GiveMeLabeledIssues: An Open Source Issue Recommendation System

Joseph Vargovich, Fabio Santos, Jacob Penney, Marco A. Gerosa, Igor Steinmacher

School of Informatics, Computing, and Cyber Systems

Northern Arizona University

Flagstaff, United States

joseph_vargovich@nau.edu, fabio_santos@nau.edu, jacob_penney@nau.edu, Marco.Gerosa@nau.edu, Igor.Steinmacher@nau.edu

Abstract—Developers often struggle to navigate an Open Source Software (OSS) project's issue-tracking system and find a suitable task. Proper issue labeling can aid task selection, but current tools are limited to classifying the issues according to their type (e.g., bug, question, good first issue, feature, etc.). In contrast, this paper presents a tool (*GiveMeLabeledIssues*) that mines project repositories and labels issues based on the skills required to solve them. We leverage the domain of the APIs involved in the solution (e.g., User Interface (UI), Test, Databases (DB), etc.) as a proxy for the required skills. *GiveMeLabeledIssues* facilitates matching developers' skills to tasks, reducing the burden on project maintainers. The tool obtained a precision of 83.9% when predicting the API domains involved in the issues. The replication package contains instructions on executing the tool and including new projects. A demo video is available at https://www.youtube.com/watch?v=ic2quUue7i8

Index Terms—Open Source Software, Machine Learning, Label, Tag, Task, Issue Tracker

I. INTRODUCTION

Contributors struggle to find issues to work on in Open Source Software (OSS) projects due to difficulty determining the skills required to work on an issue [1]. Labeling issues manually can help [2], but manual work increases maintainers' effort, and many projects end up not labeling their issues [3].

To facilitate issue labeling, several studies proposed mining software repositories to tag the issues with labels such as bug/non-bug and good-first-issue (to help newcomers find a suitable task) [4, 5, 6, 7]. However, newcomers to a project have a variety of skill levels and these approaches do not indicate the skills needed to work on the tasks.

APIs encapsulate modules with specific functionality. If we can predict APIs used to solve an issue, we can guide new contributors on what to contribute. However, since projects include thousands of different APIs, adding each API as a label would harm the user experience. To keep the number of labels manageable, we classified the APIs into domains for labeling [2, 8]

In our study, API domains represent categories of APIs such as "UI," "DB," "Test," etc. In our previous work, we used 31 distinct API domains [8]. However, we acknowledge that the set of possible API-domain labels varies by project, as each project requires different skills and expertise. In another work, Santos et al. [9] have evaluated the API domain labels in an empirical experiment with developers. The study showed that the labels enabled contributors to find OSS tasks that matched their expertise more efficiently. Santos et al. study showed promising results when classifying issues by API-Domain, with precision, recall, and F-measure scores of 0.755, 0.747, and 0.751, respectively.

Following this idea, we implemented *GiveMeLabeledIssues* to classify issues for potential contributors. *GiveMeLabeledIssues* is a web tool that indicates API-domain labels for open issues in registered projects. Currently, *GiveMeLabeledIssues* works with three open-source projects: JabRef, PowerToys, and Audacity. The tool enables users to select a project, input their skill set (in terms of our categories), and receive a list of open issues that would require those skills.

We trained the tool with the title and body text of closed issues and the APIs utilized within code that solve the issue, as gathered from pull requests. We evaluated the tool using closed issues as the ground truth and found an average of 83.8% precision and 78.5% recall when training and testing models for individual projects.

II. GIVEMELABELEDISSUES ARCHITECTURE

GiveMeLabeledIssues leverages prediction models trained with closed issues linked with merged pull requests. On top of these models, we built a platform that receives user input and queries the open issues based on the users' skills and desired projects. In the following, we detail the process of building the model, the issue classification process, and the user interface that enables users to get the recommendation. *GiveMeLabeledIssues* is structured in two layers: the frontend web interface[1] and the backend REST API[2].

A. Model training

In the current implementation, a model must be built for each project using data collected from project issues linked with merged pull requests. The tool maps the issue text data to the APIs used in the source code that solved the issues (Figure 1).

---

[1]https://github.com/JoeyV55/GiveMeLabeledIssuesUI
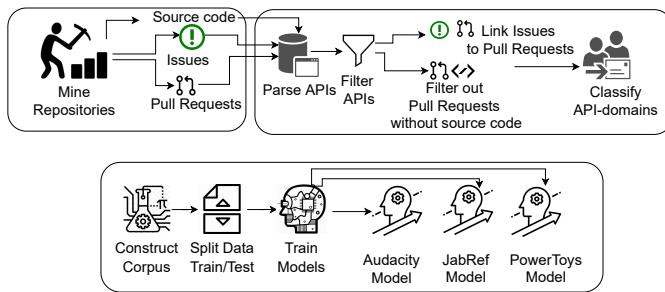[2]https://github.com/JoeyV55/GiveMeLabeledIssuesAPI

Fig. 1. The Process of Training a Model

1) Mining repositories: We collected 18,482 issues and 3,129 pull requests (PRs) from JabRef, PowerToys, and Audacity projects up to November 2021. We used the GitHub REST API v3 to collect the title, body, comments, closure date, name of the files changed in the PR, and commit messages.

2) APIs parsing: We built a parser to process all source files from the projects to identify the APIs used in the source code affected by each pull request. In total, we found 3,686 different APIs in 3,108 source files. The parser looked for specific commands, i.e., import (Java), using (C#), and include (C++). The parser identified all classes, including the complete namespace from each import/using/include statement.

3) Dataset construction: We kept only the data from issues linked with merged and closed pull requests since we needed to map issue data to the APIs that are used in the source code files changed to close the issue. To find the links between pull requests and issues, we searched for the symbol `#issue_number` in the pull request title and body and checked the URL associated with each link. We also filtered out issues linked to pull requests without at least one source code file (e.g., those associated only with documentation files) since they do not provide the model with content related to any API.

4) API categorization: We use the API domain categories defined by Santos et al. [8]. These 31 categories were defined by experts to encompass APIs from several projects (e.g., UI, IO, Cloud, DB, etc.—see our replication package[3]).

5) Corpus construction: We used the issue title and body as our corpus to train our model since they performed well in our previous analysis [9]. Similar to other studies [10, 11], we applied TF-IDF as a technique for quantifying word importance in documents by assigning a weight to each word following the same process described in the previous work [9]. TF-IDF returns a vector whose length is the number of terms used to calculate the scores. Before calculating the scores, we convert each word to lowercase and removed URLs, source code, numbers, and punctuation. After that, we remove templates and stop-words and stem the words. These TF-IDF scores are then passed to the Random Forest classifier (RF) as features for prediction. RF was chosen since it obtained the best results in previous work [9]. The ground truth has a

binary value (0 or 1) for each API domain, indicating whether the corresponding domain is present in the issue solution.

We also offer the option of using a BERT model in *GiveMeLabeledIssues*. We created two separate CSV files to train BERT: an input binary with expert API-domain labels paired with the issue corpus and a list of the possible labels for the specific project. BERT directly labels the issue using the corpus text and lists possible labels without needing an additional classifier (such as Random Forest).

6) Building the model: The BERT model was built using the Python package fast-bert[4], which builds on the Transformers[5] library for Pytorch. Before training the model, the optimal learning rate was computed using a LAMB optimizer [12]. Finally, the model was fit over 11 epochs and validated every epoch. The BERT model was trained on an NVIDIA Tesla V100 GPU with an Intel(R) Xeon(R) Gold 6132 CPU within a computing cluster.

TF-IDF and BERT models were trained and validated for every fold in a ShuffleSplit 10-fold cross-validation. Once trained, the models were hosted on the backend. The replication package contains instructions on registering a new project by running the model training pipeline that feeds the demo tool. The models can then output predictions quickly without continually retraining with each request.

For the RandomForestClassifier (TF-IDF), the best classifier, we kept the following parameters: criterion = 'entropy,' max depth = 50, min samples leaf = 1, min samples split=3, and n estimators = 50.

B. Issue Classification Process

GiveMeLabeledIssues classifies currently open issues for each registered project. The tool combines the title and body text and sends it to the classifier. The classifier then determines which domain labels are relevant to the gathered issues based on the inputted issue text. The labeled issues are stored in an SQLite database for future requests, recording the issue number, title, body, and domain labels outputted by the classifier.

The open issues for all projects registered with *GiveMeLabeledIssues* are reclassified daily to ensure that the database is up to date as issues are opened and closed. Figure 2 outlines the daily classification procedure.


Fig. 2. The Process of Classifying and Storing Issues

C. User Interface

GiveMeLabeledIssues outputs the labeled issues to the User Interface. The user interface is implemented using the Angular web framework. To use the tool, users provide the project

name and select API-domain labels they are interested in. This information is sent to the backend REST endpoint via a GET request. The backend processes the request, recommending a set of relevant issues for the user.

The backend REST API is implemented using the Django Rest Framework. It houses the trained TF-IDF and BERT text classification models and provides an interface to the labeled issues. When receiving the request, the backend queries the selected project for issues that match the user's skills. Once the query is completed, the backend returns the labeled issues to the user interface. Each labeled issue includes a link to the open issue page on GitHub and the issue's title, number, and applicable labels. The querying process is shown in Figure 3.


Fig. 3. The Process of Outputting Issues

Figure 4 shows JabRef selected as the project and "Utility," "Databases," "User Interface," and "Application" as the API domains provided by the user. Figure 5 shows the results of this query, which displays all JabRef open issues that match those labels.
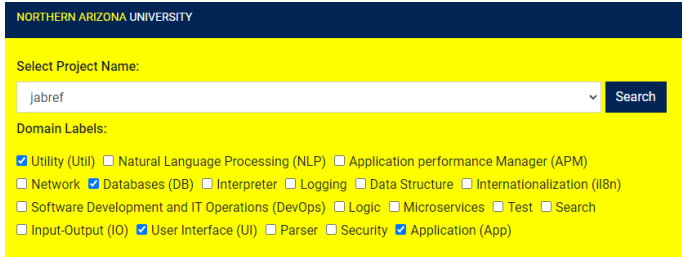

Fig. 4. Selection of a project and API domains

III. EVALUATION

We have evaluated the performance of the models used to output API-domain labels using a dataset comprised of 18,482 issues, 3,129 PRs, 3,108 source code files, and 3,686 distinct APIs, chosen from active projects and diverse domains: Audacity (C++), PowerToys (C#), and JabRef (Java).[6] Audacity is an audio editor, PowerToys is a set of utilities for Windows, and JabRef is an open-source bibliography manager. Table I shows the results with Precision, Recall, F-Measure, and Hamming Loss values. We trained models to predict API-domain labels using individual issue datasets from each project and a single dataset that combined the data from all the projects.

As shown in Table I, TF-IDF overcame BERT both in the per-project analysis and for the complete dataset. The

[6]The model training replication package is available at https://zenodo.org/record/7726323#.ZA5oy-zMIeY

Fig. 5. Labeled Issues Outputted for JabRef with the Utility, Databases, User Interface, and Application skills Selected

TABLE I
OVERALL METRICS FROM MODELS -AVERAGES. RF-RANDOM FOREST*
HLA - HAMMING LOSS**

Model averages	One/Multi projects	Precision	Recall	F-meas -ure	Hla**
RF* TF-IDF	O	**0.839**	0.799	0.817	**0.113**
BERT	O	0.595	0.559	0.568	0.269
RF* TF-IDF	M	0.659	0.785	0.573	0.153
BERT	M	0.593	0.725	0.511	0.219
Izadi et al. [13]	M	0.796	0.776	0.766	NA
Kallis et al. [14]	M	0.832	**0.826**	**0.826**	NA
Santos et al. [9]	O	0.755	0.747	0.751	0.142

difference was quite large when using single projects. The results were closer when we used the combined dataset that included data from all projects (3,736 linked issues and pull requests). We hypothesize that the sample size influenced the classifiers' performance. This aligns with previous research on issue labeling that showed that BERT performed better than other language models for datasets larger than 5,000 issues [15]. TF-IDF performs very well when the dataset is from a single project because the vocabulary used in the project is very contextual, and the frequency of terms can identify different aspects of each issue. When we include the dataset from all the projects, the performance of TF-IDF drops as the context is not unique. These results outperformed the results from the API-domain labels case study conducted by Santos et

TABLE II
OVERALL METRICS FROM MODELS - BY PROJECTS. RF-RANDOM
FOREST* HLA - HAMMING LOSS**

Model by project	Project	Precision	Recall	F-meas -ure	Hla**
RF* TF-IDF	Audacity	**0.872**	**0.839**	**0.854**	0.103
RF* TF-IDF	JabRef	0.806	0.782	0.793	0.143
RF* TF-IDF	PowerToys	0.84	0.776	0.805	**0.094**
BERT	Audacity	0.382	0.511	0.434	0.42
BERT	JabRef	0.791	0.606	0.686	0.192
BERT	PowerToys	0.619	0.643	0.626	0.187

al. [9]. The project metrics (Table II) varied less than 6% (e.g., the recall: 0.839 (Audacity) - 0.776 (PowerToys). Audacity had the best scores for all metrics except Hamming Loss.

IV. Related Work

The existing literature explores strategies to help newcomers find tasks within OSS projects. Steinmacher et al. [16] proposed a portal to help newcomers find relevant documentation. Despite pointing the contributor to existing resources, a newcomer may have difficulty relating the documentation to the skills required to solve a task.

There are also several approaches designed to label issues automatically. However, most of them only try to distinguish bug reports from non-bug reports [4, 17, 18]. Zhou et al. [17] built a Naive Bayes (NB) classifier to predict the priorities of bug reports. Xia et al. [18] tagged questions using preexisting labels. Other work [6, 7] is also restricted to existing labels while others [13, 14] proposed other labels. Kallis et al. [14], for instance, employed the textual description of the issues to classify the issues into types.

Labeling also attracted the attention of software requirement researchers. Pérez-Verdejo et al. [19] and Quba et al. [20] proposed to categorize documents with functional and non-functional requirements labels to solve software engineers' problems. Non-functional requirements labels included 12 general domains like "Performance" or "Availability" [20]. Such approaches use higher-level labels that would not guide contributors to choose appropriate tasks given their skills. "Availability," for instance, may be related to a database, network, or a cloud module and can be challenging for a newcomer to realize where to start due to the extent of modules to analyze to find the root cause of a bug.

APIs are also often investigated in software engineering. Recent work focuses on providing API recommendation [21, 22], giving API tips to developers [23], defining the skill space for APIs, developers, and projects [24] or identifying experts in libraries [25]. The APIs usually implement a concept (database access, thread management, etc.). Knowing the API involved in a potential solution allows newcomers to pick an issue that matches their skillset. Therefore, unlike the presented related studies, our tool labels issues based on API domains [9].

V. Conclusion and Future Work

GiveMeLabeledIssues provides OSS developers with issues that can potentially match their skill sets or interests. Such a tool may facilitate the onboarding of newcomers and alleviate the workload of project maintainers.

Future work can explore different domain labels, such as those offered by accredited standards (i.e., ACM/IEEE/ISO/SWEBOK). As a future step to evaluate the tool's impact, we will conduct a study to receive feedback from contributors and assess how the tool influences their choices by means of controlled experiments. Future work can also incorporate the use of social features [8] and integrate the tool into GitHub Actions, or bots [26].

References

[1] I. Steinmacher, M. A. G. Silva, M. A. Gerosa, and D. F. Redmiles, "A systematic literature review on the barriers faced by newcomers to open source software projects," *Information and Software Technology*, vol. 59, 2015.

[2] F. Santos, B. Trinkenreich, J. F. Pimentel, I. Wiese, I. Steinmacher, A. Sarma, and M. A. Gerosa, "How to choose a task? mismatches in perspectives of newcomers and existing contributors," in *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, 2022.

[3] A. Barcomb, K.-J. Stol, B. Fitzgerald, and D. Riehle, "Managing episodic volunteers in free/libre/open source software communities," *IEEE Transactions on Software Engineering*, 2020.

[4] N. Pingclasai, H. Hata, and K.-i. Matsumoto, "Classifying bug reports to bugs and other requests using topic modeling," in *2013 20Th asia-pacific software engineering conference (APSEC)*, vol. 2.   IEEE, 2013, pp. 13–18.

[5] Y. Zhu, M. Pan, Y. Pei, and T. Zhang, "A bug or a suggestion? an automatic way to label issues," *arXiv preprint arXiv:1909.00934*, 2019.

[6] F. El Zanaty, C. Rezk, S. Lijbrink, W. van Bergen, M. Côté, and S. McIntosh, "Automatic recovery of missing issue type labels," *IEEE Software*, 2020.

[7] Q. Perez, P.-A. Jean, C. Urtado, and S. Vauttier, "Bug or not bug? that is the question," in *2021 IEEE/ACM 29th International Conference on Program Comprehension (ICPC)*.   IEEE, 2021, pp. 47–58.

[8] F. Santos, J. Penney, J. F. Pimentel, I. Wiese, B. Trinkenreich, I. Steinmacher, and M. A. Gerosa, "Tell me who are you talking to and I will tell you what issues need your skills," in *2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR)*, 2023.

[9] F. Santos, I. Wiese, B. Trinkenreich, I. Steinmacher, A. Sarma, and M. A. Gerosa, "Can I solve it? identifying apis required to complete oss tasks," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*.   IEEE, 2021.

[10] D. Behl, S. Handa, and A. Arora, "A bug mining tool to identify and analyze security bugs using naive bayes and tf-idf," in *2014 International Conference on Reliability Optimization and Information Technology (ICROIT)*.   IEEE, 2014, pp. 294–299.

[11] S. L. Vadlamani and O. Baysal, "Studying software developer expertise and contributions in Stack Overflow and GitHub," in *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*.   IEEE, 2020, pp. 312–323.

[12] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bho-janapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh, "Large batch optimization for deep learning: Training bert in 76 minutes," in *International Conference on Learning Representations*, 2020.

[13] M. Izadi, K. Akbari, and A. Heydarnoori, "Predicting the objective and priority of issue reports in software repositories," *Empirical Software Engineering*, vol. 27, no. 2, pp. 1–37, 2022.

[14] R. Kallis, A. Di Sorbo, G. Canfora, and S. Panichella, "Ticket tagger: Machine learning driven issue classification," in *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2019, pp. 406–409.

[15] J. Wang, X. Zhang, and L. Chen, "How well do pre-trained contextual language representations recommend labels for github issues?" *Knowledge-Based Systems*, vol. 232, p. 107476, 2021.

[16] I. Steinmacher, T. U. Conte, C. Treude, and M. A. Gerosa, "Overcoming open source project entry barriers with a portal for newcomers," in *Proceedings of the 38th International Conference on Software Engineering*, 2016, pp. 273–284.

[17] Y. Zhou, Y. Tong, R. Gu, and H. Gall, "Combining text mining and data mining for bug report classification," *Journal of Software: Evolution and Process*, vol. 28, no. 3, pp. 150–176, 2016.

[18] X. Xia, D. Lo, X. Wang, and B. Zhou, "Tag recommendation in software information sites," in *Mining Software Repositories*. USA: IEEE, 2013.

[19] J. M. Pérez-Verdejo, Á. J. Sánchez-García, J. O. Ocharán-Hernández, E. Mezura-Montes, and K. Cortes-Verdin, "Requirements and github issues: An automated approach for quality requirements classification," *Programming and Computer Software*, vol. 47, 2021.

[20] G. Y. Quba, H. Al Qaisi, A. Althunibat, and S. AlZu'bi, "Software requirements classification using machine learning algorithm's," in *2021 International Conference on Information Technology (ICIT)*. IEEE, 2021.

[21] H. Zhong and H. Mei, "An empirical study on api usages," *IEEE Transactions on Software Engineering*, vol. 45, no. 4, pp. 319–334, 2019.

[22] Y. Zhou, C. Wang, X. Yan, T. Chen, S. Panichella, and H. C. Gall, "Automatic detection and repair recommendation of directive defects in java api documentation," *IEEE Transactions on Software Engineering*, pp. 1–1, 2018.

[23] S. Wang, N. Phan, Y. Wang, and Y. Zhao, "Extracting api tips from developer question and answer websites," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*, 2019, pp. 321–332.

[24] T. Dey, A. Karnauch, and A. Mockus, "Representation of developer expertise in open source software," *arXiv preprint arXiv:2005.10176*, 2020.

[25] J. E. Montandon, L. Lourdes Silva, and M. T. Valente, "Identifying experts in software libraries and frameworks among github users," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*, 2019, pp. 276–287.

[26] T. Kinsman, M. Wessel, M. A. Gerosa, and C. Treude, "How do software developers use github actions to automate their workflows?" in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 2021, pp. 420–431.