

EGAD: A moldable tool for GitHub Action analysis

Pablo Valenzuela-Toledo¹, Alexandre Bergel², Timo Kehrer¹, Oscar Nierstrasz³

¹*Software Engineering Group, University of Bern, Bern, Switzerland*

²*RelationalAI, Bern, Switzerland*

³*feenk GmbH, Wabern, Switzerland*

Abstract—GitHub Actions (GA) enjoy increasing popularity in many software development projects as a means to automate repetitive software engineering tasks by enabling programmable event-driven workflows. Researchers and developers typically analyze GA at the raw data level using batch tools to mine and analyze actions, jobs, and steps within GA workflows. Although this approach is widely applicable, it ignores the specific context of the GA workflow domain. Consequently, researchers and developers do not reason directly about the domain abstractions.

We present our preliminary steps in building EGAD (Explorable GitHub Action Domain Model), a moldable domain-specific tool to depict and analyze detailed GA workflow data. EGAD consists of an explorable domain model of GA workflows augmented with custom, domain-specific views, and live narratives. We illustrate EGAD in action using it to explore “sticky commits” in GitHub repositories.

Index Terms—GitHub Actions, software evolution, moldable development

I. INTRODUCTION

GitHub Actions (GA) have been increasingly adopted in software development projects [1]. Although GA was publicly released only in November 2019, it is already the dominant continuous integration service on GitHub [2]. GA allows actions to be triggered automatically based on events such as commits, comments, issues, pull requests, and schedules. For example, GA enables the automation of testing, code reviews, continuous integration, communication, dependency management, and security monitoring.

GA supports automation through *workflows*. Workflows are specified in YAML files that describe the actions to be triggered by specific events (*e.g.*, the automatic treatment of pull requests [3]). However, we observe that developers make frequent mistakes in implementing the workflows and commit multiple changes before running them correctly [4]. The only tool they have for specifying workflows is the text editor, which does not provide features for authoring, analyzing or debugging. Therefore, developers are forced to validate their workflows by pushing multiple versions to the repository.

To conduct studies into GA, researchers work at the raw data level using batch tools to mine and analyze data of interest [1], [3], [4] (see section III). Although this approach is widely applicable, it ignores the contextual nature of the mining and does not provide explorable domain models [5].

To support researchers and developers in studying GA workflows, we present Explorable GitHub Action Domain Model (EGAD), a domain-specific tool to depict and analyze GA workflows and their evolution. EGAD offers an explorable

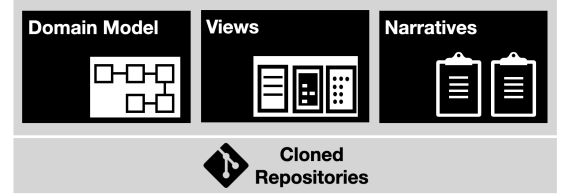


Fig. 1. Architecture overview of EGAD.

domain model, custom views and live narratives that enable researchers and developers to inspect and analyze workflows.

EGAD provides an approach to answer specific questions about GA workflows. The procedure consists of four steps, namely: (i) goal definition, (ii) wrapping the GA workflow data in domain model entities, (iii) exploring the domain model, and (iv) interpreting the results.

We give an overview of EGAD’s architecture in section II, and illustrate its use by investigating so-called *sticky commits* in GitHub software repositories in section III.

II. TOOL ARCHITECTURE AND IMPLEMENTATION

Figure 1 depicts the architecture of EGAD which consists of (i) an explorable domain model, (ii) custom views, and (iii) live narratives that access the raw repository data.

A. Domain model

The domain model wraps the GA workflow data. The model includes: (i) the implementation of the workflow history, including all the commits associated with a GA workflow, and (ii) the representation of the workflow, containing components such as events, jobs and steps.

A class diagram focusing on the most important entities is shown in Figure 2. A `WEHistories` instance contains a collection of `WEHistory` objects, each of which represents the history of workflows of a dedicated project. A `WEFileCommit` object represents a commit revising a dedicated workflow by a new version (`WEWorkflow`), which in turn consists of events, jobs and steps (`WEEvent`, `WEJob` and `Step`).

As we shall see in the next section, we populate an explorable domain model by extracting data from GA workflow files obtained from cloned software repositories.

B. Custom views

Custom views enable the exploration and navigation of the domain model from multiple perspectives. Instead of

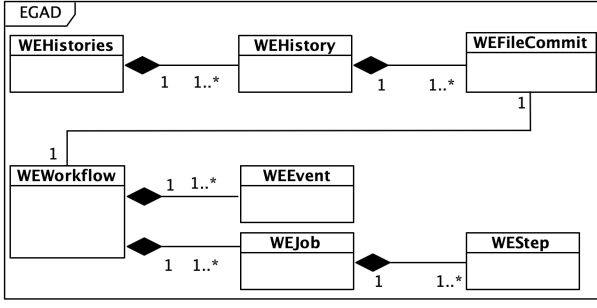


Fig. 2. Representation of GA workflows and their evolution.

presenting the data generically, custom views provide critical insights into the domain model, and make it possible to pose questions and explore hypotheses. In order to answer domain-specific questions, the views can be easily extended. This feature allows the tool to be adapted to face new scenarios to explore. We shall see several examples in the next section.

C. Narratives

Narratives link documentation, source code, and running objects to specific questions or hypotheses of the domain model. We use narratives to explain use cases, scenarios, and features of the GA domain. The creation of narratives is accomplished by (i) telling stories, or (ii) following a path of custom views.

D. EGAD implementation

EGAD is developed on top of the Glamorous Toolkit¹ (GT), a moldable environment for building live and explorable domain models using hyperlinked notebooks, live domain objects, and customizable, domain-specific views [6]. GT is built on Pharo,² a modern, open-source Smalltalk environment.

III. EGAD IN ACTION

In this demonstration, we will work with a dataset comprising selected GitHub repositories that (i) currently use GA and contain at least one workflow file in the `.github/workflow` directory, (ii) have been created after 2019-01-11 (GA official release date) and before 2022-12-14, and (iii) have at least ten stars and 500 commits. We used the GitHub Search tool to select repositories (excluding forks) meeting these criteria [7]. We cloned the repositories on 2022-12-14. For the purpose of this paper, we limit our dataset to the first 50 repositories found.

For the sake of analysis and exploration, EGAD supports a four-step process: (i) goal definition, (ii) wrapping the GA workflow data in domain model entities, (iii) inspecting the domain model, and (iv) interpreting the results.

We illustrate the EGAD approach in the context of sticky commits. Figure 3 introduces the sticky commits narrative in a single notebook page.

¹<https://gtoolkit.com>

²<https://pharo.org>

Sticky Commits narrative

Our goal is to detect one or more sequences of sticky commits through the exploration of the history of the `pythonpackage.yml`.

We wrap the history of a workflow in a `WEHistory`

```
history := WEHistory> fromRepoPath: '../repositories/rich'
forYML: 'pythonpackage.yml'
```

Similarly, we replicate the Sticky Commits candidates analysis considering all the GitHub Actions workflows in a set of repositories.

First, we collect all the repositories:

```
FileLocator> imageDirectory> parent / 'repositories'
```

Then, we create a collection of workflow histories:

```
histories := WEHistories> fromRepoPath: '../repositories'
```

We wrap these histories in the next object example:

```
WEHistoriesExamples>>#historiesExample>
```

a WEHistories

Workflow Histories	Initialized Histories	Raw	Print	Meta
Repo	YML file			Size
../repositories/OpenBBTerm	gh-pages.yml			14
../repositories/OpenBBTerm	draft.yml			3
../repositories/OpenBBTerm	test.yml			165
../repositories/OpenBBTerm	issue.yml			2
../repositories/OpenBBTerm	m1_macos_build.yml			8
../repositories/OpenBBTerm	windows10_build.yml			9
../repositories/OpenBBTerm	mac.yml			15
../repositories/OpenBBTerm	docker.yml			6

Fig. 3. The sticky commits narrative.

Sticky commits consist of a sequence of commits performed immediately, one after another. The commits are made by the same developer, and are intended to correct errors in workflow specifications made in the previous commit. Sticky commits suggest that the developer is having difficulty correctly implementing some GA behavior. We call them “sticky commits” because they are continuously pushed, as if they were sticky.

As an example of sticky commits, let’s consider the history of the GA workflow file `pythonpackage.yml` [8]. This file is part of the `Rich` GitHub repository [9]. The history of this file consists of seventy commits.

A. Goal definition

Our goal is to detect one or more sequences of sticky commits through the exploration of the history of the `pythonpackage.yml` file.

B. Wrapping the GA workflow data in the domain model

We wrap the GA workflow data in a `WEHistory` object. Figure 3a shows the associated code snippet. Executing this snippet yields a `WEHistory` object that we can explore.

The *Workflows* view of this object lists the GA workflow files of the `Rich` repository (Figure 4a). Each custom view, such as this one, has been developed in just a few lines of code, to offer a new, dedicated perspective of the domain objects.

(a) Workflow view

Icon	Name	Size	Creation
📁	..	--	2023-01-15 22:48:57
📄	codespell.yml	333 B	2023-01-15 22:48:57
📄	comment.yml	531 B	2023-01-15 22:48:57
📄	newissue.yml	760 B	2023-01-15 22:48:57
📄	pythonpackage.yml	1.52 kB	2023-01-15 22:48:57
📄	readmechanged.yml	902 B	2023-01-15 22:48:57

(b) Commits view

Index	Author	Timestamp	Delta	Comment	Category
1	Jeremy Howard	2019	na	Initial commit	nil
2	Jeremy Howard	2020	274:23:37.12	Create docs.yml	nil
3	Jeremy Howard	2020	0:00:03:15	Update docs.yml	nil
4	Jeremy Howard	2020	0:00:37:15	Update docs.yml	nil
5	Jeremy Howard	2020	0:02:41:53	Update docs.yml	nil
6	Jeremy Howard	2020	0:03:29:06	Update docs.yml	nil
7	Jeremy Howard	2020	0:00:07:29	Update docs.yml	nil
8	Jeremy Howard	2020	0:00:07:50	Update docs.yml	nil
9	Jeremy Howard	2020	0:00:08:22	Update docs.yml	nil

(c) Sticky Group view

Index	Author	Start	End	Number of commits
1	Will McGugan	2020-07-27T13:33:51-03:00	2019-11-10T12:28:10-03:00	38
2	Hedy Li	2020-08-21T06:22:34-03:00	2020-08-21T06:20:36-03:00	3
3	Will McGugan	2020-10-09T11:09:19-03:00	2020-10-09T10:59:19-03:00	2
4	Sondre Lillebø Gundersen	2020-11-30T05:43:53-03:00	2020-11-28T10:51:12-03:00	2
5	Will McGugan	2021-05-19T13:02:53-03:00	2021-05-19T12:50:09-03:00	4
6	Sondre Lillebø Gundersen	2021-05-24T09:11:42-03:00	2021-05-24T09:11:04-03:00	2
7	Yash Rathi	2021-09-15T00:13:51-03:00	2021-09-14T14:31:42-03:00	2
8	Will McGugan	2022-02-28T14:15:09-03:00	2021-10-06T05:07:09-03:00	3
9	Darren Burns	2022-08-09T09:47:26-03:00	2022-07-14T11:12:50-03:00	9
10	Dave Pearson	2022-10-11T07:00:49-03:00	2022-10-10T11:47:23-03:00	3

Fig. 4. Different views of a WEHistory object: (a) Workflow view, (b) Commits view, (c) Sticky Group view.

The *Commits* view lists all the workflow commits in chronological order (Figure 4b), of which only the first seven commits are shown in the figure.

We now move to the *Sticky Groups* view, which lists groups of sticky commits (Figure 4c). Sticky groups are computed as a series of commits that are (i) related to GA workflows, (ii) consecutive, and (iii) performed by the same author.

C. Inspecting the domain model

To analyse the presence of sticky commits, we inspect the first group, which has 38 commits (Figure 4c). As a result, we get a WESTickyGroup object (Figure 5) listing all the commits within the group.

Inspecting the WESTickyGroup object (Figure 5), we notice in the *Duration* column that the time between commits 4 through 9 is always less than 8 minutes. This is of interest since we have a sequence of commits performed repetitively in a short period of time.

To investigate this sequence of commits in depth, we navigate through each commit and review the changes. Each commit is a WEFileCommit object (Figure 6). This object shows the code that was modified (the *Diff* view), highlighting added and deleting lines of code in green and red, respectively.

D. Interpreting the results

Reviewing each commit modification, we realize that these commits correspond to minor modifications that could have been handled in a single commit. To record this particular behaviour, we use the editable *Category* column, entering “yes”, which means there are sticky commit candidates that could be empirically validated (Figure 5).

Index	Author	Date time	Duration	Comment	Category
1	Will McGugan	2019-11-10T12:28:10-03:00	na	Initial commit	nil
2	Will McGugan	2019-12-25T12:27:53-03:00	44:23:59.43	GH action to run tests on push	nil
3	Will McGugan	2019-12-25T12:37:28-03:00	0:00:09:35	use mypy	nil
4	Will McGugan	2019-12-25T12:39:21-03:00	0:00:01:53	disable typecheck for now	yes
5	Will McGugan	2019-12-25T12:41:49-03:00	0:00:02:28	add pytest-cov	yes
6	Will McGugan	2019-12-25T12:45:12-03:00	0:00:03:23	install with poetry	yes
7	Will McGugan	2019-12-25T12:50:04-03:00	0:00:04:52	don't use makefile	yes
8	Will McGugan	2019-12-25T12:57:26-03:00	0:00:07:22	add pip install	yes
9	Will McGugan	2019-12-25T12:59:14-03:00	0:00:01:48	no can do 3.5	yes
10	Will McGugan	2019-12-27T07:08:22-03:00	1:18:09:08	add typecheck to CI	nil

Fig. 5. Commit view of the WESTickyGroup object.

Details	Changes	Diff	Diff (two panes)	Raw	Print	Meta
name: Python package						
on: [push]						
jobs:						
build:						
runs-on: ubuntu-latest						
strategy:						
max-parallel: 4						
matrix:						
python-version: [3.5, 3.6, 3.7, 3.8]						
steps:						
- uses: actions/checkout@v1						
- name: Set up Python \${{ matrix.python-version }}						
uses: actions/setup-python@v1						
with:						
python-version: \${{ matrix.python-version }}						
- name: Install dependencies						
run:						
python -m pip install --upgrade pip						
pip install -r requirements.txt						
- name: Typecheck with mypy						
run:						
pip install mypy						
# make typecheck						
make typecheck						
- name: Test with pytest						
run:						
pip install pytest						
make test						

Fig. 6. Diff view of the WEFileCommit object.

E. What's next?

Are sticky commits pervasive in workflows of other repositories? To answer this question, we replicate the analysis considering all the GA workflows in a set of repositories. A broader view of the presence of sticky commits in GA workflows may guide the refinement of the heuristics for computing sticky groups.

We replicate the analysis by wrapping the histories of GA workflows from multiple repositories in the WEHistories object (Figure 3b). To make this behavior easier to access, we create the historiesExample example object in our narrative (Figure 3c). The historiesExample object lists all the histories from all the workflows of the 50 repositories.

To investigate the presence of sticky commits in these histories we inspect the test.yml workflow file from the OpenBBTerminal repository. The history of this file has 165 WEFileCommit (Figure 3c) and 32 WESTickyGroup objects (Figure 7a).

Then, we inspect the WESTickyGroup with index 14 (Figure 7a), which contains 11 WEFileCommit objects (Figure 7b). We notice in the *Duration* column that the time between commits 3 through 5, and between commits 7 through 11 is always less than 8 minutes. As with our initial exploration, this is of interest since we have a sequence of sticky commit candidates.

To address our finding, we navigate the last five WEFileCommit objects of the list using the *Diff* view (Figure 7c). We observe a sequence of commits performing

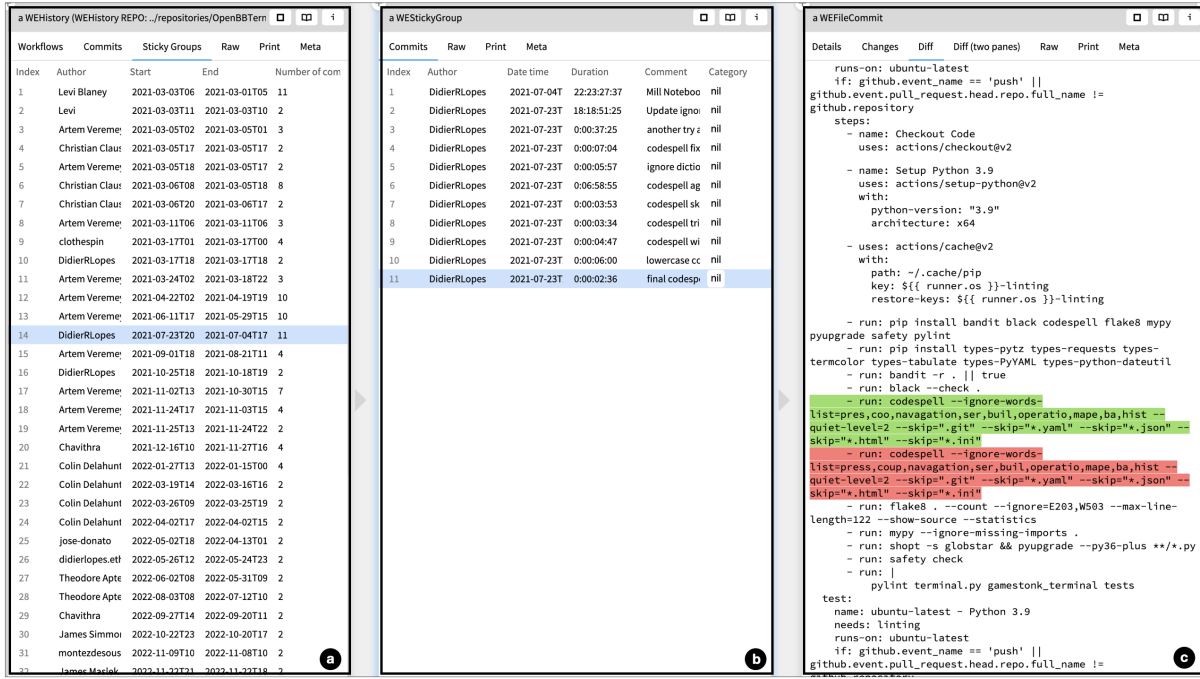


Fig. 7. Following a path of views for WEStickyGroup object inspection.

minor modifications one after another, within a short period of time. Although not a confirmation, this finding suggests that the presence of sticky commits is plausible in the context of our narrative.

Then, as a next step, we could be interested in inspecting the workflow structure in detail and, for example, being able to check the implementation of events and jobs in the workflow. This is a feature that will be implemented in the next release of the tool.

We provide a full replication package [10] which includes documentation, installation instructions, and an example dataset for running the tool.

IV. RELATED WORK

Seminal works from earlier days of mining software repositories relied on artifacts that have been designed to answer specific research questions [11], [12], while the reusability of tools has been discussed largely at an architectural level [13], [14], proposing a blackboard architectural style which we adopt in EGAD.

To date, the well-known GHTorrent [15], [16] and Boa [17] provide queryable offline mirrors to deal with massive amounts of GitHub data. Other tools aim at helping developers in curating their own datasets [18], providing facilities for sampling [7], filtering [19], or finding similar software projects [20]. Besides, there are initiatives for providing archives of these data, fostering long-term preservation and reproducibility [21], [22]. All of these efforts are orthogonal to ours in the sense that EGAD can be combined with any of them for populating its explorable domain model.

Next to curating large datasets serving as general research platforms, a number of tools aim at supporting specific mining

tasks [20], [23], [24], potentially making use of bots to extract the relevant information from software repositories [25]–[27]. The development of most of these tools has been initiated prior to the public release of GA in November 2019, and the analysis of GA workflows has not been considered.

Recently, researchers have started to mine GA using batch tools at the raw data level. Kinsman *et al.* [3] studied how developers use GA, creating a dataset that they later processed to carry out their study. Decan *et al.* [1] investigated trends and adoption patterns of GA through the generation of multiple datasets for answer specific research questions. Similarly, Valenzuela-Toledo and Bergel [4] investigated the evolution of GA workflows, generating a single dataset that they then used to conduct their research. In our work, we shift the focus by introducing a reusable and extensible tool that provides explorable domain models to conduct GA mining studies.

V. CONCLUSION AND FUTURE WORK

We have proposed EGAD, a moldable domain-specific tool. We have illustrated our approach by means of the sticky commit narrative. Our tool enables the study of GA through a domain model, and provides an approach to replicate empirical studies on this subject. Our efforts aid in the representation and analysis of GA workflow data, and the creation of an expandable bench for future research.

In future work, we plan to extend EGAD. We envision the visualization of the evolution of GA to enable the discovery of unseen patterns and the automated detection of specific commit groups such as the sticky ones. Knowing the types of commit groups made by developers can guide the development of tools to help developers being more efficient in producing correct workflow specifications.

REFERENCES

- [1] A. Decan, T. Mens, P. R. Mazrae, and M. Golzadeh, "On the use of GitHub actions in software development repositories," in *2022 IEEE International Conference on Software Maintenance and Evolution (IC-SME)*. IEEE, 2022, pp. 235–245.
- [2] M. Golzadeh, A. Decan, and T. Mens, "On the rise and fall of CI services in GitHub," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2022, pp. 662–672.
- [3] 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.
- [4] P. Valenzuela-Toledo and A. Bergel, "Evolution of GitHub action workflows," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2022, pp. 123–127.
- [5] A. Chis, T. Gırba, J. Kubelka, O. Nierstrasz, S. Reichhart, and A. Syrel, "Moldable tools for object-oriented development," in *Present and Ulterior Software Engineering*. Springer, 2017, pp. 77–101.
- [6] O. Nierstrasz and T. Gırba, "Making systems explainable," in *2022 Working Conference on Software Visualization (VISSOFT)*. IEEE, 2022, pp. 1–4.
- [7] O. Dabic, E. Aghajani, and G. Bavota, "Sampling projects in GitHub for MSR studies," in *18th IEEE/ACM International Conference on Mining Software Repositories, MSR 2021*. IEEE, 2021, pp. 560–564.
- [8] Textualize. (2019, Oct.) pythonpackage.yml file. [Online]. Available: <https://github.com/Textualize/rich/blob/master/.github/workflows/pythonpackage.yml>
- [9] ——. (2020, Oct.) pythonpackage.yml file. [Online]. Available: <https://github.com/Textualize/rich>
- [10] P. Valenzuela-Toledo, A. Bergel, T. Kehrer, and O. Nierstrasz, "EGAD: a Moldable Tool for GitHub Action Analysis," 1 2023. [Online]. Available: <https://github.com/pavt/egad>
- [11] J. Śliwerski, T. Zimmermann, and A. Zeller, "When do changes induce fixes?" *ACM sigsoft software engineering notes*, vol. 30, no. 4, pp. 1–5, 2005.
- [12] C. Weiss, R. Premraj, T. Zimmermann, and A. Zeller, "How long will it take to fix this bug?" in *Fourth International Workshop on Mining Software Repositories (MSR'07: ICSE Workshops 2007)*. IEEE, 2007, pp. 1–1.
- [13] M. Fischer, M. Pinzger, and H. Gall, "Populating a release history database from version control and bug tracking systems," in *International Conference on Software Maintenance, 2003. ICSM 2003. Proceedings*. IEEE, 2003, pp. 23–32.
- [14] T. Mens, S. Demeyer, M. D'Ambros, H. Gall, M. Lanza, and M. Pinzger, "Analysing software repositories to understand software evolution," *Software evolution*, pp. 37–67, 2008.
- [15] G. Gousios, "The ghtorrent dataset and tool suite," in *2013 10th Working Conference on Mining Software Repositories (MSR)*. IEEE, 2013, pp. 233–236.
- [16] G. Gousios, B. Vasilescu, A. Serebrenik, and A. Zaidman, "Lean ghtorrent: GitHub data on demand," in *Proceedings of the 11th working conference on mining software repositories*, 2014, pp. 384–387.
- [17] R. Dyer, H. A. Nguyen, H. Rajan, and T. N. Nguyen, "Boa: A language and infrastructure for analyzing ultra-large-scale software repositories," in *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 2013, pp. 422–431.
- [18] S. Romano, M. Caulo, M. Buompastore, L. Guerra, A. Mounsif, M. Telesca, M. T. Baldassarre, and G. Scanniello, "G-repo: a tool to support msr studies on github," in *2021 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2021, pp. 551–555.
- [19] N. Munaiah, S. Kroh, C. Cabrey, and M. Nagappan, "Curating github for engineered software projects," *Empirical Software Engineering*, vol. 22, no. 6, pp. 3219–3253, 2017.
- [20] E. Bogomolov, Y. Golubev, A. Lobanov, V. Kovalenko, and T. Bryksin, "Sosed: a tool for finding similar software projects," in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, 2020, pp. 1316–1320.
- [21] R. Di Cosmo and S. Zacchiroli, "Software heritage: Why and how to preserve software source code," in *iPRES 2017-14th International Conference on Digital Preservation*, 2017, pp. 1–10.
- [22] V. Markovtsev and W. Long, "Public git archive: a big code dataset for all," in *Proceedings of the 15th International Conference on Mining Software Repositories*, 2018, pp. 34–37.
- [23] E. Lyulina, A. Birillo, V. Kovalenko, and T. Bryksin, "Tasktracker-tool: A toolkit for tracking of code snapshots and activity data during solution of programming tasks," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, 2021, pp. 495–501.
- [24] N. Sviridov, M. Evtikhiev, and V. Kovalenko, "Tnm: A tool for mining of socio-technical data from git repositories," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 2021, pp. 295–299.
- [25] S. Surana, S. Detroja, and S. Tiwari, "A tool to extract structured data from github," *arXiv preprint arXiv:2012.03453*, 2020.
- [26] T. F. Bissyandé, F. Thung, D. Lo, L. Jiang, and L. Réveillere, "Orion: A software project search engine with integrated diverse software artifacts," in *2013 18th international conference on engineering of complex computer systems*. IEEE, 2013, pp. 242–245.
- [27] A. Abdellatif, K. Badran, and E. Shihab, "Msrbot: Using bots to answer questions from software repositories," *Empirical Software Engineering*, vol. 25, no. 3, pp. 1834–1863, 2020.