

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2023

What do users ask in open-source AI repositories? An empirical study of GitHub issues

Zhou YANG

Chenyu WANG

Jieke SHI

Thong HOANG

Pavneet Singh KOCHHAR

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), [Databases and Information Systems Commons](#), and the [Software Engineering Commons](#)

Citation

YANG, Zhou; WANG, Chenyu; SHI, Jieke; HOANG, Thong; KOCHHAR, Pavneet Singh; LU, Qinghua; XING, Zhenchang; and David LO. What do users ask in open-source AI repositories? An empirical study of GitHub issues. (2023). *Proceedings of the 20th IEEE/ACM International Conference on Mining Software Repositories, Melbourne, Australia 2023 May 15-16*. 79-91.

Available at: https://ink.library.smu.edu.sg/sis_research/8571

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Author

Zhou YANG, Chenyu WANG, Jieke SHI, Thong HOANG, Pavneet Singh KOCHHAR, Qinghua LU, Zhenchang XING, and David LO

What Do Users Ask in Open-Source AI Repositories? An Empirical Study of GitHub Issues

Zhou Yang*, Chenyu Wang*, Jieke Shi*, Thong Hoang[†], Pavneet Kochhar[‡], Qinghua Lu[†], Zhenchang Xing^{†§}, David Lo*

*Singapore Management University, Singapore

[†]Data61, Eveleigh, CSIRO, Sydney, Australia

[‡]Microsoft Research, Vancouver, Canada

[§]Australian National University, Canberra, Australia

{zyang, chenyuwang, jkshi, davidlo}@smu.edu.sg

{James.Hoang, Qinghua.Lu, Zhenchang.Xing}@data61.csiro.au

Pavneet.Kochhar@microsoft.com

Abstract—Artificial Intelligence (AI) systems, which benefit from the availability of large-scale datasets and increasing computational power, have become effective solutions to various critical tasks, such as natural language understanding, speech recognition, and image processing. The advancement of these AI systems is inseparable from open-source software (OSS). Specifically, many benchmarks, implementations, and frameworks for constructing AI systems are made open source and accessible to the public, allowing researchers and practitioners to reproduce the reported results and broaden the application of AI systems. The development of AI systems follows a data-driven paradigm and is sensitive to hyperparameter settings and data separation. Developers may encounter unique problems when employing open-source AI repositories.

This paper presents an empirical study that investigates the issues in the repositories of open-source AI repositories to assist developers in understanding problems during the process of employing AI systems. We collect 576 repositories from the `PapersWithCode` platform. Among these repositories, we find 24,953 issues by utilizing GitHub REST APIs. Our empirical study includes three phases. First, we manually analyze these issues to categorize the problems that developers are likely to encounter in open-source AI repositories. Specifically, we provide a taxonomy of 13 categories related to AI systems. The two most common issues are runtime errors (23.18%) and unclear instructions (19.53%). Second, we see that 67.5% of issues are closed. We also find that half of these issues resolve within four days. Moreover, issue management features, e.g., label and assign, are not widely adopted in open-source AI repositories. In particular, only 7.81% and 5.9% of repositories label issues and assign these issues to assignees, respectively. Finally, we empirically show that employing GitHub issue management features and writing issues with detailed descriptions facilitate the resolution of issues. Based on our findings, we make recommendations for developers to help better manage the issues of open-source AI repositories and improve their quality.

Index Terms—Open Source Software, Mining Software Repositories, Best Development Practices, AI Repositories

I. INTRODUCTION

Artificial Intelligence (AI) has revolutionized many critical tasks in our modern lives, such as natural language understanding [1], [2], image processing [3], speech recognition [4], and code understanding [5], [6]. The success of these AI techniques is heavily dependent on the open-source culture.

Popular open-source repositories like `TensorFlow`¹ and `PyTorch`² have greatly facilitated the development of AI applications. Additionally, research papers published at flagship AI conferences often release their replication packages and models, making it easier for researchers and practitioners to replicate results, run models on other tasks, or conduct further research. Furthermore, open-source benchmarks such as `CodeXGLUE`³ and `ImageNet`⁴ have played a crucial role in the advancement of AI by providing a common platform for researchers to evaluate their models and compare their results.

Despite the benefits of open-source AI repositories, developers often encounter challenges when utilizing them to build AI applications. A recent article [7] points out that many researchers have complained about a ‘replication crisis’ in AI repositories. This issue can be attributed to various factors, such as incomplete instructions for configuring development environments [8] and a lack of information about the datasets used to train AI models. Additionally, buggy code can also be a significant obstacle for developers. These problems can lead to frustration and hinder the use of open-source AI repositories.

To address the problems above, developers should communicate with repository owners and other developers to discuss and address these problems together. There are several channels for such communication. For example, developers can post questions in `StackOverflow`⁵ (SO). However, SO is unsuitable for finding answers to questions specific to a repository as there may be fewer experts on this topic. Another channel is email. Researchers or the repositories’ owners usually publish their contact information, e.g., email addresses and affiliations, in their papers or the descriptions of their repositories, respectively. The developers can also inquire about questions via email; however, this communication is often private and inaccessible for analysis. The issue trackers,

¹<https://www.tensorflow.org/>

²<https://pytorch.org/>

³<https://github.com/microsoft/CodeXGLUE>

⁴<https://www.image-net.org/>

⁵<https://stackoverflow.com/>

such as GitHub Issues⁶ or JIRA,⁷ can be appropriate resources to analyze the problems faced by the AI repository users.

We obtain a list of papers published between 2013 and 2022 from ten top-tier AI conferences and retrieve their replication packages from the PapersWithCode⁸ platform. After cleaning the data (e.g., removing the repositories that are not provided by the authors of papers), we identify 576 open-source software (OSS) AI repositories hosted on GitHub and 24,953 GitHub issues using the GitHub REST APIs. We then conduct open card sorting to develop a taxonomy of the issues in open-source AI repositories, categorizing them into 13 categories. Our findings reveal that the most commonly reported issue is `Runtime Error`. The second-largest group of issues falls under `Unclear Instructions`, primarily caused by inadequate documentation.

Our study examines how developers manage and address issues in their open-source AI repositories. The results indicate that 67.53% of issues are resolved, with 50% of issues closed within four days. Additionally, we found a positive correlation between the number of issues in a repository and the average time to address the issues (i.e., the time between the open and close date of issues), as well as a negative correlation between the number of issues and the closed issue rate (i.e., the ratio of closed issues to all the issues in a repository). We also discovered that developers often neglect the features provided by GitHub to manage issues, with only 7.81% of repositories using labels to categorize issues and only 5.90% of repositories assigning issues to specific individuals. Further analysis of the relationship between various features and the closure of issues in open-source AI repositories reveals that an issue’s label(s) and assignee(s) have a statistically significant impact on the closure of issues. These findings suggest that repository maintainers should actively utilize GitHub issue management tools to effectively manage issues in their repositories. Additionally, issues with longer descriptions and code blocks are more likely to be closed, indicating that issue raisers should provide more details, especially code blocks, to convey information clearly.

We make the following contributions in this paper:

- We conduct an empirical study to systematically investigate the issues in open-source AI repositories hosted on GitHub. The study contributes to the community with a dataset of 576 AI repositories that are manually confirmed to be the official implementations of papers from top-tier AI conferences and 24,953 issues in these repositories.
- We categorize the issues in open-source AI repositories to help developers better understand the problems users encounter when using these systems.
- Our empirical study shows that repository maintainers and developers should follow some good practices to help better address the issues in open-source AI repositories. We encourage the maintainers to actively adopt issue management features, e.g., the labeling and assigning

features provided by GitHub. We also suggest that the developers provide detailed information, especially code blocks, and format the issues properly.

The rest of this paper is organized as follows. Section II explains the background of open-source AI repositories and their issues. In Section III, we describe the process of collecting and cleaning datasets. Section IV reports research questions and experiment results. We discuss the difference between academic and industry repositories, as well as the threats to validity in Section V. Section VI reports the works relevant to our work. We conclude the paper and present future work in Section VII.

II. BACKGROUND

A. Open-Source AI Repositories

The continued success of open-source software (OSS) in the field of AI, as exemplified by projects such as TensorFlow, has been widely acknowledged by both practitioners and researchers. The critical role that OSS plays in the advancement and development of AI techniques has been emphasized in various studies [9], [10]. In particular, the following five AI-related components should be made open source: design, implementation, frameworks, benchmarks, and models [10].

Artificial intelligence (AI) papers and technical reports that describe the design of algorithms and models are often made available to the public in the form of open-source materials. For example, authors are encouraged to publish their papers as open-access or release preprint versions on platforms such as arXiv. Additionally, the review process for many AI conferences is conducted publicly on platforms like OpenReview. Furthermore, the implementation of AI systems is commonly made open source as well. Developers release their code on platforms such as GitHub and BitBucket, enabling researchers and practitioners to replicate results and conduct further extensions. The key frameworks that support the implementation of AI software, such as TensorFlow and PyTorch, are also made open source to facilitate the development of various AI models. There are emerging open-source frameworks for specific tasks, such as Carla for autonomous driving [11] and OpenAI Gym for reinforcement learning [12].

Open-source public benchmark platforms, e.g., CodaLab⁹ and Kaggle,¹⁰ provide developers with opportunities to participate in competitions and share their AI solutions. Platforms also offer large pre-trained AI models that can be fine-tuned for downstream tasks. For instance, HuggingFace¹¹ is a platform that shares a variety of Transformer-based models and collections of datasets for various tasks, including natural language processing, audio processing, and speech recognition. Similarly, ModelZoo¹² is another platform that shares pre-trained AI models. Giant tech companies like IBM and

⁶<https://github.com/features/issues>

⁷<https://www.atlassian.com/jira>

⁸<https://paperswithcode.com/>

⁹<https://codalab.org/>

¹⁰<https://www.kaggle.com/>

¹¹<https://huggingface.co/>

¹²<https://modelzoo.co/>

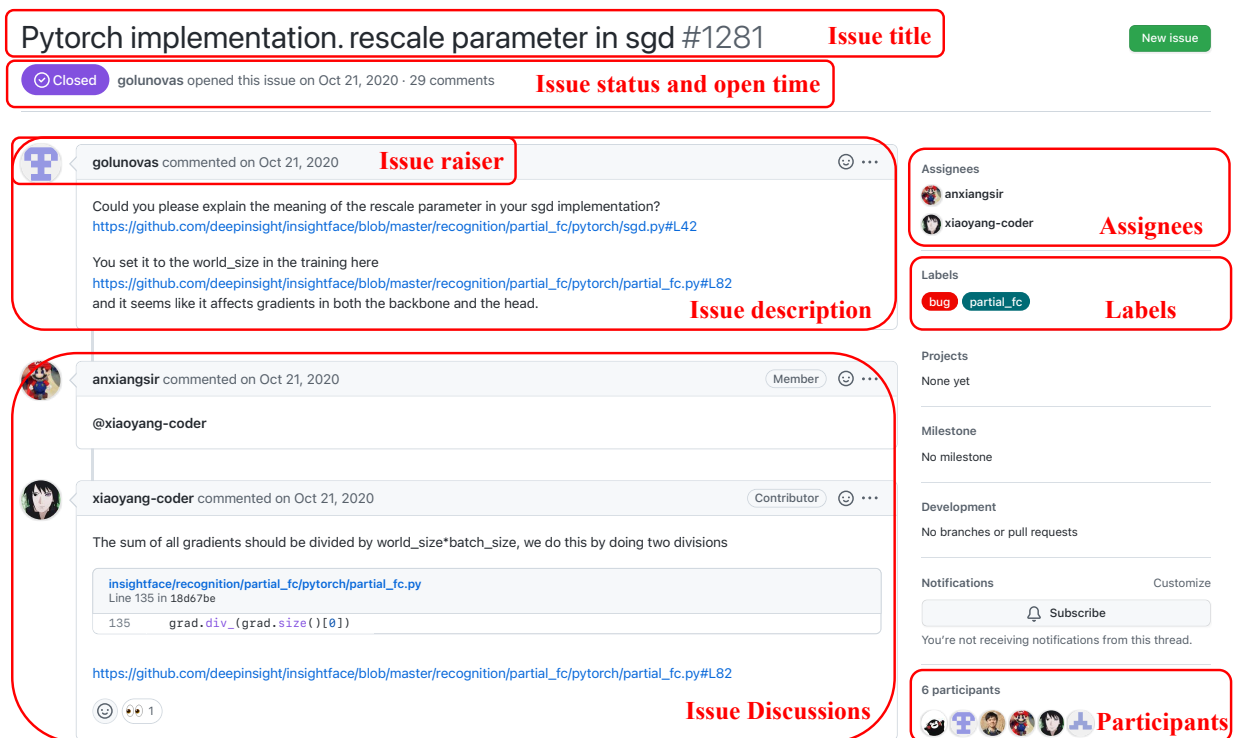


Figure 1. An example of an issue in an AI repository hosted on GitHub.

Google also actively participate in the open-source community by releasing their AI-related projects [13]–[16]. Among these open-source platforms, GitHub is one of the most popular and hosts numerous AI repositories [17], [18]. Therefore, in this study, we conduct an analysis of open-source AI repositories on GitHub to gain insights into their usage and popularity.

B. Tracking Issues in GitHub

The software system is not always reliable [7], [8]. Developers of open-source software (OSS) repositories may encounter difficulties when employing these repositories. Introduced in 2009 for the first time, GitHub Issues [19] is a built-in issue-tracking system integrated into GitHub that allows developers to raise questions, request features, etc.

Figure 1 illustrates an example of an issue in an open-source AI repository hosted on GitHub.¹³ The person raising the issue, referred to as the *issue raiser*, provides a title and a description of the issue. The issue can be in either an *open* or *closed* status, indicating whether it has been resolved or not. Any GitHub users who can access the repository can provide *comments* to discuss the issue, such as adding extra details or suggesting solutions. The issue tracker has two basic functions for issue management: a *label(s)* and assign issues to an *assignee(s)*. Anyone with write access to the repository can create a label, and anyone with triage access to the repository can apply and dismiss labels. Only those with write access to the repository can specify the assignee(s). All the users involved in an

issue, including the issue raiser, the commenter(s), and the assignee(s) are referred to as *participants*.

There are various challenges when employing open-source AI repositories. For example, missing datasets and descriptions can make it difficult for users to evaluate AI models and replicate their results. Additionally, pre-trained AI models should be provided so that users can fine-tune them for other downstream tasks. Furthermore, frameworks used to develop AI software applications are constantly evolving; therefore, installing compatible development environments can be challenging for new users. Moreover, the documentation of open-source AI repositories may only mention the names of installed packages without specifying their versions, which can lead to runtime errors. Lastly, AI models are sensitive to hyperparameters [20], [21], and attempts to replicate the results may fail when using random hyperparameters. These issues can impede the usability of AI repositories. In this paper, we aim to understand these challenges to enable users to make better use of open-source AI repositories when developing their software applications.

III. DATA COLLECTION

This section presents the process of collecting and cleaning the dataset, i.e., GitHub issues of open-source software AI repositories, for our empirical study.

A. Scope

Following a previous study on analyzing open-source software (OSS) AI repositories [22], we include five prestigious

¹³<https://github.com/deepinsight/insightface/issues/1281>

AI conferences in our experiments: NeurIPS, ICML, CVPR, ICCV, and ECCV. Out of these five conferences, CVPR, ICCV, and ECCV are computer vision-oriented. We also collect papers and repositories from other highly regarded AI conferences, such as ICLR, ACL, SIGKDD, AAAI, and AAMAS. All the selected conferences are categorized as A* in the CORE Ranking.¹⁴ In total, our dataset is constructed from ten AI conferences that cover a wide range of AI topics, including natural language processing, reinforcement learning, etc. It is worth noting that frameworks and applications of AI techniques are constantly evolving. Our dataset includes repositories related to AI conferences from 2013 to 2022.

B. Data Curation

Figure 2 illustrates an overview of the workflow for collecting and cleaning our dataset. We use PapersWithCode¹⁵ (PwC), a platform for storing AI-related papers, implementations, and datasets, to obtain AI papers published at conferences and their corresponding repositories. According to the contributing guidelines, PwC users can connect a repository to a paper by adding GitHub, GitLab, or BitBucket URLs. Furthermore, users can monitor all edits on a Slack¹⁶ channel, where everyone is encouraged to review their contributions to ensure data quality. If contributors acknowledge that a repository is the implementation provided by the authors of the corresponding paper, they mark this repository as *official*.

In this paper, we adopt a conservative strategy for selecting repositories. Specifically, we only consider papers whose linked repositories are labeled as *official*. Intuitively, users tend to ask questions in the official repository maintained by the authors, and such official repositories are more likely to contain more issues from which more actionable suggestions can be derived. We use PwC APIs¹⁷ to obtain papers from different venues and their metadata, including a list of authors and their repository links. We then conduct a series of queries from the PwC APIs with the following constraints:

- **Venue:** Papers from the ten selected conferences.
- **Year:** The dates of these conferences are between 1/1/2013 and 31/12/2022.
- **Authority:** All the repositories that are labeled as *official* on PapersWithCode.

In our initial collection, we obtain 652 open-source AI repositories and their corresponding papers from PapersWithCode. Among these repositories, 639 (98.01%) are published on GitHub, nine (1.38%) are hosted on CodaLab, and only four (0.61%) are hosted on Bitbucket. Given the prevalence of GitHub as a platform for hosting open-source AI repositories and the convenience of utilizing GitHub APIs, we decided to focus on the repositories hosted on GitHub for our analysis, resulting in a set of 639 open-source AI repositories.

Subsequently, we use the `git clone` command to download these 639 open-source AI repositories and manually

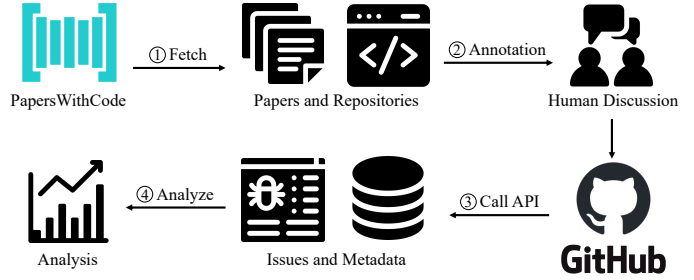


Figure 2. The workflow of collecting and cleaning the data. We obtain the paper list and the corresponding repository list from PapersWithCode. After manually removing the unofficial repositories, we query the GitHub API to obtain the metadata and issue information.

check whether the downloaded repositories are indeed the *official* implementation of the corresponding papers. To verify the relationship between repositories and their corresponding papers, we utilize three strategies. Firstly, the README.md file of an AI repository should contain the title of the corresponding paper. Secondly, the paper must mention the repository URL. Lastly, the repository owners should be listed as the authors of the corresponding paper. To evaluate the accuracy of these strategies, we asked two annotators to manually examine each pair of a repository and its corresponding paper to evaluate the accuracy of our strategies. Following prior studies [23], we employ Cohen’s Kappa [24] to measure the inter-annotator agreement, which is 0.72. According to [25], this inter-annotator agreement value indicates substantial agreement among our annotators. In cases of disagreement, a third annotator was consulted to finalize the decision. The three annotators, who are Ph.D. students with over five years of experience in using GitHub, were able to confirm 576 pairs of repositories and their corresponding papers as *official* implementations.

We also employ the GitHub REST API¹⁸ to fetch the metadata and issue information of the collected repositories. The metadata of these repositories includes information such as contributors, the number of stars, the number of forks, the number of watchers, the creation date, and the last update date. We also obtain the issue list of each repository and compute the number of open and closed issues. For each issue, we extract its title, discussion history, status, creation date, close date (if the issue has been closed), assignees, and labels. Our final dataset consists of 576 pairs of repositories and their corresponding papers. Additionally, we find that 24,953 issues belong to these repositories.

IV. RESULTS

In this section, we present our analysis of the issues in open-source software (OSS) AI repositories. We conduct experiments to answer the following three research questions:

- *RQ1. What do users discuss in the issues of open-source AI repositories?*

¹⁸<https://docs.github.com/en/rest?apiVersion=2022-11-28>

¹⁴<https://www.core.edu.au/conference-portal>

¹⁵<https://paperswithcode.com>

¹⁶<https://slack.com/intl/en-au/>

¹⁷<https://paperswithcode.com/api/v1/docs/>

- RQ2. How are issues in open-source AI repositories managed and addressed?
- RQ3. What are the relationships between different features and the closure of issues in open-source AI repositories?

A. RQ1. What do users discuss in the issues of open-source AI repositories?

Motivation. The development of AI systems follows a data-driven paradigm [26], [27], in which AI models learn knowledge from a large amount of data automatically, as opposed to traditional software systems where developers encode their programming language through control flow and data flow graphs [28], [29]. Additionally, AI systems often require specialized hardware, such as a GPU or Google TPU, to accelerate the training and testing processes. Therefore, the issues that developers face in open-source AI repositories may differ from those in traditional OSS repositories. In order to better understand the primary issues encountered in the development and usage of AI applications, we aim to construct a taxonomy of the issues in open-source AI repositories.

Categorization Process. We obtain 24,953 issues from the 576 repositories using the GitHub REST API. Among these repositories, 122 (21.18%) have no GitHub issues. Previous studies [30], [31] have used the card sorting method, a technique for creating mental models and deriving taxonomies from data, to categorize software engineering-related discussions. As the issues in open-source AI repositories can be unique, we adopt open card sorting, which uses no predefined groups to create the taxonomy for issues in AI repositories. The taxonomy is intended to categorize *the primary reason why developers raise an issue* rather than the *root cause of this issue*, the latter of which is usually difficult to determine without sufficient information. For example, a developer may raise an issue asking why the results are impossible to reproduce, but the root cause can be due to many reasons, such as differences in library versions, hardware, or randomness. For this issue, we place it in the `Failure to replicate` category. In this paper, we followed three steps to categorize issues in the open-source AI repositories. These steps are described as follows:

- (i) *Preparation*: We randomly pick 384 issues¹⁹ from all the collected issues, then we import these issues into a Notion database and create a card for each issue.
- (ii) *Execution*: Two annotators discuss with each other how to categorize the cards into meaningful groups. We adopt the *open card* sorting method, which lets the groups emerge and evolve during the sorting process. In this step, the two authors are encouraged to create groups. This step takes about seven hours.
- (iii) *Analysis*: Two annotators who conduct card sorting, as well as a third annotator, discuss together how to merge

¹⁹It is statistically representative sample size using a popular sample size calculator (<https://www.surveysystem.com/sscalc.htm>) with a confidence level of 99% and a confidence interval of 10.

Table I
THE DISTRIBUTION OF ISSUES OF DIFFERENT CATEGORIZES.

ID	Category	No.	Rate
1	Missing Content	28	7.29%
	· Missing Dataset	10	2.60%
	· Missing Model	9	2.34%
	· Missing Code/Module	9	2.34%
2	Unclear Instructions	75	19.53%
	· Unclear Code Usage Instruction	56	14.58%
	· Unclear Data Information	12	3.13%
	· Unclear Environment Configuration	7	1.82%
3	Discuss Methodology	10	2.60%
4	Extension	42	10.94%
5	Runtime Error	89	23.18%
6	Discuss Implementation	35	9.11%
7	Fail to Replicate	13	3.38%
8	Enhancement	16	4.17%
9	Abnormal Behavior	50	13.02%
10	Paper-Code Misalignment	5	1.30%
11	Supplementary Information	4	1.04%
12	Re-implementation	15	3.91%
13	Others	2	0.52%

groups of relevant topics into more general categories. This step takes approximately 1.5 hours.

Categorization Results. We obtain 13 categories; two of them are further divided into three subcategories. Table I shows the categories and subcategories in the sampled issues. In our taxonomy, each issue is assigned to one category, which reflects the primary reason for raising this issue. We discuss the details of these categories as follows:

- **Missing Content**: Users raise issues as they are unable to find the contents used to employ open-source AI repositories. There are 28 issues that fall into this category, which is subsequently divided into three different subcategories, such as *missing dataset* (10 issues), *missing model* (9 issues), and *missing code/module* (9 issues). For example, the issue²⁰ titled “where to download the Paris StreetView Dataset” requests the Paris StreetView dataset for running the Context Encoders model to generate the contents of an arbitrary image [32].
- **Unclear Instructions**: The category is mentioned when users are unable to comprehend the instructions used to employ open-source AI repositories. This category contains 75 issues (19.53%), which is the second-largest category. We define three subcategories based on information from the instructions. The first subcategory is *unclear environment configuration*, e.g., the version of a programming language library in an AI repository.²¹ The second subcategory is *unclear data information*, e.g., how to split a dataset into training and testing sets. The third subcategory is *unclear code usage instruction*, e.g., hyperparameters used to train AI models.
- **Discuss Methodology**: This category, which includes ten issues, focuses on discussing the methodology

²⁰<https://github.com/pathak22/context-encoder/issues/24>

²¹<https://github.com/ShichenLiu/CondenseNet/issues/22>

of AI models. For example, an issue²² asks the input size of an AI model.

- **Extension:** Users raise their issues to ask AI repositories' owners whether AI models can be extended beyond their original scope. For example, users may want to apply AI models to a new dataset or do some modifications to these models.
- **Runtime Error:** We define this category for users who encounter runtime errors when using AI repositories. The category includes 89 issues (23.18%) which is the largest population of all categories. There are various reasons for the runtime errors, e.g., the users may not have the required packages installed, the users employ the wrong path, or the hardware is incompatible. We find that it is challenging to further divide these issues into subcategories based on the root cause of the runtime errors. It is difficult to further divide these issues into subcategories based on the cause of the runtime errors.
- **Discuss Implementation:** This category aims to discuss the implementation details of AI repositories. For example, users may ask whether we should apply dropout techniques to AI models. Some users may also require the AI repositories' owners to clarify the details of their implementation.
- **Failure to Replicate:** Users often submit their issues as they fail to employ an AI repository to replicate results reported in its corresponding paper. We put these issues in this category.
- **Enhancement:** Issues in this category aim to improve the quality of AI repositories, e.g., by adding new features, fixing bugs, or updating libraries.
- **Abnormal Behaviour:** This category focuses on users who may encounter performance issues, e.g., the training time of AI models takes too long.²³
- **Paper-Code Misalignment:** Issues related to this category focus on the difference between the code in AI repositories and the descriptions of their corresponding papers. For example, users may ask why the dimensions of hidden layers in an AI repository differ from the descriptions on its paper.²⁴
- **Supplementary Information:** These issues request supplementary information for AI repositories, e.g., slides, documents, and videos of AI models.
- **Re-implementation:** Users aim to re-implement AI repositories by employing other programming languages; hence, they may raise issues to discuss with the AI repositories' owners and other developers.
- **Others:** The remaining issues that are unable to categorize into any of the above categories. For example, some issues have insufficient information to be categorized.²⁵

²²<https://github.com/irwinherrmann/stochastic-gates/issues/1>

²³<https://github.com/open-mmlab/mmdetection/issues/5237>

²⁴<https://github.com/timmeinhardt/trackformer/issues/58>

²⁵https://github.com/matterport/Mask_RCNN/issues/2046

Answer to RQ1

We categorize the issues in open-source AI repositories into 13 classes, including: Missing Content, Unclear Instructions, Discuss Methodology, Extension, Runtime Error, Discuss Implementation, Failure to Replicate, Enhancement, Abnormal Behaviour, Paper-Code Misalignment, Supplementary Information, Re-implementation, and Others.

B. RQ2. How are the issues managed and addressed in open-source AI repositories?

Motivation. The issues raised by users in open-source AI repositories serve as an important indicator of problems encountered during usage and also serve as a means of communication with repository maintainers. Effective management and addressing of these issues can greatly impact the user experience and the overall quality of the AI repositories. As the first research question highlights, users raise issues for a variety of reasons, such as missing content, unclear instructions, or runtime errors. Ignoring these issues can prevent users from successfully replicating the results reported in the repositories. While some issues can be easily resolved by users, such as installing required packages, other issues such as runtime errors or unclear instructions require better management and addressing in order for users to effectively employ the AI repositories for building their software applications. Considering the rapid growth of the AI field, it is crucial to investigate these issues to aid maintainers in effectively managing open-source AI repositories.

How are the issues addressed? We aim to understand how to address the issues present in open-source AI repositories. These issues are categorized into two groups, *open* and *closed* issues. Open issues are those that have yet to be addressed, while closed issues are those that have been resolved. Utilizing the GitHub REST API, we collected 24,953 issues from open-source AI repositories. Of these issues, 8,102 (32.47%) are open issues, and 16,851 (67.53%) are closed issues. As shown in Figure IV-B (a), the distribution of the *closed issue rate* (CIR) for each repository reveals that 35.76% of repositories have closed all of their issues. However, it is worth noting that these repositories with a CIR higher than 0.9 have an average of 83.83 issues, indicating that these repositories have a larger number of issues in total.

We discover that 40% of closed issues are “*self-closed*,” indicating that the individuals who raised the issues are able to find solutions to their problems through discussion. Additionally, we identify another type of issue referred to as *ignored issues*. These are issues that only the individuals who raised them are interested in and typically receive no response from other users or developers. These ignored issues comprise 11.79% of all issues. We also define the *ignored issue rate* as the ratio of ignored issues for each repository. This metric allows us to evaluate the level of active engagement by

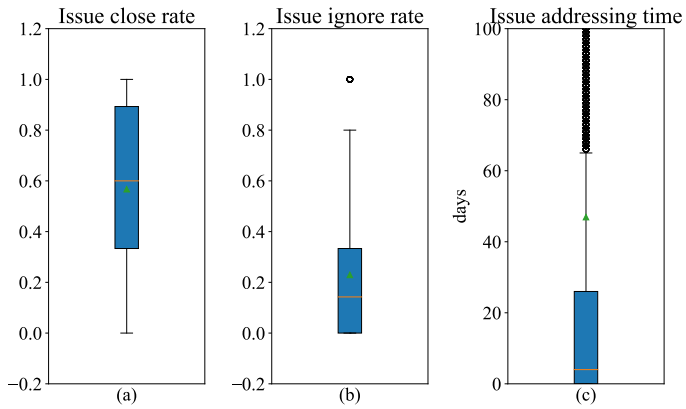


Figure 3. The distribution of closed issue rates (a), addressing time (b), and ignored issue rates (c) in investigated repositories.

repository maintainers in addressing issues. The results of the ignored issue rate are presented in Figure IV-B (b).

Additionally, we investigate the addressing time of issues, i.e., the time between the creation of an issue and its closure. We found that the average addressing time and standard deviation are 46.94 days and 135.34 days, respectively, in the 16,851 closed issues. However, the median addressing time is only four days, while the longest addressing time is over two years. Figure IV-B (c) illustrates the distribution of the addressing time. We analyze what factors affect the addressing time by considering two features at the repository level: (1) the number of contributors and (2) the total number of issues. We hypothesize that more contributors can reduce the addressing time. However, Spearman’s rank correlation [33] shows that the relationship between the number of contributors and the addressing time is not statistically significant (p -value > 0.05). We also hypothesize that more issues can increase the average addressing time; intuitively, more issues mean more workload on the maintainers. Spearman’s rank correlation validates this hypothesis by finding a weak correlation between the number of issues and average addressing time (p -value < 0.01).

How are the issues managed? GitHub provides several methods to help maintainers manage issues. We consider two methods: (1) using labels to tag issues and (2) assigning issues to specific people to address. In the remaining part of this section, we use the terms `labeling` and `assigning` to refer to these two methods, respectively. However, our results show that the two methods are not widely adopted to manage issues in open-source AI repositories. Only 7.81% of the repositories use labels to categorize issues, and only 5.90% of repositories assign issues to assignees. We find that 11.42% of issues are tagged with labels and 16.45% of issues have assignees, which suggests that *repositories with more issues tend to use these two methods to manage issues more often*.

We analyze the use of `labeling`. GitHub provides a set of nine default labels, such as `bug`, `duplicate`, `enhancement`, `help wanted`, `invalid`, `question`, `wontfix`, and `good first issue`. One sign of actively managing issues is to define new labels




Figure 4. The distribution of the top 10 most frequently used labels. The numbers represent the labels, where 1: bug, 2: enhancement, 3: awaiting response, 4: community help wanted, 5: help wanted, 6: good first issue, 7: community discussion, 8: reimplementation, 9: question, 10: feature request.

beyond just using the default ones. We find that among all the repositories that use labels, 14.41% of them describe new labels. Including the default labels, we discover 109 unique labels being used in all the repositories. We then count the occurrences of each label. Table 4 shows the ranked list of the most frequently used default labels.

Answer to RQ2

In the investigated open-source AI repositories, 67.53% of issues are closed, and 50% of issues are closed within four days. Only 7.81% of the repositories use labels to categorize issues, and only 5.9% of them assign issues to specific assignees. We find that repositories with more contributors do not necessarily address issues faster, but repositories with more issues tend to take a longer time to close issues.

C. RQ3. What are the relationships between different features and the closure of issues in open-source AI repositories?

Motivation. This research question aims to explore the relationships between the various features of issues and the closure of those issues. The answer to this research question can provide practical suggestions for both repository maintainers and issue raisers to improve the management and resolution of issues in open-source AI repositories.

Investigated Features. We study the relationships between nine features and the closure of issues. The first eight features are related to the issues and the 9th feature, i.e., `num-contributors` is related to the repositories. We present the nine features as follows:

- `has-label`: Whether an issue is labeled. In GitHub, anyone with triage access to a repository can tag issues with labels. A labeled issue usually means that it has been noticed and read by the maintainers and, as a result, it may be more likely to be addressed.

- `has-assignee`: Whether an issue is assigned. Users who have a write access permission to a repository can assign issues to developers. If the issue is assigned, the assignee is responsible for addressing the issue, which may lead to a successful resolution.
- `title-length`: The length of the issue title. An informative title may help the repository maintainers better understand the issue at a glance. Usually, a longer title carries more information that can help catch the maintainers’ attention, leading to a faster solution.
- `body-length`: The length of the issue body. Similarly, an informative description can help the repository maintainers better understand the issue. In some guidelines, the issue raisers are encouraged to provide a detailed description of the issue, which can help the maintainers better understand the issue.
- `sentiment`: The sentiment of the issue description. If an issue is impolite, e.g., includes negative sentiment, a repository maintainer may be reluctant to address it.
- `is-English`: Whether the issue is written in English. We noticed that 5.94% of issues are written in a language other than English, e.g., Chinese, leading the repository maintainers to barely understand these issues to provide their solutions.
- `has-code`: Whether the issue description contains code blocks. Code blocks are useful for the repository maintainers to understand the problem. For example, issue raisers may include code snippets corresponding to error tracks in the issue description.
- `has-url`: Whether the issue description contains URLs. GitHub uses Markdown syntax to format the issue description. An URL can be used to refer to another issue, a website, or an image. This information may help repository maintainers better understand the issue.
- `num-contributors`: The number of contributors in the repository. Intuitively, the more contributors in a repository, the more likely the issue is to be addressed.

Feature Measurement and Statistical Test. We obtain issues by querying the GitHub API. We compute the length of the issue title and body by converting them into Python strings and using the results returned by the `len()` function. Following the previous work [34], we employ Senti4SD [35], a sentiment analysis tool specifically tuned for software engineering tasks, to compute the sentiment of the issue description. Senti4SD produces three sentiment scores: *positive*, *negative*, and *neutral*. As we train Senti4SD on texts written in English, we discard the issues that contain non-English characters. For issues that contain code snippets and URLs that are not relevant to analyzing the sentiment of the issue description, we remove them using regular expressions.

We first tokenize the issue title and body into characters. We then check whether the issue contains characters that are in the American Standard Code for Information Interchange²⁶

Table II
THE SIGNIFICANCE OF THE DIFFERENCES BETWEEN THE FEATURE VALUES OF OPEN AND CLOSED ISSUES.

Feature	<i>p</i> -value	Cohen’s δ
<code>has-label</code>	< 0.01	0.26 (S)
<code>has-assignee</code>	< 0.01	0.50 (M)
<code>title-length</code>	> 0.05	0.02 (N)
<code>body-length</code>	< 0.01	0.08 (N)
<code>sentiment</code>	> 0.05	0.01 (N)
<code>is-Non-English</code>	> 0.05	0.05 (N)
<code>has-code</code>	< 0.01	0.14 (N)
<code>has-pics</code>	> 0.05	0.01 (N)
<code>num-contributors</code>	< 0.01	0.54 (M)

(ASCII). If the issue contains non-ASCII characters, we consider it `non-English`. The GitHub issues use Markdown syntax, which allows users to format their issue description using code blocks (i.e., code snippet in ```` and `'''`) and URLs (i.e., `[text]` (URL)). We employ regular expressions to detect whether the issue description contains code blocks or URLs using Markdown syntax.

We group the issues into two groups: *open* and *closed* issues. For each issue, we measure its eight features, such as `has-label`, `has-assignee`, or `title-length`. The number of contributors in a repository is obtained by querying the GitHub API to fetch metadata about the repository. Then, we adopt the Wilcoxon rank-sum test [36] to compute the significance of the differences between the feature values of open and closed issues. We also estimate the Cohen’s effect sizes δ of the difference between the open and closed issues. According to the guideline for using the effect size, when $|\delta|$ is less than 0.2, between 0.2 and 0.5, between 0.5 and 0.8, and larger than 0.8, the effect size is considered negligible, small, medium, and large, respectively.

Table II shows the results of the Wilcoxon rank-sum test and Cohen’s effect sizes δ for the nine features. Five features, such as `has-label`, `has-assignee`, `body-length`, `has-code`, and `num-contributors`, show a statistically significant difference between the open and closed issues. The `num-contributors` feature also has a medium effect size (M), indicating that the more contributors who join the repository, the better we can address the issues. The remaining features, i.e., `title-length`, `sentiment`, `is-English`, and `has-url`, indicate that there is no significant difference between the open and closed issues.

We classify the issue-related features into two groups, i.e., features related to the practices of raising issues and features related to the practices of handling issues. Our results show that issues are more likely to be closed if they have assignees. Moreover, issues with labels are more likely to be closed than those without labels. The effect sizes of `has-label` and `has-assignee` are small and medium, respectively. The two features, i.e., `body-length` and `has-code`, are related to how users raise issues. Although their effect sizes are smaller than that of `has-label` and `has-assignee`, they still demonstrate statistically significant differences.

²⁶<https://en.wikipedia.org/wiki/ASCII>

Table III
THE DIFFERENCE BETWEEN INDUSTRY AND ACADEMIC REPOSITORIES.

Class	Academic	Industry	p -value	δ
No. Issues	22.57	257.62	<0.01	Small
No. Stars	232.82	2749.10	<0.01	Medium
No. Contributors	2.25	41.06	<0.01	Small
Address Time	45.47	47.90	<0.01	Negligible
Close Rate	57%	58%	>0.05	Negligible
Assign Rate	1.35%	5.15%	<0.05	Small
Label Rate	2.45%	6.45%	>0.05	Small

Answer to RQ3

Our results show that issues are more likely to be closed if they have assignees. Moreover, issues with labels are more likely to be closed than those without labels. Besides, issues with more details (e.g., longer descriptions and code blocks) have a higher chance of being closed.

V. DISCUSSION

A. Industry vs. Academic Repositories

Developers from industry and academia work closely to promote the development of AI systems together. On the one side, industry practitioners also use academic AI repositories to build their products. On the other side, academic researchers use popular AI frameworks originating from industry to create their models. We are also interested in the difference between issues in the industry and academic repositories. Similar to the process in [37], we ask two annotators to manually split collected repositories into two types: repositories owned by industry and those owned by academia. The inter-annotator agreement value is 0.62, indicating substantial agreement among our annotators [25]. We invite another annotator to resolve the disagreement between the two annotators. In total, we have 63 repositories owned by industry and 513 repositories owned by academia. We compare the two groups of repositories from the perspectives of the number of issues, addressing time, close rate, label distribution, assigning rate, and labeling rate. We conduct a Wilcoxon rank-sum test to check whether the difference is statistically significant. Table III presents our results.

On average, each industry repository includes 257.62 issues, while each academic repository contains 22.57 issues. However, it does not necessarily indicate that industry repositories have low quality. The average number of stars in industry repositories (2749.10) is one order of magnitude greater than that in academic repositories (232.82), meaning that industry repositories attract much more attention and, therefore, the number of issues in these repositories is larger. Note that the ratio of the number of issues to the number of stars is close: 0.094 for industry repositories and 0.097 for academic ones.

The industry repositories have over 18 times more average contributors than the academic ones (41.06 vs. 2.25). The ratio of the number of issues to the number of contributors is 6.27

for industry repositories and 10.03 for academic repositories. However, a smaller issue-contributor ratio does not lead to faster issue-addressing time; industry repositories even have a slightly (but statistically significant) longer addressing time (47.90 days) than academic repositories (45.47 days). Besides, the issue close rate is also similar between the two types of repositories (0.58 and 0.57). One possible explanation is that the academic repositories are less active after the code is released, meaning that the contributors mainly focus on addressing the issues. However, industry repositories require their contributors to put more effort into both developing new features and addressing issues.

The two types of repositories also demonstrate different behaviors in using the issue management tools. Although both the labeling and assigning functions are not widely adopted in the open-source AI repositories, there are more issues in industry repositories that are labeled and assigned. The assign rates are 5.15% and 1.35% for industry and academic repositories, respectively; the difference is also statistically significant (p -value < 0.05). The label rates are 6.45% and 2.45% for industry and academic repositories, respectively. However, the difference is not statistically significant (p -value > 0.05). Overall, the industry repositories are more active in using the issue management features than the academic ones.

B. Implications

Based on the analysis of each research question, we summarize the implications of our study as the suggestions to open-source AI repository developers and users.

In this study, we find that the most prevalent issue among the open-source AI repositories is the `Runtime Error` category, which can be caused by a variety of factors related to other categories of issues, such as missing datasets or missing models. The second most common category of issues falls under the `Unclear Instructions`, which is primarily caused by a lack of detailed descriptions. To address these issues and improve the quality of open-source AI repositories, we propose two suggestions for the maintainers: (1) ensure that all necessary datasets, models, and other files are provided to allow for accurate replication of results, and (2) provide clear and detailed instructions on how to use the repository, including information on dataset splitting, setting hyperparameters, and managing library dependencies.

Our experiment results show that issues are more likely to be closed if they have assignees and labels. Therefore, we recommend that repository maintainers actively leverage the features provided by GitHub to manage issues by (1) tagging issues with labels and (2) assigning specific people to be responsible for addressing the issues. Besides, the length of issues descriptions and whether the issues have code blocks cause a statistically significant difference to the closure of issues. Based on the results, we make two suggestions to the issue raisers: (1) provide more details to precisely describe the issue and (2) use the Markdown syntax to formulate the issue description with code snippets.

C. Threats to Validity

Threats to Internal Validity. We collect a list of repositories from the PaperWithCode platform, which is a community-driven platform where users can link papers with their corresponding repositories. However, there is a risk that some of these repositories may not be official implementations, which can be of low quality and attract less attention from users. To mitigate this threat, we only keep the repositories that are marked as “official implementations.” Additionally, we compare the repositories and paper information, such as owner and author information, to ensure that the repositories are official. This process is done manually, with two authors involved, and any disagreements are resolved with the involvement of another annotator. Furthermore, we manually group the repositories by their ownership, such as industry and academia. However, considering that academia and industry work closely together in AI research, it may be difficult to unanimously agree on the ownership of a repository. In such cases, another annotator is involved in the process to lead the discussion and resolve any disagreements.

Threats to External Validity. Our study collects papers and repositories from the flagship AI conferences within a ten-year timeframe. However, it is important to note that there is a risk that the conclusions drawn from our study may not be generalizable to a broader range of open-source AI repositories. The field of AI research encompasses a wide range of subtopics, some of which are more popular than others. This leads to an imbalanced distribution of repositories tackling various topics in our dataset, which may affect the generalizability of our conclusions to other AI subtopics or interdisciplinary AI research. Additionally, while we analyze open-source AI repositories hosted on GitHub, which is the most popular platform for open-source AI projects, other platforms such as GitLab or BitBucket, which employ different issue trackers, such as JIRA, also contain open-source AI repositories. These platforms may have different issue management strategies, which could lead to diverse conclusions.

AI is a fast-evolving field. Future work that replicates our study in a wider scope of open-source AI repositories, e.g., more venues, longer time periods, more AI subtopics, or other hosting platforms, can further validate and generalize our findings. We also make the replication package publicly available to facilitate such replication.

VI. RELATED WORK

A. Challenges and Practices in AI Development

There has been a series of works investigating the challenges and best practices of various stages in the development of AI software, e.g., data collection, system designing, model testing and deployment, etc.

Whang and Lee [38] discussed the challenges of data collection for AI systems. Amershi et al. [39] conducted surveys and interviews to understand the process of developing AI software in Microsoft teams. They organized the response from developers into a set of best practices. For example, in the

data collection stage, they suggested reusing the data as much as possible to reduce duplicated effort. Paleyes et al. [40] provided suggestions on using flow-based programming to better discover and collect data for AI systems. Researchers also investigate the code smells [41] and data smells [42], [43] of AI software systems. Serban et al. [44] evaluated the adoption and effects of using conventional software engineering practices in AI software development. Wan et al. [45] analyzed differences between the development of machine learning systems and the development of non-machine-learning systems, which derives recommendations for AI developers.

Song et al. [46] explored the practices of testing machine learning software via conducting an interactive rapid review with industry practitioners. Software engineering researchers designed a list of methods to test and improve different AI systems (e.g., speech recognition [47]–[49], reinforcement learning [50], language models [51], etc) from various perspectives beyond correctness, e.g., fairness [23], [52]–[54], robustness [55]–[57], security [58], [59], etc. Shneiderman [60] discussed the guidelines for addressing the ethics and reliability issues in human-centered AI systems. Paleyes et al. [61] presented a survey of case studies of the challenges and practices in deploying AI systems. Fan et al. [22] mined academic AI repositories and provided some suggestions to improve the quality and popularity of open-source AI repositories.

B. Mining GitHub Issues

Researchers have investigated the bugs in deep learning systems [62], [63]. This paper mainly focuses on the issues raised in AI repositories, which according to our taxonomy in RQ1 covers contents beyond bugs, e.g., paper discussion. We also introduce the works on mining GitHub issues.

GitHub Issues [19] is an issue-tracking system integrated into GitHub that allows users to track tasks, enhancements, and bugs [64] for repositories hosted on GitHub. A series of empirical studies have been established to analyze software repositories and their issues. In 2013, Bissyande et al. [65] conducted a large-scale investigation of issue trackers from tens of thousands of GitHub projects and obtained interesting findings, i.e., the correlation between the numbers of issue raisers and the addressing time. *Ticket Tagger* was presented to automatically assign labels to issues in the GitHub issue tracker [66], [67]. In this paper, we analyze the influence of labels in managing GitHub issues of AI repositories. Our results show that overall the label function is not frequently used in open-source AI repositories. Moreover, there is a significant difference in label usage distribution between open and closed issues. We believe tools for automatically tagging issues in AI repositories can boost the efficiency of issue resolution [68], [69]. Recently, Hata et al. [70] analyzed the early usage of GitHub Discussion, a new feature of GitHub that allows users to discuss issues and topics. Their approach, i.e., the sentiment of the content, is also similar to our study. However, their paper does not focus on issue tracking systems.

VII. CONCLUSION AND FUTURE WORK

This paper presents the first systematic empirical study of issues in open-source AI repositories. Using the PapersWithCode platform, we collected 576 open-source AI repositories hosted on GitHub and 24,953 issues from these repositories. Through manual analysis of a representative sample of issues, we developed a taxonomy of 13 categories, with the two most common being runtime error (23.18%) and unclear instructions (19.53%). Our findings also showed that 67.5% of issues were closed, with half of them closed within five days. Additionally, we discovered that GitHub issue management features were not widely adopted in open-source AI repositories, with only 7.81% and 5.9% of repositories using labels and assignees, respectively. We empirically show that employing GitHub issue management features and writing issues with detailed descriptions (e.g., using more code blocks) can help resolve issues. Our study provides recommendations for repository maintainers to improve the management and quality of open-source AI repositories.

In the future, we intend to conduct surveys and interviews to uncover the reasons behind the lack of widespread adoption of issue management features in open-source AI repositories and to explore ways to increase their usage. Additionally, we will investigate the topics of discussion surrounding the development of AI systems, such as AI testing and data management, on platforms like StackOverflow and GitHub.

Reproducibility. The code and documentation, along with the obtained data, have been made open-source for reproducibility: <https://github.com/soarsmu/Mining-AI-repos-issues>.

ACKNOWLEDGMENT

This research is supported by the Ministry of Education, Singapore under its Academic Research Fund Tier 3 (Award ID: MOET32020-0004). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore.

REFERENCES

- [1] M. Danilevsky, K. Qian, R. Aharonov, Y. Katsis, B. Kawas, and P. Sen, "A survey of the state of explainable ai for natural language processing," *arXiv preprint arXiv:2010.00711*, 2020.
- [2] K. Chowdhary, "Natural language processing," *Fundamentals of artificial intelligence*, pp. 603–649, 2020.
- [3] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (xai): Toward medical xai," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 11, pp. 4793–4813, 2020.
- [4] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [5] C. Yang, B. Xu, J. Khan, G. Uddin, D. Han, Z. Yang, and D. Lo, "Aspect-based api review classification: How far can pre-trained transformer model go?" in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. Los Alamitos, CA, USA: IEEE Computer Society, mar 2022, pp. 385–395. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SANER53432.2022.00054>
- [6] B. Alsulami, E. Dauber, R. Harang, S. Mancoridis, and R. Greenstadt, "Source code authorship attribution using long short-term memory based networks," in *Computer Security - ESORICS 2017*. Springer Verlag, 2017, pp. 65–82.

- [7] W. D. Heaven, "Ai is wrestling with a replication crisis," Apr 2022. [Online]. Available: shorturl.at/djMN3
- [8] O. E. Gundersen, Y. Gil, and D. W. Aha, "On reproducibility ai: Towards reproducible research, open science, and digital scholarship in ai publications," *AI magazine*, vol. 39, no. 3, pp. 56–68, 2018.
- [9] S. Sonnenburg, M. L. Braun, C. S. Ong, S. Bengio, L. Bottou, G. Holmes, Y. LeCun, K.-R. Muller, F. Pereira, C. E. Rasmussen *et al.*, "The need for open source software in machine learning," *Journal of Machine Learning Research*, vol. 8, no. 81, pp. 2443–2466, 2007. [Online]. Available: <http://jmlr.org/papers/v8/sonnenburg07a.html>
- [10] J. Spohrer, "The role of open-source software in artificial intelligence," *AI Magazine*, vol. 42, no. 1, pp. 93–94, Apr. 2021. [Online]. Available: <https://ojs.aaai.org/index.php/aimagazine/article/view/7488>
- [11] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [12] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [13] E. Strickland, "Ibm watson, heal thyself: How ibm overpromised and underdelivered on ai health care," *IEEE Spectrum*, vol. 56, no. 4, pp. 24–31, 2019.
- [14] A. Carvalho, A. Levitt, S. Levitt, E. Khaddam, and J. Benamati, "Off-the-shelf artificial intelligence technologies for sentiment and emotion analysis: a tutorial on using ibm natural language processing," *Communications of the Association for Information Systems*, vol. 44, no. 1, p. 43, 2019.
- [15] M. A. Goralski and T. K. Tan, "Artificial intelligence and sustainable development," *The International Journal of Management Education*, vol. 18, no. 1, p. 100330, 2020.
- [16] N. Tomašev, J. Cornebise, F. Hutter, S. Mohamed, A. Picciariello, B. Connelly, D. Belgrave, D. Ezer, F. C. v. d. Haert, F. Mugisha *et al.*, "Ai for social good: unlocking the opportunity for positive impact," *Nature Communications*, vol. 11, no. 1, pp. 1–6, 2020.
- [17] B. Zhang, "An explorative study of github repositories of ai papers," *arXiv preprint arXiv:1903.01555*, 2019.
- [18] D. Gonzalez, T. Zimmermann, and N. Nagappan, "The state of the ml-universe: 10 years of artificial intelligence & machine learning software development on github," in *Proceedings of the 17th International Conference on Mining Software Repositories*, 2020, pp. 431–442.
- [19] T. Preston-Werner, "Github issue tracker!" Apr 2009. [Online]. Available: <https://github.blog/2009-04-15-github-issue-tracker/>
- [20] Y. Yang, W. Youyou, and B. Uzzi, "Estimating the deep replicability of scientific findings using human and artificial intelligence," *Proceedings of the National Academy of Sciences*, vol. 117, no. 20, pp. 10762–10768, 2020.
- [21] D. Zhuang, X. Zhang, S. Song, and S. Hooker, "Randomness in neural network training: Characterizing the impact of tooling," *Proceedings of Machine Learning and Systems*, vol. 4, pp. 316–336, 2022.
- [22] Y. Fan, X. Xia, D. Lo, A. E. Hassan, and S. Li, "What makes a popular academic ai repository?" *Empirical Software Engineering*, vol. 26, no. 1, pp. 1–35, 2021.
- [23] M. H. Asyrofi, Z. Yang, I. N. B. Yusuf, H. J. Kang, F. Thung, and D. Lo, "Biasfinder: Metamorphic test generation to uncover bias for sentiment analysis systems," *IEEE Transactions on Software Engineering*, vol. 48, no. 12, pp. 5087–5101, 2022.
- [24] M. L. McHugh, "Interrater reliability: the kappa statistic," *Biochemia medica*, vol. 22, no. 3, pp. 276–282, 2012.
- [25] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *biometrics*, pp. 159–174, 1977.
- [26] E. Ntoutsi, P. Fafalios, U. Gadiraju, V. Iosifidis, W. Nejdl, M.-E. Vidal, S. Ruggieri, F. Turini, S. Papadopoulos, E. Krasanakis *et al.*, "Bias in data-driven artificial intelligence systems—an introductory survey," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 10, no. 3, p. e1356, 2020.
- [27] C. T. Wolf, "Ai models and their worlds: Investigating data-driven, ai/ml ecosystems through a work practices lens," in *International conference on information*. Springer, 2020, pp. 651–664.
- [28] F. E. Allen, "Control flow analysis," *ACM Sigplan Notices*, vol. 5, no. 7, pp. 1–19, 1970.
- [29] K. M. Kavi, B. P. Buckles, and U. N. Bhat, "A formal definition of data flow graph models," *IEEE Transactions on computers*, vol. 35, no. 11, pp. 940–948, 1986.

- [30] A. Begel and T. Zimmermann, "Analyze this! 145 questions for data scientists in software engineering," in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 12–23. [Online]. Available: <https://doi.org/10.1145/2568225.2568233>
- [31] S. Breu, R. Premraj, J. Sillito, and T. Zimmermann, "Information needs in bug reports: Improving cooperation between developers and users," in *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*, ser. CSCW '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 301–310. [Online]. Available: <https://doi.org/10.1145/1718918.1718973>
- [32] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2536–2544.
- [33] P. Sedgwick, "Spearman's rank correlation coefficient," *Bmj*, vol. 349, 2014.
- [34] H. Hata, N. Novielli, S. Baltes, R. G. Kula, and C. Treude, "Github discussions: An exploratory study of early adoption," *Empirical Softw. Engg.*, vol. 27, no. 1, jan 2022. [Online]. Available: <https://doi.org/10.1007/s10664-021-10058-6>
- [35] F. Calefato, F. Lanubile, F. Maiorano, and N. Novielli, "Sentiment polarity detection for software development," in *Proceedings of the 40th International Conference on Software Engineering*, ser. ICSE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 128. [Online]. Available: <https://doi.org/10.1145/3180155.3182519>
- [36] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bulletin*, vol. 1, no. 6, pp. 80–83, 1945. [Online]. Available: <http://www.jstor.org/stable/3001968>
- [37] R. Widyasari, Z. Yang, F. Thung, S. Q. Sim, F. Wee, C. Lok, J. Phan, H. Qi, C. Tan, Q. Tay, and D. Lo, "Niche: A curated dataset of engineered machine learning projects in python," in *Proceedings of the 20th International Conference on Mining Software Repositories*, ser. MSR '23.
- [38] S. Whang and J.-G. Lee, "Data collection and quality challenges for deep learning," *Proceedings of the VLDB Endowment*, vol. 13, pp. 3429–3432, 08 2020.
- [39] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, and T. Zimmermann, "Software engineering for machine learning: A case study," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 2019, pp. 291–300.
- [40] A. Paleyes, C. Cabrera, and N. D. Lawrence, "Towards better data discovery and collection with flow-based programming," *arXiv preprint arXiv:2108.04105*, 2021.
- [41] H. Zhang, L. Cruz, and A. van Deursen, "Code smells for machine learning applications," in *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, ser. CAIN '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 217–228. [Online]. Available: <https://doi.org/10.1145/3522664.3528620>
- [42] A. Shome, L. Cruz, and A. van Deursen, "Data smells in public datasets," in *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, ser. CAIN '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 205–216. [Online]. Available: <https://doi.org/10.1145/3522664.3528621>
- [43] H. Foidl, M. Felderer, and R. Ramler, "Data smells: Categories, causes and consequences, and detection of suspicious data in ai-based systems," in *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, ser. CAIN '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 229–239. [Online]. Available: <https://doi.org/10.1145/3522664.3528590>
- [44] A. Serban, K. van der Blom, H. Hoos, and J. Visser, "Adoption and effects of software engineering best practices in machine learning," in *Proceedings of the 14th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, ser. ESEM '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: <https://doi.org/10.1145/3382494.3410681>
- [45] Z. Wan, X. Xia, D. Lo, and G. C. Murphy, "How does machine learning change software development practices?" *IEEE Transactions on Software Engineering*, vol. 47, no. 9, pp. 1857–1871, 2021.
- [46] Q. Song, M. Borg, E. Engström, H. Ardö, and S. Rico, "Exploring ml testing in practice: Lessons learned from an interactive rapid review with axis communications," in *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, ser. CAIN '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 10–21. [Online]. Available: <https://doi.org/10.1145/3522664.3528596>
- [47] Z. Yang, J. Shi, M. H. Asyrofi, B. Xu, X. Zhou, D. Han, and D. Lo, "Prioritizing speech test cases," 2023. [Online]. Available: <https://arxiv.org/abs/2302.00330>
- [48] D. H. X. Yuen, A. Y. C. Pang, Z. Yang, C. Y. Chong, M. K. Lim, and D. Lo, "Asdf: A differential testing framework for automatic speech recognition systems," in *16th IEEE Conference on Software Testing, Verification and Validation, ICST*. IEEE, 2023.
- [49] M. H. Asyrofi, Z. Yang, and D. Lo, "Crossar++: A modular differential testing framework for automatic speech recognition," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 1575–1579. [Online]. Available: <https://doi.org/10.1145/3468264.3473124>
- [50] Q. Pang, Y. Yuan, and S. Wang, "Mdpfuzz: Testing models solving markov decision processes," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2022. New York, NY, USA: Association for Computing Machinery, 2022, p. 378–390. [Online]. Available: <https://doi.org/10.1145/3533767.3534388>
- [51] M. T. Ribeiro, T. Wu, C. Guestrin, and S. Singh, "Beyond accuracy: Behavioral testing of NLP models with CheckList," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, Jul. 2020, pp. 4902–4912. [Online]. Available: <https://aclanthology.org/2020.acl-main.442>
- [52] Z. Wang, Y. Zhou, M. Qiu, I. Haque, L. Brown, Y. He, J. Wang, D. Lo, and W. Zhang, "Towards fair machine learning software: Understanding and addressing model bias through counterfactual thinking," 2023. [Online]. Available: <https://arxiv.org/abs/2302.08018>
- [53] Z. Yang, M. H. Asyrofi, and D. Lo, "Biasrv: Uncovering biased sentiment predictions at runtime," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 1540–1544. [Online]. Available: <https://doi.org/10.1145/3468264.3473117>
- [54] Z. Yang, H. Jain, J. Shi, M. H. Asyrofi, and D. Lo, "Biasheal: On-the-fly black-box healing of bias in sentiment analysis systems," in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2021, pp. 644–648.
- [55] C. Gong, Z. Yang, Y. Bai, J. Shi, A. Sinha, B. Xu, D. Lo, X. Hou, and G. Fan, "Curiosity-driven and victim-aware adversarial policies," in *Proceedings of the 38th Annual Computer Security Applications Conference*, ser. ACSAC '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 186–200. [Online]. Available: <https://doi.org/10.1145/3564625.3564636>
- [56] Z. Yang, J. Shi, M. H. Asyrofi, and D. Lo, "Revisiting neuron coverage metrics and quality of deep neural networks," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. Los Alamitos, CA, USA: IEEE Computer Society, mar 2022, pp. 408–419.
- [57] L. S. Khoo, J. Q. Bay, M. L. K. Yap, M. K. Lim, C. Y. Chong, Z. Yang, and D. Lo, "Exploring and repairing gender fairness violations in word embedding-based sentiment analysis model through adversarial patches," in *2023 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE Computer Society, 2023.
- [58] Z. Yang, B. Xu, J. M. Zhang, H. J. Kang, J. Shi, J. He, and D. Lo, "Stealthy backdoor attack for code models," 2023. [Online]. Available: <https://arxiv.org/abs/2301.02496>
- [59] C. Gong, Z. Yang, Y. Bai, J. He, J. Shi, A. Sinha, B. Xu, X. Hou, G. Fan, and D. Lo, "Mind your data! hiding backdoors in offline reinforcement learning datasets," 2022. [Online]. Available: <https://arxiv.org/abs/2210.04688>
- [60] B. Shneiderman, "Bridging the gap between ethics and practice: guidelines for reliable, safe, and trustworthy human-centered ai systems," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 10, no. 4, pp. 1–31, 2020.
- [61] A. Paleyes, R.-G. Urma, and N. D. Lawrence, "Challenges in deploying machine learning: A survey of case studies," *ACM Comput. Surv.*, vol. 55, no. 6, dec 2022. [Online]. Available: <https://doi.org/10.1145/3533378>

- [62] N. Humbatova, G. Jahangirova, G. Bavota, V. Riccio, A. Stocco, and P. Tonella, "Taxonomy of real faults in deep learning systems," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, ser. ICSE '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1110–1121.
- [63] F. Thung, S. Wang, D. Lo, and L. Jiang, "An empirical study of bugs in machine learning systems," in *2012 IEEE 23rd International Symposium on Software Reliability Engineering*, 2012, pp. 271–280.
- [64] S. Wang and D. Lo, "Version history, similar report, and structure: Putting them together for improved bug localization," in *Proceedings of the 22nd International Conference on Program Comprehension*, ser. ICPC 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 53–63. [Online]. Available: <https://doi.org/10.1145/2597008.2597148>
- [65] T. F. Bissyandé, D. Lo, L. Jiang, L. Réveillère, J. Klein, and Y. L. Traon, "Got issues? who cares about it? a large scale investigation of issue trackers from github," in *2013 IEEE 24th International Symposium on Software Reliability Engineering (ISSRE)*, 2013, pp. 188–197.
- [66] R. Kallis, A. Di Sorbo, G. Canfora, and S. Panichella, "Predicting issue types on github," *Science of Computer Programming*, vol. 205, p. 102598, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167642320302069>
- [67] 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)*, 2019, pp. 406–409.
- [68] 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.
- [69] J. Wang, X. Zhang, L. Chen, and X. Xie, "Personalizing label prediction for github issues," *Information and Software Technology*, vol. 145, p. 106845, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0950584922000192>
- [70] H. Hata, N. Novielli, S. Baltes, R. G. Kula, and C. Treude, "Github discussions: An exploratory study of early adoption," *Empirical Software Engineering*, vol. 27, no. 1, pp. 1–32, 2022.