

A Preliminary Investigation of MLOps Practices in GitHub

Fabio Calefato
University of Bari
Bari, Italy
fabio.calefato@uniba.it

Filippo Lanubile
University of Bari
Bari, Italy
filippo.lanubile@uniba.it

Luigi Quaranta
University of Bari
Bari, Italy
luigi.quaranta@uniba.it

ABSTRACT

Background. The rapid and growing popularity of machine learning (ML) applications has led to an increasing interest in MLOps, that is, the practice of continuous integration and deployment (CI/CD) of ML-enabled systems. **Aims.** Since changes may affect not only the code but also the ML model parameters and the data themselves, the automation of traditional CI/CD needs to be extended to manage model retraining in production. **Method.** In this paper, we present an initial investigation of the MLOps practices implemented in a set of ML-enabled systems retrieved from GitHub, focusing on GITHUB ACTIONS and CML, two solutions to automate the development workflow. **Results.** Our preliminary results suggest that the adoption of MLOps workflows in open-source GitHub projects is currently rather limited. **Conclusions.** Issues are also identified, which can guide future research work.

KEYWORDS

CI/CD, automated workflows, ML-enabled systems, GitHub Actions, CML

1 INTRODUCTION

ML-enabled systems, that is, software systems incorporating machine learning (ML) models, are receiving more and more attention from both researchers and practitioners [8, 12, 13].

Accordingly, there is also an increasing interest in MLOps [10, 14], that is, the development of solutions for the rapid delivery and deployment of ML models. Because ML models are the core part of larger software systems, MLOps builds on DevOps and GitOps practices [3, 9], and introduces additional actions that are specific to machine learning.

Despite the increasing popularity of MLOps, studies on the adoption of automation and its impact on changes of ML-enabled systems are lacking in the academic literature. As such, in this paper we aim to gather an initial understanding of how MLOps solutions are used to automate the execution of tasks to build and deploy ML-enabled systems, starting from open source projects available in GitHub. In doing so, we make sure to exclude all ML-related projects that cannot be directly classified as ML-enabled systems or components: i.e., we discard all ML libraries, frameworks, and utilities, as they ultimately represent instances of traditional software projects for which the CI/CD practices have been already investigated (e.g., [5]).

Then, we ask the following research questions:

RQ1. *How common is workflow automation in ML-enabled systems hosted on GitHub?*

RQ2. *What type of events are used to trigger MLOps workflows?*

RQ3. *What are the most frequently enacted tasks?*

We answer these questions by means of an explorative archival study that mixes quantitative and qualitative analysis, focusing on GITHUB ACTIONS and CML. The former is a popular continuous integration and continuous delivery (CI/CD) platform, backed by GitHub, to automate software workflows,¹ while the latter is a complementary solution to support continuous CI/CD practices that are specific to ML-enabled systems.²

The remainder of the paper is structured as follows. In Sect. 2, we briefly illustrate how workflow automation works in GITHUB ACTIONS and CML. In Sect. 3, we present the design of the study, namely how we built the experimental datasets and the analyses conducted. The results are presented in Sect. 4 and their limitations in Sect. 5. We discuss the findings and draw conclusions in Sect. 6.

2 WORKFLOW AUTOMATION WITH GITHUB ACTIONS AND CML

GITHUB ACTIONS³ is an event-driven API that automates development workflows in GitHub. Workflows are defined by storing YAML files (e.g., `build.yml` or `deploy.yml`) checked-in to a project's `.github/workflows/` directory, which will typically run when triggered by an *event* occurring in the repository, such as a developer creating a pull request, opening an issue, or pushing a commit to a repository.

A workflow contains one or more jobs. A *job* is a set of steps that are executed in order on the same dedicated runner (typically, a container or a virtual machine) and are dependent on each other. Each step lets developers either execute a script or run an *action*, a custom application for the GITHUB ACTIONS platform that performs a complex but frequently repeated task (e.g., set up the correct toolchain for a build environment, set up the cloud provider authentication, etc.). As such, using actions helps reduce the amount of repetitive code written in workflow files. In addition, actions can be even published on the GitHub Marketplace for others to reuse.

Figure 1 shows an example of GITHUB ACTIONS as a YAML file (Figure 1a) and as graphically displayed on GitHub (Figure 1b). The workflow is triggered on push events on the repository; it contains one job (i.e., `check-bats-version`) consisting of four consecutive steps, to be executed on an Ubuntu container as the chosen runner: step (1) and (2) use actions that, respectively, check out the source code and set up a NodeJS ver. 14 environment; then, step (3) runs a

¹<https://github.com/features/actions>

²<https://cml.dev/doc>

³<https://docs.github.com/en/actions/learn-github-actions/understanding-github-actions>

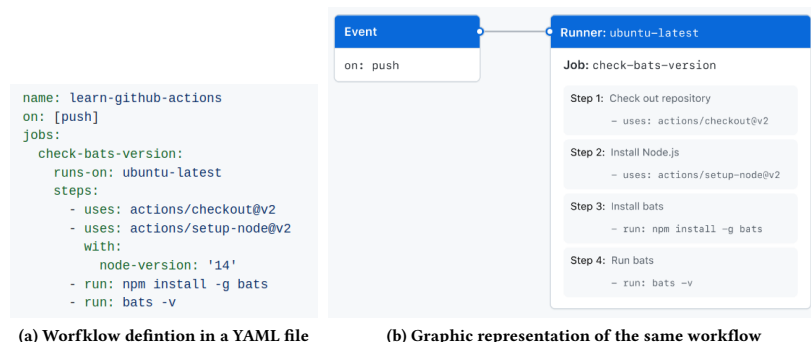


Figure 1: An example of a GitHub Actions workflow.

custom command to install the bats library via the npm package manager; finally, step (4) checks the version of the installed library.

CML⁴ is an open-source, command-line interface tool for implementing CI/CD in machine learning projects hosted on GitHub or GitLab. In GitHub, CML builds upon the GitHub Actions infrastructure, therefore the way it works is fairly similar; in fact, developers must define jobs within a `cml.yml` file to be stored in the same `.github/workflows/` directory. A typical CML workflow can comprise shell commands, custom scripts, and specific CML *commands*. CML commands can be used, for instance, to commit changes to a new branch and open a Pull Request (PR), add comments to a PR, and start a runner.

In the rest of the paper, we use the term *task* as an abstraction of Action in GitHub Actions and command in CML.

3 STUDY DESIGN

In this section, we describe the design of our study. First, we describe the process adopted for building the two experimental datasets of MLOps workflows in GitHub, specific for GitHub Actions and CML. Then, we provide details about the data analysis method.

3.1 Dataset Construction

3.1.1 Identification of GitHub Actions Workflows. The first step in our study was assembling an experimental dataset of GitHub Actions workflows taken from ML-related repositories. We decided to focus on GitHub Actions as it is the default workflow automation tool offered by GitHub, the largest and the most popular code hosting platform for open-source projects. Figure 2a summarizes the steps of the dataset creation process.

We started from the dataset used in a recent work on the use of GitHub Actions [7], comprising 446,862 GitHub repositories.

In GitHub, users mark with a *star* the repositories they like or are interested to; these become part of the list of their favorites, aka their *stars*. In the GitHub jargon, users that mark a repository with a star are known as the repository *stargazers*. Many studies consider the number of stargazers to be evidence of repository popularity and – as such – a viable proxy of their quality [2, 11]. Hence, as a measure to keep low-quality projects out of our dataset, we decided

to analyze only repositories having more than one star on GitHub (i.e., 183,127 out of the 446,862 repositories from the original set).

Moreover, Kalliamvakou et al. showed that a large portion of GitHub repositories are typically found to be inactive [6]. In line with the goals of our study, we decided to consider only repositories that had been active for at least six months after the public release of GitHub Actions – happened in November 2019. Specifically, we filtered out all repositories whose last-commit date preceded May 2020. In addition, to restrict our selection to ML-related repositories, we adopted a strategy similar to the one used by Biswas et al. in [2]. Specifically, we leveraged the GitHub API to retrieve the descriptions and topics of the repositories resulting from the filtering above. Then, we sought ML systems-related keywords contained therein (e.g., ‘*machine learning*’, ‘*deep learning*’, ‘*neural network*’, ‘*image processing*’, etc.)⁵ Roughly, we used the same set of keywords adopted by Biswas et al., but also added a couple of noticeable absentees (i.e., ‘*AI*’ and ‘*prediction model*’) and removed the subset of keywords referred to specific data science technologies (e.g., ‘*spark*’, ‘*hadoop*’, and ‘*cafe*’). After applying the keyword-based filtering, we obtained a set of 2,516 repositories.

For each candidate repository, we then checked the adoption of GitHub Actions by verifying the existence of YAML files (with extension `.yaml` or `.yml`) in the `.github/workflows` directory. Out of the 2,516 selected repositories, only 155 contain at least one GitHub Actions workflow. Contextually, by leveraging the GitHub API, we downloaded all the available 399 workflows: almost all of them were valid YAML files, with only 2 exceptions.

In conclusion, our first dataset – hereafter referred to as the *GitHub Actions dataset* – consists of 397 valid workflows extracted from 155 GitHub repositories.

3.1.2 Identification of CML Workflows. To build the second dataset of projects containing CML workflows, we used the advanced search web interface to search globally across all of GitHub for repositories containing `.github/workflows/cml.yml` files. We decided to focus on CML because it is designed as a natural, ML-oriented extension of GitHub Actions and it is also tightly integrated with DVC,⁶ one

⁴<https://cml.dev>

⁵A complete list of the keywords we considered is available in the file `/settings.json` as part of the replication package of this study.

⁶<https://dvc.org>

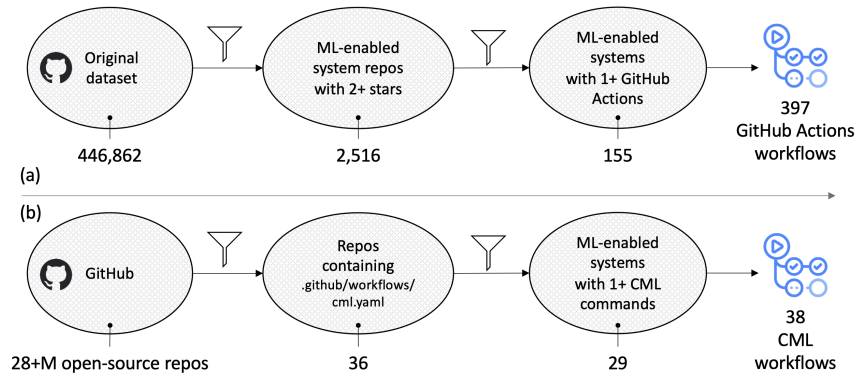


Figure 2: Steps applied to create the two datasets of GitHub ACTIONS (a) and CML (b) workflows.

of the most popular data version control tools available to date [1]. Figure 2b illustrates the steps for the dataset building process.

The query returned a list of only 36 candidates. Given the few hits, in this case we did not apply any further filtering (e.g., number of stars as a proxy of quality). Still, two authors manually vetted the list to exclude non-relevant results. In particular, six repositories were discarded because they contain no ML models (2), no code (3), and an empty workflow file (1); we also discarded one repository that is merely the implementation of one of the tutorials available on the CML website.

The complete list of repositories is available as supplemental material in the replication package.⁷ We then used the GitHub API to download GitHub ACTIONS workflows from the retained repositories.

In conclusion, our second dataset – hereafter referred to as the *CML dataset* – contains 41 CML workflows extracted from 29 GitHub repositories, of which 38 valid.

3.2 Data Analysis Method

To answer RQ1 (i.e., “How common is workflow automation in ML-enabled systems hosted on GitHub?”), we characterized the general adoption of GitHub ACTIONS and CML in our datasets by computing the number of repositories containing related workflows as well as the median number of workflows per repository. Then, two of the paper authors engaged in a manual inspection of the selected repositories to verify that they actually qualify as ML-enabled systems, discarding false positives from the filtering process.

To answer RQ2 (i.e., “What type of events are used to trigger MLOps workflows?”), we computed the frequency of GitHub events used to trigger the workflows in our dataset.

Regarding RQ3 (i.e., “What are the most frequently enacted tasks?”), to infer the most frequent tasks, we sought patterns among recurring Actions and shell commands from the collected workflows. Specifically, with the term ‘task’ here we refer to any logical unit of work that is typically found within an ML development workflow (e.g., ‘launching a model-retraining job,’ ‘deploying an ML-enabled component to production,’ or ‘saving a newly created model into

a model registry’). Within GitHub ACTIONS, such tasks can be implemented by applying one or more pre-defined actions, by running a sequence of shell commands, or even using a combination of both. Therefore, to uncover patterns that may possibly constitute high-level tasks, we started by analyzing the most frequently used actions. Accordingly, we first computed the descriptive statistics of the actions found in our datasets, while also identifying the occurrence of those containing the words ‘cml’ or ‘docker’ in their slugs; indeed, we consider the use of these technologies a trace of typical MLOps practices (experiment tracking and ML models deployment, respectively). Then, we leveraged the *Apriori* algorithm to compute the set of frequently co-occurring actions; to do so, we grouped action slugs (stripped of their version tags) by the workflows in which they appear, thus generating a collection of transactions.⁸

As in GitHub ACTIONS the steps of a workflow job may consist of either actions (denoted by the attribute *uses*) or shell commands/scripts (denoted by *run*), we applied a similar analytical procedure – including the application of the *Apriori* algorithm – to particular subsets of shell commands from the run steps. Specifically, we studied the (co-)occurrence of shell commands containing the keywords *cml* and *docker*.

Finally, we performed a qualitative analysis of the workflows by manually inspecting the most interesting ones, i.e., those belonging to repositories that contain ML-enabled software components.

4 RESULTS

In this section, we present and discuss the results of our preliminary investigation of MLOps practices in GitHub.

4.1 RQ1. How common is workflow automation in ML-enabled systems hosted on GitHub?

4.1.1 Results from the GitHub ACTIONS dataset. Our GitHub ACTIONS dataset comprises 397 valid GitHub Actions workflows distributed across 155 repositories. The median number of workflows

⁸In frequent pattern mining, a transaction is defined as a collection of items that have been observed together; in our case, each transaction corresponds to the collection of actions observed together in the context of a single workflow.

⁷<https://github.com/collab-uniba/mlops4aisystems>

contained in each repository is 2, while the third quartile of the distribution is 3 and the maximum number of workflows found in a repository is 14.

Having retrieved only a limited number of repositories, we could verify the results of the filtering process with a manual inspection. Two authors inspected each repository on GitHub and checked that it actually contained an ML-enabled system or an ML-enabled software component to be integrated in a larger system. The vast majority of the repositories (105) have been misclassified as ML-related projects; the remaining repositories (50) are ML libraries or frameworks (e.g., [dmlc/xgboost](#), [mlpack/mlpack](#), [h2oai/h2o-dev](#)), mostly written in Python or Java. We decided to discard these kinds of repositories as – apart from their ML-oriented goal – they actually look no different from traditional software projects, with which they share the typical engineering requirements and needs (e.g., the use of traditional unit testing, the adoption of code quality assurance tools, etc.)[13]. By looking at the misclassified instances, we noticed that the ‘*machine learning*’ and ‘*artificial intelligence*’ topic labels are often used as buzzwords due to the current hype.

Only one repository can be considered an actual example of ML-enabled system (see [tesseract-ocr/tesseract](#)). However, the related GITHUB ACTIONS workflows are devoted to testing and benchmarking the application as a whole and none of them seems to address any ML task directly.

Given the young age of mature ML techniques, we speculate that most of the professional collaboration around ML-enabled components and systems might currently be happening within companies and private research institutions. In such contexts, high-performance ML models constitute a valuable economical asset and sharing them as open-source software would likely spoil their value. Consequently, we argue that ML-enabled systems may be more frequently found in private GitHub repositories rather than in public ones.

Having not found relevant material for our study in the GITHUB ACTIONS dataset, after our manual inspection, we decided to discard it. As such, the analytical results hereafter presented are exclusively referred to the CML dataset discussed next.

4.1.2 Results from the CML dataset. The CML dataset comprises 38 valid GitHub Actions workflows distributed across 29 repositories. The median number of workflows contained in each repository is 1 and the third quartile of the distribution is still 1; the maximum number of workflows found in a repository is 4.

Also in this case, two authors manually vetted the dataset, individually examining the contents of the retrieved repositories. As a result of this manual inspection, we classified most of the repositories from this collection (24) as *test-driving* (i.e., repositories containing proof-of-concept ML projects, likely used by GitHub users to test-drive the CML library). A couple of repositories are explicitly linked to educational material, i.e., a university ML course ([tue-5ARA0/mlOps-demo-live](#)) and a technical blog post ([amitvulkarni/Bring-DevOps-to-Machine-Learning-with-CML](#)); as such, we classify them as *educational*. Finally, in 3 repositories we found small-size ML projects (i.e., up to three contributors), to which we refer as *ML component*: the first ([Ascend](#)

[NTNU/perception_testing_21](#)) is a workspace for building models for ROS, a robot operating system; the second ([cheesama/morphine](#)) an entity classifier in Korean for the pynori Python library; the third ([mozartofmath/AmharicSpeechToText](#)) a speech-to-text engine for the Amharic language.

Also in this case, we observe a substantial lack of production-grade ML projects containing traces of MLOps practices. However, all repositories from this dataset contain CML workflows, each of which is specifically defined to accomplish one or more MLOps tasks. As such, despite many of them having testing or educational purposes, we deemed all repositories from this dataset worth of being analyzed.

4.2 RQ2. What type of events are used to trigger MLOps workflows?

Most of the workflows from our dataset are triggered by `git push` and `pull_request` events: they are used in 78.95% and 18.42% of the workflows, respectively. Moreover, these events co-occur as triggers in 13.16% of the workflows. Instead, GitHub-managed events are rarely used to activate workflows. Only 3 of them start as a consequence of a new `issue_comment` and just 2 are triggered by GitHub’s release event. The `schedule` event – allowing workflows activation at specific times – is also used only twice.

While it is not surprising that the native `git` events `push` and `pull_request` are the most commonly used triggers for projects hosted on GitHub, we would have expected a prevalence of workflows activated by `pull_request` and `release` events. Indeed, typical MLOps tasks are computationally heavy and time consuming (e.g., model re-training and testing) and, as such, presumably expensive to run every time a new push operation is performed.

4.3 RQ3. What are the most frequently enacted tasks?

To identify potential ML-related tasks accomplished within the collected workflows, we analyzed the pre-defined actions used in each of them. The total number of distinct actions retrieved in the CML dataset is 15 (14 if we ignore version tags). Typically, actions are used slightly more than run commands: the average number of workflow steps using actions over the total number of workflow steps is 0.53. The median number of actions in each workflow is 3 (mean: 2.55, standard deviation: 1.20); no workflow uses more than 5 actions. Of the 14 distinct actions adopted by the workflows in our dataset, 11 (~79%) are available in the GitHub Marketplace and 8 of them are by creators verified by GitHub.⁹ Moreover, we were able to find only one action containing the substring ‘*cml*’ in its slug (i.e., `setup-cml`), while no slug contained the word ‘*docker*’.

By using the *Apriori* algorithm, we computed the frequent sets of actions with support higher than 0.07 (i.e., appearing in at least 3 workflows); the results are reported in Table 1. The top-4 frequent itemsets combine `actions/checkout` with `actions/setup-cml` and `actions/setup-python`. Actually, `actions/checkout` is found in almost every frequent itemset, as it is commonly required in all jobs aimed at processing the contents of a repository.

⁹In the GitHub Marketplace, Actions get this label when they are developed by creators that have an established relationship with GitHub; these Actions are typically the result of the joint work of GitHub and their creators.

Table 1: Frequently co-occurring actions with support > 0.07 (i.e., actions co-occurring in at least 3 workflows).

Ⓢ Action 1	Ⓢ Action 2	Ⓢ Action 3	Support
actions/checkout	iterative/setup-cml		0.47 (18 workflows)
actions/checkout	actions/setup-python		0.37 (14 workflows)
actions/setup-cml	actions/setup-python		0.34 (13 workflows)
actions/checkout	actions/setup-python	actions/setup-cml	0.34 (13 workflows)
actions/checkout	aws-actions/configure-aws-credentials		0.11 (4 workflows)
actions/checkout	actions/setup-go		0.08 (3 workflows)

Table 2: Actions found within CML workflows.

Ⓢ Action	Frequency
actions/checkout	28
iterative/setup-cml	17
actions/setup-python	13
iterative/setup-dvc	2
ros-industrial/industrial_ci	1
aws-actions/configure-aws-credentials	1

If we limit our analysis to the workflows containing ‘*cml*’ in their filename, we find that only 6 distinct actions are used. Table 2 presents each of them together with their frequencies. In the following, we briefly report their definition:

- Ⓢ **actions/checkout** Checks-out the contents of a repository in the virtual environment in which a job is running;
- Ⓢ **iterative/setup-cml** Enables CML functionalities within GITHUB ACTIONS workflows.
- Ⓢ **actions/setup-python** Sets up the requested version of Python, to be used in the subsequent steps of a job.
- Ⓢ **iterative/setup-dvc** Sets up DVC (Data Version Control) functionalities within GITHUB ACTIONS workflows.
- Ⓢ **ros-industrial/industrial_ci** Configures a CI process specifically tailored for packages powered by ROS (Robot Operating System).
- Ⓢ **aws-actions/configure-aws-credentials** Configures AWS credentials and region environment variables to be used for AWS API calls.

Every CML workflow uses `actions/checkout` to make the contents of the repository available in the Actions runner; then, about half of them set up CML and Python using pre-defined actions. Noticeably, `iterative/setup-cml` is not an action verified by GitHub and is not yet available in the Marketplace. We observe that none of these actions can be used alone to perform a high-level MLOps task. Instead, they are primarily used to set up the CI environment and accomplish preliminary configuration steps.

After characterizing the use of actions in our dataset, we delved into the analyses of shell commands extracted from `run` attributes. Of the 38 GITHUB ACTIONS workflows from the CML dataset, 28 present shell commands containing ‘*cml*’ as a substring. CML offers a few commands to support the implementation of CI/CD in ML projects; here we report a brief definition of those found in our workflows:

- Ⓢ **send-comment** adds a Markdown report as a comment on a commit or pull/merge request.
- Ⓢ **publish** uploads and publicly hosts an image (e.g., a .png plot) for displaying in a CML report.
- Ⓢ **runner** starts a runner in any supported cloud compute provider (or locally, in case of an on-premise CML deployment).
- Ⓢ **tensorboard-dev** returns a link to `tensorboard.dev`, a managed TensorBoard¹⁰ platform that lets users upload and share their ML experiment results.

All 28 workflows containing a CML command use `send-comment` to decorate a commit or pull-request message with model training and evaluation metrics. In 17 cases (~63%), the `publish` command is used along with `send-comment` to add an image to the Markdown-formatted message. Only in 2 distinct cases we observed the use of `runner` and `tensorboard-dev`. As the `send-comment` and `publish` are used just for reporting purposes, we find particularly interesting the adoption of `runner` in (`ibrahimkaratas88/cml_cloud_try`); by manually inspecting the workflow, we found that it is used to set up an AWS virtual machine to be used as a self-hosted runner in a subsequent job. In the following, the workflow uses DVC to reproduce the data-transfer task on the cloud machine and report the related results via the `send-comment` command. However, as far as we can tell, model training or re-training tasks are not involved in this example.

In search for end-to-end MLOps pipelines, we also sought the use of Docker commands. Within our dataset, Docker is used just once (in `2796gaurav/automate`, classified as a *test-driving* repository). Specifically, the `build` and `push` Docker commands are used in `dockerize.yml` to build and upload a Docker image to AWS Elastic Container Registry using Terraform, a multi-cloud tool to implement infrastructure as code. This workflow is executed only when a message containing the string ‘*dockerize*’ is published as a comment in an issue or pull request.

Along with this and a few other GITHUB ACTIONS workflows, the same repository contains a CML workflow, named `cml_report.yml`, which uses DVC to reproduce a full ML pipeline (from data preparation to model evaluation) and then makes training results available on GitHub via the `cml-send-comment` command.

¹⁰TensorBoard is a visualization tool for tracking ML experiments. It is typically used in combination with TensorFlow, a popular open-source experimentation platform for ML.

5 LIMITATIONS

In this section, we illustrate the main limitations of our study and our plans to overcome them. For the sake of transparency and repeatability, we have made our scripts and data publicly available.

Being based on the analysis of a very limited sample of ML-enabled projects, in its current form, our study lacks external validity. To overcome this, we plan on expanding our work along different paths. Upon inspecting the sample of open-source GitHub repositories from the dataset used in [7], we could not find any relevant project for our study. Nevertheless, we believe that other datasets assembled from GitHub might be worth investigating as well: for instance, we are currently analyzing the corpus of ML projects gathered by Gonzalez et al. [4]. Moreover, we intend to take into account additional code hosting platforms, such as GitLab and Bitbucket, as they might also represent a valuable source for ML-enabled systems. Furthermore, as already discussed in Sect. 4.1, we argue that many interesting repositories containing ML-enabled components or systems might not be currently available as open-source software. To take into account the adoption of MLOps practices in the industry, we are now designing a couple of case studies to be investigated at specialized companies.

Another limitation of our study is the focus on two particular software solutions (i.e., GITHUB ACTIONS and CML). The landscape of CI/CD tools and MLOps platforms is large and constantly evolving. In a future study, we plan to take into consideration further popular CI/CD solutions, such as GitLab CI/CD, TravisCI, and CircleCI. Also, we will take into account alternatives to CML as a specialized MLOps solution: indeed, despite having been publicly available for about one year now, CML appears to be still scarcely adopted in GitHub and mostly for test-driving purposes.

Finally, in our study, we considered the use of Docker a viable proxy to investigate deployment practices for ML-enabled components. In our future work, we plan on considering additional deployment strategies and tools, e.g., by studying deployment solutions offered by commercial cloud providers and their integration into MLOps workflows.

6 CONCLUSIONS

In this paper, we present a preliminary investigation of MLOps practices in GitHub. Overall – despite the popularity of the code hosting platform – we highlight a substantial lack of open-source projects concerning ML-enabled systems or components which leverage GITHUB ACTIONS; arguably, it is possible that such kind of projects might be hosted as private repositories.

Conversely, many ML-related open-source libraries and frameworks leverage GITHUB ACTIONS. However, these cannot be deemed relevant for the study of MLOps practices as they are not different from traditional software systems, apart from the ML-related goal or use-case.

Searching for projects that leverage CML – a specialized MLOps tool – also yielded a few results. Therefore, researchers interested in studying MLOps practices might want to consider broadening the scope of potential sources by taking into account further code hosting platforms (e.g., GitLab and Bitbucket), CI/CD services (e.g., TravisCI and CircleCI), and specialized MLOps tools.

REFERENCES

- [1] Amine Barrak, Ellis E. Eghan, and Bram Adams. 2021. On the Co-evolution of ML Pipelines and Source Code - Empirical Study of DVC Projects. In *2021 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, Honolulu, HI, USA, 422–433. <https://doi.org/10.1109/SANER50967.2021.00046>
- [2] Sumon Biswas, Md Johirul Islam, Yijia Huang, and Hriday Rajan. 2019. Boa Meets Python: A Boa Dataset of Data Science Software in Python Language. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, Montreal, QC, Canada, 577–581. <https://doi.org/10.1109/MSR.2019.00086>
- [3] Christof Ebert, Gorka Gallardo, Josune Hernantes, and Nicolas Serrano. 2016. DevOps. *Ieee Software* 33, 3 (2016), 94–100. Publisher: IEEE.
- [4] Danielle Gonzalez, Thomas Zimmermann, and Nachiappan Nagappan. 2020. 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*. ACM, Seoul Republic of Korea, 431–442. <https://doi.org/10.1145/3379597.3387473>
- [5] Michael Hilton, Timothy Tunnell, Kai Huang, Darko Marinov, and Danny Dig. 2016. Usage, costs, and benefits of continuous integration in open-source projects. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*. ACM, Singapore Singapore, 426–437. <https://doi.org/10.1145/2970276.2970358>
- [6] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M. German, and Daniela Damian. 2014. The promises and perils of mining GitHub. In *Proceedings of the 11th Working Conference on Mining Software Repositories - MSR 2014*. ACM Press, Hyderabad, India, 92–101. <https://doi.org/10.1145/2597073.2597074>
- [7] Timothy Kinsman, Mairieli Wessel, Marco A. Gerosa, and Christoph Treude. 2021. How Do Software Developers Use GitHub Actions to Automate Their Workflows?. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, Madrid, Spain, 420–431. <https://doi.org/10.1109/MSR52588.2021.00054>
- [8] Grace A. Lewis, Stephany Bellomo, and Ipek Ozkaya. 2021. Characterizing and Detecting Mismatch in Machine-Learning-Enabled Systems. In *2021 IEEE/ACM 1st Workshop on AI Engineering - Software Engineering for AI (WAIN)*. IEEE, Madrid, Spain. <https://doi.org/10.1109/WAIN52551.2021.00028>
- [9] Thomas A. Limoncelli. 2018. GitOps: A Path to More Self-Service IT: IaC + PR = GitOps. *Queue* 16, 3 (June 2018), 13–26. <https://doi.org/10.1145/3236386.3237207> Place: New York, NY, USA Publisher: Association for Computing Machinery.
- [10] Sasu Makinen, Henrik Skogstrom, Eero Laaksonen, and Tommi Mikkonen. 2021. Who Needs MLOps: What Data Scientists Seek to Accomplish and How Can MLOps Help?. In *2021 IEEE/ACM 1st Workshop on AI Engineering - Software Engineering for AI (WAIN)*. IEEE, Madrid, Spain, 109–112. <https://doi.org/10.1109/WAIN52551.2021.00024>
- [11] Nuthan Munaiah, Steven Kroh, Craig Cabrey, and Meiyappan Nagappan. 2017. Curating GitHub for engineered software projects. *Empirical Software Engineering* 22, 6 (Dec. 2017), 3219–3253. <https://doi.org/10.1007/s10664-017-9512-6>
- [12] Nadia Nahar, Shurui Zhou, Grace Lewis, and Christian Kästner. 2021. Collaboration Challenges in Building ML-Enabled Systems: Communication, Documentation, Engineering, and Process. *arXiv:2110.10234 [cs]* (Dec. 2021). <http://arxiv.org/abs/2110.10234> arXiv: 2110.10234.
- [13] Ipek Ozkaya. 2020. What is really different in engineering ai-enabled systems? *IEEE Software* 37, 4 (2020), 3–6. Publisher: IEEE.
- [14] Mark Treveil, Nicolas Omont, Clément Stenac, Kenji Lefevre, Du Phan, Joachim Zentici, Adrien Lavoillotte, Makoto Miyazaki, and Lynn Heidmann. 2020. *Introducing MLOps*. O'Reilly Media.