA), †William & Mary (USA), ‡George Mason University (USA)

junayed.mahmud@ucf.edu, james.chen@ucf.edu, terry.achille@ucf.edu, ca840220@ucf.edu,
darrenbansil@ucf.edu, patrick.ijeh@ucf.edu, sa898457@ucf.edu, kgdesilva@wm.edu, oscarch@wm.edu,
amarcus7@gmu.edu, kpmoran@ucf.edu

Abstract—This paper introduces *LadyBug*, a GitHub bot that automatically localizes bugs for Android apps by combining UI interaction information with text retrieval. *LadyBug* connects to an Android app’s GitHub repository, and is triggered when a bug is reported in the corresponding issue tracker. Developers can then record a reproduction trace for the bug on a device or emulator and upload the trace to *LadyBug* via the GitHub issue tracker. This enables *LadyBug* to utilize both the text from the original bug description, and UI information from the reproduction trace to accurately retrieve a ranked list of files from the project that most likely contain the reported bug.

We empirically evaluated *LadyBug* using an automated testing pipeline and benchmark called *RedWing* that contains 80 fully-localized and reproducible bug reports from 39 Android apps. Our results illustrate that *LadyBug* outperforms text-retrieval-based baselines and that the utilization of UI information leads to a substantial increase in localization accuracy. *LadyBug* is an open-source tool, available at <https://github.com/LadyBugML/ladybug>.

A video showing the capabilities of *Ladybug* can be viewed here: <https://youtu.be/hI3tzBRK0Cw>

I. INTRODUCTION

The management of bug reports is a time-consuming process for developers [1, 2]. One of the most difficult bug report management tasks is localizing a failure described in a bug report to the code of the corresponding project. This requires reasoning between the natural language description of a given bug and the various programming languages with which a project is written. This task is further complicated by the often inadequate or incomplete information in bug reports [3, 4].

Researchers have proposed various techniques to automate the bug localization process. However, a large body of this prior work tends to formulate this problem as a text retrieval (TR) [5] problem, wherein a bug report is treated as a query and different granularities of source code (files, classes, methods, etc.) are ranked and presented to the programmer based on their likelihood of containing the described bug.

The primary assumption of TR-based bug localization is that bug reports and the corresponding source code will share terms; however, there exists a semantic gap between the information present in bug reports and the information available in source code [6]. Previous research proposed different techniques to mitigate this semantic gap in existing TR-based bug localization techniques. For instance, researchers have introduced techniques to preprocess the text in bug reports or source code, while other

techniques focus on query reformulation utilizing information from various resources [7]. Other lines of research focused on boosting the ranking of buggy code elements using execution information (e.g., stack traces [8]), code dependencies (e.g., static code analysis [9]), or app version history information (e.g., mining from GitHub repositories [8]).

In this paper, we propose a novel tool for bug localization of Android applications called *LadyBug*, which ranks buggy source code files using the bug descriptions from reports *and* a new, largely underexplored source – information from an app’s graphical user interface (GUI). The key intuition that underlies *LadyBug* is that manually or automatically collected bug reproduction traces which capture GUI-level interactions for recreating a bug can be utilized to help *refine* the localization process (i.e., filtering out files unlikely to contain the bug while boosting the rankings of files that are more likely to contain the bug.) As such *LadyBug* uses GUI interaction data in conjunction with a neural text embedding approach (i.e., UniXCoder [10]), to generate a ranked list of files based on the likelihood of containing bugs.

We constructed *LadyBug* as a GitHub bot that can be easily added to GitHub repositories of Android applications and used as an automated assistant for bug localization. *LadyBug*’s GitHub bot architecture consists of three key components: (i) a repository connector, built on top of the ProBot [11] framework, which allows for *Ladybug* to post/receive information to/from GitHub issues, (ii) a backend that indexes a given repository and implements the text retrieval engine and (iii) *MetaMorph*, a desktop application which allows a developer to easily record reproduction traces of a described bug. *LadyBug* can also be used in conjunction with recent automated reproduction tools [12] which automates the localization task end-to-end.

We empirically evaluated *LadyBug* using a benchmark called *RedWing* that contains 80 bug reports from 39 Android applications compared to a text retrieval baseline. Our results illustrate that (i) using GUI information improves the accuracy of file-level bug localization for Android apps, and (ii) *LadyBug* is able to outperform a text-only baseline. *LadyBug* is an open-source tool [13], and is compatible with any Android app. Further details about *LadyBug*’s algorithms and an expanded evaluation can be found in its original research paper [14].

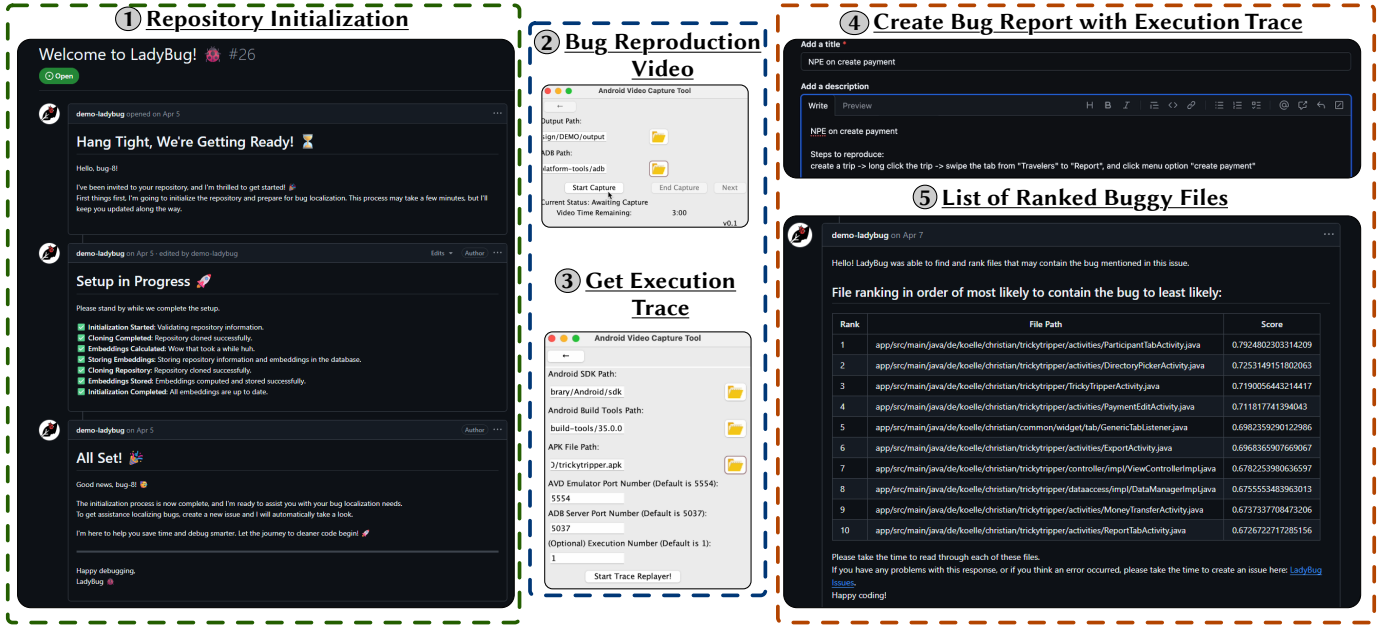


Fig. 1: The Typical User Workflow of the LadyBug GitHub Bot

II. LADYBUG: A GUI BUG LOCALIZATION TOOL FOR ANDROID APPLICATIONS

LadyBug is a GitHub bot that automates bug localization by analyzing Java source code and Native Android GUI interactions, reporting results directly within GitHub issue threads. LadyBug workflow is structured into three primary phases, illustrated in Figure 1. The first is the *Repository Initialization* phase (highlighted in green), where the Ladybug bot is installed on the target GitHub repository and granted the necessary permissions to monitor repository activities. Second, during the *Bug Reproduction* phase (highlighted in blue), the MetaMorph desktop application assists a developer with recording and replaying a reproduction trace of a given bug, generating an execution trace capturing detailed GUI interactions. Lastly, the *Issue Reporting* phase (highlighted in orange) involves developers submitting new GitHub issues to trigger the bug localization process, with an option to include the previously recorded execution trace to aid localization.

A. Repository Initialization

The first thing that a developer will need to do is add the Ladybug bot to their GitHub repository through the GitHub front-end (see Figure 1-①). After adding LadyBug to a given repository, LadyBug will create a new issue giving the developer a welcome message and incremental updates on the initialization process. When all these steps are completed, a confirmation message will be posted to the issue thread, indicating that LadyBug is ready for further use.

B. Bug Reproduction

When the MetaMorph application is opened (Figure 1-②), the developer will be tasked with providing system paths to the Android emulator's files, as well as system paths to the Android application. After the developer has completed providing the information requested, they are asked to open the application

inside of their emulator and press "Start Capture". Then, the developer will retrace the steps described in the original bug report, performing actions until the bug occurs in the app on the emulated device – and then hit "Stop Capture". Afterwards, the developer will supply the necessary information MetaMorph needs to retrace the steps the developer took in the bug report, and click on "Start Trace Replayer" (Figure 1-③). After this is done, MetaMorph writes this information into an execution.json file, which contains the necessary GUI information that LadyBug needs to perform GUI-enhanced bug localization.

C. Issue Reporting

With all the necessary GUI information, the developer can create a new GitHub issue with the trace and bug report (Figure 1-④). LadyBug will respond to the issue with status messages updating the user on the bug localization process. After localizing the bug, the LadyBug bot will reply with the ranked list of potential buggy files (Figure 1-⑤).

III. LADYBUG'S ARCHITECTURE & IMPLEMENTATION

Ladybug is comprised of three main components shown in Figure 2: (i) the MetaMorph desktop application for Mac/Windows (ii) the GitHub connector that facilitates communication between the Localization engine and GitHub, and (iii) the backend server that implements the bug localization engine.

A. The MetaMorph Bug Reproduction Tool

While LadyBug can be used in conjunction with automated tools that reproduce bugs on an Android device given a bug report [12], these tools are not perfect, and some developers may prefer to reproduce bugs manually. As such, LadyBug is packaged with MetaMorph, a Mac and Windows Desktop application that facilitates the recording of a reproduction trace from an Android device or emulator. MetaMorph has two main components, a *recorder* and a *replayer*. The *recorder*

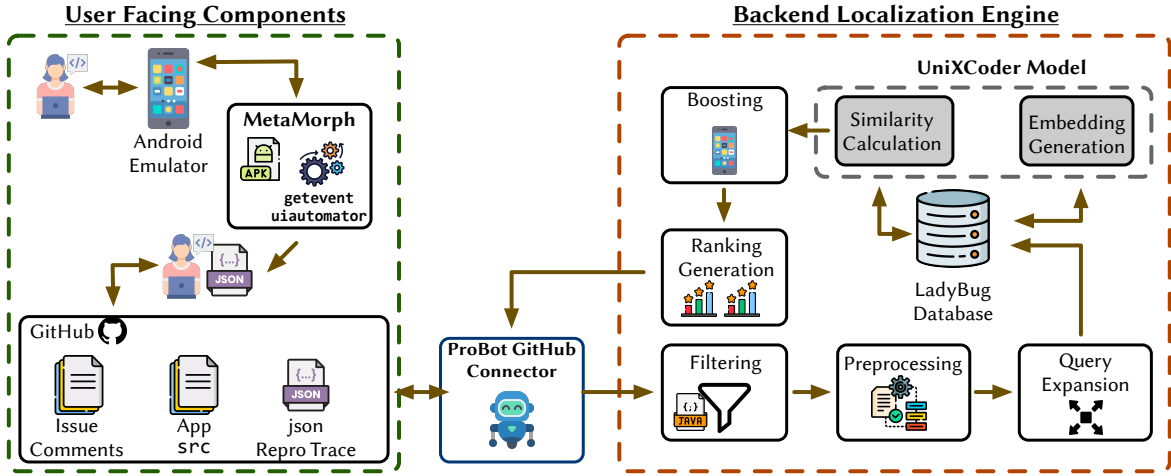


Fig. 2: Ladybug’s Architecture

makes use of the **getevent** utility, included as a part of the Android kernel, that is capable of capturing fine-grained touch information with millisecond-level accuracy and the Android **screencap** utility to record both a trace and video of a given bug. Then the **replayer** takes as input the **getevent** trace, and automatically replays it on the device, this time capturing fine-grained GUI information of each widget that the user interacted with to reproduce the bug using the Android **uiautomator** framework. The **replayer** exports a json file that contains all of the GUI interaction information needed by LadyBug.

B. ProBot GitHub Repository Connector

LadyBug’s Repository connector is built on the Probot framework, which offers a high-level abstraction for developing GitHub Apps using Node.js and the Octokit library. Probot listens for webhook events—such as repository installations and issue creations—and routes them through an internal handler, allowing the app to respond programmatically via GitHub’s REST API. LadyBug uses this infrastructure to authenticate as a GitHub App, monitor repository events, and interface directly with GitHub issues. The backend logic, written in JavaScript, handles routing, repository tracking, and progress updates.

When LadyBug is first installed on a repository, it records the repository identifier ("owner/repo") and the latest commit SHA, then creates a welcome issue to initiate setup. This SHA and the repository’s source code are stored in the LadyBug database in the backend server, and embeddings are computed using the UniXCoder model in preparation for retrieval of source code files pending an incoming issue. We chose to use the UniXCoder model as it was one of the better performing models reported in the original technical paper [14]. However, LadyBug is extensible to any text-embedding format, including from Large Language Models (LLMs). When a user opens a new GitHub issue to report a bug, LadyBug validates the request and posts an initial comment to confirm that the report has been received. If the report includes a reproduction trace, it is parsed and sent to the backend along with the issue content and repository metadata. Before proceeding, LadyBug verifies whether the commit SHA in the report matches the stored version. If it does not, the initialization phase is re-triggered to ensure the GitHub source code files and computed embeddings are current. The updated issue data is then compared against the stored

embeddings to generate a ranked list of likely buggy files. The final ranked output is posted as a comment to the GitHub issue.

C. Localization Engine

LadyBug’s localization engine is a Flask server fully written in Python that handles bug localization. It listens for two requests sent from Probot: initialization and report. Probot sends an ‘initialization’ request when LadyBug is added to a repository and a ‘report’ request whenever a GitHub Issue is created on the target repository.

Upon an ‘initialization’ request, LadyBug, using its Probot credentials, clones the repository’s default branch. LadyBug filters the repository to retain only Java source files—any files or directories not matching the .java extension are discarded. Each remaining Java files go through preprocessing for uniform representation and their respective embeddings are stored in LadyBug’s embedding database. LadyBug performs the following preprocessing steps: (i) Sanitization: Strips comments, imports, and special characters not relevant to semantics. (ii) Tokenization: Breaks the source text into syntactic tokens (identifiers, keywords, literals). (iii) Normalization: Removes language-specific stop tokens (e.g., punctuation, boilerplate keywords) and applies lemmatization to identifiers and literals. (iv) Segmentation: Splits long files into semantically coherent segments if they exceed 500 tokens.

When the localization engine receives a ‘report’ request, LadyBug will begin by extracting relevant GUI information from the attached execution.json file. Using regex, resource-ids of individual GUI components are stored and classified as **Screen Component terms**. Activity/Fragment names are stored and classified as **GUI Screen terms**. Screen Component terms are appended to the end of the bug report and then used for **Query Reformulation**. After this, a corpus of source code file embeddings is built using a brute-force approach: If the source code file contained any Screen Component Term, **Filtering** is performed to remove files not related to the relevant screens. Lastly, **Boosting** is performed by boosting source code files that contain GUI Screen terms to the top of the ranked list, preserving relative order. This configuration of Filtering, Boosting, and Query Reformulation were derived from the best performing configuration in LadyBug’s expanded evaluation in its technical paper [14].

IV. EVALUATION

We conducted an empirical study to understand the impact of using GUI interaction information on text-retrieval-based bug localization. We formulated the following research question:

RQ: *Does the use of GUI Interaction information improve text-retrieval-based bug localization for Android apps?*

A. Data

To answer RQ, we developed *Redwing*, an automated testing suite designed specifically to assess the effectiveness of LadyBug’s localization methods, and to ensure that we were able to replicate the results of the larger empirical study performed in LadyBug’s technical paper [14]. We used the same set of 80 reproducible bug reports from 39 different open-source Android applications used in LadyBug’s technical paper. Each bug report includes the description of the bug, metadata for the bug, and the recorded steps/scenarios for reproducing the bug. For a full description of the data collection methodology, please see the technical paper [14].

B. Methodology

Redwing leverages the core bug localization pipeline, adapting it to systematically process a predefined dataset using known buggy files as ground truth. By comparing the pipeline’s predicted outcomes against these established ground truths, Redwing calculates several widely recognized evaluation metrics, including (i) Hits@K: indicates the frequency with which the correct buggy file appears within the top K results. A higher value suggests that the system consistently narrows down bugs to a manageable subset of files. (ii) Mean Reciprocal Rank (MRR): Measures how quickly the first buggy file is identified. A higher score (closer to 1) indicates better performance, meaning fewer irrelevant files to review before locating the correct one. (iii) Mean Average Precision (MAP): Evaluates the overall ranking accuracy of all buggy files within the results. Higher MAP scores demonstrate the system’s precision in consistently placing the correct files near the top of the list. (iv) Effectiveness (E): Represents the rank position of the first identified buggy file, where lower ranks are preferable. This metric directly reflects the practical efficiency, indicating how many files developers must review before pinpointing the bug.

The Redwing testing suite integrates essential components from both the initialization and reporting phases of the main localization pipeline, resulting in a cohesive and streamlined evaluation process. This consolidation eliminates performance bottlenecks typically associated with external dependencies, such as API calls, data streaming, and manual user inputs, thus ensuring efficient and reproducible testing. As a command-line interface (CLI) tool, Redwing provides users flexibility and ease of use, allowing users to specify various parameters, including the dataset file path, localization type (with GUI data, without GUI data, or a comparative mode to calculate the relative improvement), and the number of asynchronous test iterations to execute. This modular and user-friendly design supports comprehensive and customizable experimentation.

TABLE I: Metrics for GUI-augmented vs. text-only (No GUI) data.

Augmentation	H@1	H@5	H@10	MAP	MRR	Effectiveness
UniXCoder-GUI	0.30	0.74	0.81	0.39	0.44	13.81
UniXCoder-No GUI	0.20	0.61	0.68	0.31	0.36	16.41

C. Results

In Table I presented above, it is evident that incorporating text-based retrieval techniques as an augmentation to the GUI-based data consistently enhances retrieval performance across all evaluated values of k in the Hits@k metric (H@K). Notably, our primary focus was on the Hits@10 metric, as this represents a manageable number of files for a developer to examine [7, 15]. By integrating text-based retrieval with GUI data, we observed a substantial improvement in the Hits@10 score, achieving a relative increase of 16.2%. These results suggest that the complementary nature of textual information in enriching the semantic understanding of GUI representations facilitates more accurate and relevant retrieval outcomes. Additionally, these results replicate the results for the same configuration of LadyBug in its corresponding technical paper, ensuring that our implementation of the localization in our LadyBug GitHub bot, matches that of the original algorithms developed in the technical paper [14].

In addition, augmenting the GUI data with additional retrieval signals has led to noticeable improvements in both the MAP and MRR metrics. Specifically, the rise in MAP suggests a greater overall precision across multiple queries. At the same time, the higher MRR values reflect a more consistent placement of the correct file near the top of the ranked list. In terms of effectiveness, we also see that with GUI data, the mean efficacy decreases. This suggests that using GUI data, we can more accurately rank files based on their relevance to the bug. These findings demonstrate that using GUI data contributes to a more effective retrieval process by reducing the prominence of irrelevant files at the top of the list.

V. RELATED WORK

Text-Retrieval-based Bug Localization (TRBL) techniques use text similarity between code and bug reports to identify potentially buggy code artifacts. A relevance score determines how likely code is to contain bugs. TRBL techniques fall into three main categories: classical, DL, and combined.

Classical IR techniques employ algorithms including the Unigram Model (UM), Cluster Based Document Model (CBDM), Vector Space Model (VSM), Latent Semantic Indexing (LSI), and Latent Dirichlet Allocation (LDA).

DL approaches utilize neural network architectures including Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Transformers, and hybrid models that combine these architectures.

Other approaches combine IR and DL techniques [16].

To improve bug localization accuracy by bridging the gap between bug reports and source code, researchers incorporate additional information sources including similar bug reports [17], code structure [18], version history [19], stack traces [15], and part-of-speech data [20].

LadyBug is novel compared to such prior work as we incorporate GUI interaction data from Android apps to bridge the lexical gap between bug reports and source code, thereby improving buggy file retrieval.

Prior work has employed various augmentation strategies to improve the ranking of buggy code artifacts in TRBL, including boosting, filtering, and query (re)formulation. Boosting increases the relevance scores of artifacts to rank buggy ones higher [21], while filtering removes irrelevant artifacts from the search space or initial results [22]. Query (re)formulation modifies queries to better capture relevant information through expansion (adding terms) [23], replacement (substituting the query) [24], and reduction (removing unhelpful terms) [25]. Building on this work, we assess the impact of incorporating GUI interaction data through boosting, filtering, query expansion, and replacement, and plan to explore its integration with query reduction in future work.

LadyBug employs boosting, filtering, and query expansion as augmentation techniques with a new source of data, namely, GUI interaction data. Our adaption of these techniques to GUI interaction data in Android apps is novel.

VI. FINAL REMARKS & FUTURE WORK

Ladybug is a tool that helps Android developers accurately localize bugs in their source code by combining text retrieval with boosting and filtering strategies utilizing GUI interaction data. By leveraging language model to encode both source files and user-reported bug descriptions, and by extracting resource-id and view-hierarchy terms from replayed interaction traces, Ladybug generates a ranked list of candidate files that more likely to be buggy than traditional techniques. In the future, we plan on expanding the amount of embedding models used to create source code and bug report embeddings. This would ensure that both code and natural-language inputs benefit from the latest advances in representation learning. We also intend to enrich Ladybug's user-facing output: (e.g., code fix suggestions, bug explanations). We will also enhance the design of LadyBug to accept the MetaMorph results at any time while an issue remains open.

ACKNOWLEDGEMENT

This research has been supported in part by the NSF CCF-2441355, CCF-2007246, and CCF-1955853. The authors also acknowledge funding support from Apple. Any opinions, findings, and conclusions expressed herein are the authors' and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] E. Murphy-Hill, T. Zimmermann, C. Bird, and N. Nagappan, "The design of bug fixes," in *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 2013, pp. 332–341.
- [2] A. Adnan, A. Saha, and O. Chaparro, "Sprint: An assistant for issue report management," in *Proceedings of the IEEE/ACM 22nd International Conference on Mining Software Repositories (MSR'25)*, 2025, pp. 86–90.
- [3] O. Chaparro, C. Bernal-Cárdenas, J. Lu, K. Moran, A. Marcus, M. Di Penta, D. Poshyanyk, and V. Ng, "Assessing the quality of the steps to reproduce in bug reports," in *Proceedings of the 27th ACM Joint Meeting on the Foundations of Software Eng.*, 2019, p. 86–96.
- [4] J. Mahmud, A. Saha, O. Chaparro, K. Moran, and A. Marcus, "Combining Language and App Ui Analysis for the Automated Assessment of Bug Reproduction Steps," in *Proceedings of the IEEE/ACM 33rd International Conference on Program Comprehension (ICPC'25)*, Apr. 2025, pp. 1–12.
- [5] S. Davies, M. Roper, and M. Wood, "Using bug report similarity to enhance bug localisation," in *2012 19th Working Conference on Reverse Engineering*. IEEE, 2012, pp. 125–134.
- [6] K. Moran, M. Linares-Vásquez, C. Bernal-Cárdenas, and D. Poshyanyk, "Auto-completing Bug Reports for Android Applications," in *Proceedings of the Joint Meeting on Foundations of Software Engineering (FSE'15)*, 2015, pp. 673–686.
- [7] M. M. Rahman and C. K. Roy, "Improving ir-based bug localization with context-aware query reformulation," in *Proceedings of the 2018 26th ACM Joint Meeting on the Foundations of Software Engineering*, 2018, p. 621–632.
- [8] K. C. Youm, J. Ahn, and E. Lee, "Improved bug localization based on code change histories and bug reports," *Information and Software Technology*, vol. 82, pp. 177–192, 2017.
- [9] B. Dit, M. Reville, M. Gethers, and D. Poshyanyk, "Feature location in source code: a taxonomy and survey," *Journal of software: Evolution and Process*, vol. 25, no. 1, pp. 53–95, 2013.
- [10] D. Guo, S. Lu, N. Duan, Y. Wang, M. Zhou, and J. Yin, "Unixcoder: Unified cross-modal pre-training for code representation," *arXiv preprint arXiv:2203.03850*, 2022.
- [11] "Probot github bot framework," <https://probot.github.io/>, 2025.
- [12] D. Wang, Y. Zhao, S. Feng, Z. Zhang, W. G. J. Halfond, C. Chen, X. Sun, J. Shi, and T. Yu, "Feedback-driven automated whole bug report reproduction for android apps," in *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 1048–1060.
- [13] "Ladybug github repo," <https://github.com/LadyBugML/ladybug>, 2025.
- [14] J. Mahmud, N. De Silva, S. A. Khan, S. H. Mostafavi, S. M. H. Mansur, O. Chaparro, A. A. Marcus, and K. Moran, "On using gui interaction data to improve text retrieval-based bug localization," in *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*.
- [15] M. Wen, R. Wu, and S. Cheung, "Locus: Locating bugs from software changes," in *Proceedings of the 31st International Conference on Automated Software Engineering (ASE'16)*, 2016, pp. 262–273.
- [16] A. N. Lam, A. T. Nguyen, H. A. Nguyen, and T. N. Nguyen, "Bug localization with combination of deep learning and information retrieval," in *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*, 2017, pp. 218–229.
- [17] M. Rath, D. Lo, and P. Mäder, "Analyzing requirements and traceability information to improve bug localization," in *Proceedings of the 15th Intl. Conference on Mining Software Repositories*, 2018, pp. 442–453.
- [18] A. Takahashi, N. Sae-Lim, S. Hayashi, and M. Saeki, "A preliminary study on using code smells to improve bug localization," in *Proceedings of the 26th Conference on Program Comprehension*, 2018, pp. 324–327.
- [19] T. Zhang, W. Hu, X. Luo, and X. Ma, "A commit messages-based bug localization for android applications," *Intl. Journal of Software Engineering and Knowledge Eng.*, vol. 29, no. 04, pp. 457–487, 2019.
- [20] Y. Zhou, Y. Tong, T. Chen, and J. Han, "Augmenting bug localization with part-of-speech and invocation," *Intl. Journal of Software Engineering and Knowledge Engineering*, vol. 27, no. 06, pp. 925–949, 2017.
- [21] Y. Lou, Q. Zhu, J. Dong, X. Li, Z. Sun, D. Hao, L. Zhang, and L. Zhang, "Boosting coverage-based fault localization via graph-based representation learning," in *Proceedings of the 29th ACM Joint Meeting on the Foundations of Software Engineering*, 2021, pp. 664–676.
- [22] D. Liu, A. Marcus, D. Poshyanyk, and V. Rajlich, "Feature location via information retrieval based filtering of a single scenario execution trace," in *Proceedings of the 22nd IEEE/ACM international conference on Automated software engineering*, 2007, pp. 234–243.
- [23] M. Kim, Y. Kim, and E. Lee, "A novel automatic query expansion with word embedding for ir-based bug localization," in *2021 IEEE 32nd International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 2021, pp. 276–287.
- [24] J. Guo, M. Gibiec, and J. Cleland-Huang, "Tackling the term-mismatch problem in automated trace retrieval," *Empirical Software Engineering*, vol. 22, pp. 1103–1142, 2017.
- [25] J. M. Florez, O. Chaparro, C. Treude, and A. Marcus, "Combining query reduction and expansion for text-retrieval-based bug localization," in *2021 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2021, pp. 166–176.