

The Impact of Code Ownership of DevOps Artefacts on the Outcome of DevOps CI Builds

Ajiromola Kola-Olawuyi
University of Waterloo
akolaola@uwaterloo.ca

Nimmi Rashinika Weeraddana
University of Waterloo
nrweeraddana@uwaterloo.ca

Meiyappan Nagappan
University of Waterloo
mei.nagappan@uwaterloo.ca

ABSTRACT

DevOps is a key element in sustaining the quality and efficiency of software development. Yet, the effectiveness of DevOps methodologies extends beyond just technological expertise. It is greatly affected by the manner in which teams handle and engage with DevOps artefacts. Grasping the intricacies of code ownership and contribution patterns within DevOps artefacts is vital for refining strategies and ensuring they deliver their full potential.

There are two main strategies to manage DevOps artefacts as suggested in prior work: (1) all project developers need to contribute to DevOps artefacts, and (2) a dedicated group of developers needs to be authoring DevOps artefacts. To analyze which strategy works best for Open-Source Software (OSS) projects, we conduct an empirical analysis on a dataset of 892,193 CircleCI builds spanning 1,689 OSS projects. We employ a two-pronged approach to our study. First, we investigate the impact of chronological code ownership of DevOps artefacts on the outcome of a CI build on a build level. Second, we study the impact of the *Skewness* of DevOps contributions on the success rate of CI builds at the project level.

Our findings reveal that, in general, larger chronological ownership and higher *Skewness* values of DevOps contributions are related to more successful build outcomes and higher rates of successful build outcomes, respectively. We further find that projects with low *Skewness* values could have high build success rates when the number of developers in the project is relatively small. Thus, our results suggest that while larger software organizations are better off having dedicated DevOps developers, smaller organizations would benefit from having all developers involved in DevOps.

CCS CONCEPTS

• **Do Not Use This Code** → **Generate the Correct Terms for Your Paper**; *Generate the Correct Terms for Your Paper*; Generate the Correct Terms for Your Paper; Generate the Correct Terms for Your Paper.

KEYWORDS

DevOps, Code Ownership, Continuous Integrations, Empirical Study.

ACM Reference Format:

Ajiromola Kola-Olawuyi, Nimmi Rashinika Weeraddana, and Meiyappan Nagappan. 2024. The Impact of Code Ownership of DevOps Artefacts on the Outcome of DevOps CI Builds. In *21st International Conference on Mining Software Repositories (MSR '24)*, April 15–16, 2024, Lisbon, Portugal. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3643991.3644924>

1 INTRODUCTION

DevOps is a development methodology that bridges the gap between Development (Dev) and Operations (Ops) [33]. It plays a central role in projects, orchestrating effective communication and collaboration throughout the software development process. DevOps, coupled with automated deployment, is made possible through a well-defined set of practices [33]. Among various DevOps practices, *Continuous Integration (CI)* is popular because it is directly associated with the automation of the release cycles of software products. These processes hold immense importance in software development, enabling teams to streamline workflows and deliver high-quality software. In fact, large software organizations like Google and Mozilla invest millions of dollars just for CI services (excluding the cost of developers who maintain CI) [30].

However, implementing CI in software repositories can substantially burden project maintainers with the growth in the number of CI builds. The time to receive feedback from the build process increases due to long queued times, and the operational costs associated with build execution also increase [34]. Therefore, it is imperative to understand better the factors affecting CI systems to optimize release pipelines while lowering operational costs. To that end, many previous studies [38, 44, 71] have been aimed at understanding the factors that affect the outcome of CI builds. For example, change-set size [44], the files changed [71], the day of the week and the time of day [38] are some factors influencing the CI build outcomes. Moreover, a comprehensive understanding of these factors and those affecting the outcomes of CI builds associated with DevOps artefacts (that define DevOps pipelines) provides a competitive advantage to project maintainers. This advantage allows them to expedite the delivery of reliable solutions to end users while achieving cost efficiencies in the current fast-paced technological environment.

Several prior studies have looked into the aspects of developers who are working on DevOps pipelines of projects. For example, Wiedemann et al. [69] highlight the importance of having a dedicated team responsible for DevOps within an organization. They mention that the expertise required to manage and automate the operations of the IT infrastructure is vastly different from that needed for a software developer, and as such, individuals who possess these skills should be specifically sought after. On the other hand, major tech companies, such as Amazon, push a contradictory ideology –

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MSR '24, April 15–16, 2024, Lisbon, Portugal

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0587-8/24/04...\$15.00

<https://doi.org/10.1145/3643991.3644924>

you build it, you run it [61]. This ideology advocates that all software developers responsible for building a product should also be responsible for everything required to run or operate it (including DevOps). They mention that developers who practice this ideology would also use the same creativity in supporting the applications they built, leading to a better overall experience for the end users.

The above two schools of thought are contradictory. Thus, we perform an empirical study to see if it is effective for everyone in a project to be involved in DevOps or if only a subset of people is authoring the DevOps artefacts in a project. In particular, our paper investigates the impact of code ownership of DevOps artefacts, i.e., the extent to which the DevOps artefacts are authored, in the outcomes of *DevOps-related CI builds*. A DevOps-related CI build is a CI build that is associated with one or more commits that modify, add, or delete at least one DevOps artefact; we refer to such builds as *DevOps CI builds* in the rest of this paper. We conduct our analysis on a large dataset of 892,193 CircleCI¹ builds triggered from the commits related to DevOps artefacts, spanning 1,689 Open-Source Software (OSS) projects. Below, we present the research questions that we answer and a preview of the corresponding results:

(RQ1) Does the code ownership of DevOps artefacts affect outcomes of DevOps CI builds in OSS?

Motivation and approach. To investigate the impact of code ownership on outcomes of DevOps CI builds, we extend the definition of code ownership by Bird et al. [7] to measure ownership over time, and we term this new metric “total chronological ownership,” i.e., the code ownership of an author at the time a build is triggered. We use a mixed-effects logistic regression model [8] to examine the link between this ownership and CI build outcomes, controlling for confounding factors, such as commit count, as discussed in prior research [44, 71].

Results: Our mixed-effects logistic regression model shows that the code ownership of DevOps artefacts is related to the outcome of DevOps CI builds. In fact, there is a statistically significant positive relationship (coefficient = 0.173893) between the total chronological code ownership of DevOps artefacts and successful DevOps CI builds, i.e., as a developer’s DevOps code ownership increases, so does the likelihood of their CI builds succeeding. This positive relationship underscores the pivotal role of a developer’s ownership within the CI process, e.g., project maintainers can prioritize builds triggered by developers with substantial code ownership and allocate additional resources to these builds; on the other hand, for builds triggered by developers with low code ownership, project maintainers may allocate more experienced developers, or those with a thorough understanding of the project, to review these builds; also, pairing more experienced developers with less experienced developers may enhance the review process. Our model also shows that certain controlled features are also positively correlated with CI build outcomes. For example, the number of prior builds a committer triggers positively relates to a successful build outcome.

(RQ2) Does the *Skewness* of DevOps contributions affect the success rate of DevOps CI builds in OSS projects?

Motivation and Approach: Understanding the impact of code ownership on DevOps CI builds on a project level is crucial for

improving the overall software development and the developer community in OSS. A more skewed distribution of DevOps contributions reflects that a majority of developers are making only minor contributions to DevOps (i.e., small code ownership) while a few developers are making significant contributions to DevOps (i.e., large code ownership), and vice versa. We use a linear regression model to study the relationship between the *Skewness* of DevOps contributions and the success rate of DevOps CI builds. We further control this model for features, such as the number of DevOps authors and the number of DevOps commits.

Results: Our results show that the *Skewness* of DevOps contributions in OSS projects has a statistically significant positive impact on projects’ success rate of DevOps CI builds. We find that the linear regression coefficient is 1.28 for modelling the rate of successful DevOps CI builds. Thus, the higher the *Skewness*, the higher the successful rate of DevOps CI builds in projects. Accordingly, confining DevOps-related changes to a specific group of developers enhances the likelihood of successful DevOps CI builds. Therefore, project maintainers may advocate for dedicated DevOps personnel, allowing other developers to focus on their main tasks.

Our study reveals that code ownership of DevOps artefacts strongly impacts the outcome of DevOps CI builds and informs project maintainers to take cognizance of our findings while forming development teams. The three major contributions of our study are (1) an in-depth analysis of build-level code ownership of DevOps artefacts, (2) an in-depth analysis of project-level code ownership of DevOps artefacts, (3) a large dataset of 892,193 DevOps CI builds with the build-level and project-level features and the scripts to reproduce this dataset and our statistical analyses (RQ1 and RQ2).²

2 KEY DEFINITIONS

Below, we define the terminologies that we use in this paper.

DevOps Artefact. A file pertaining to a DevOps tool,³ e.g., the `Dockerfile` is a DevOps artefact used for a containerization tool: Docker.⁴

DevOps Commit. A git commit that modifies/adds or deletes at least one DevOps artefact, e.g., a commit (hash `#da500aa`) that contains changes to four files, including a `Dockerfile`.

DevOps Build. A CI build associated with at least one DevOps commit, e.g., a CI build that contains a few commits, including the above commit (`#da500aa` that modifies a `Dockerfile`).

3 EXPERIMENTAL DESIGN

This section describes our process for collecting and curating the dataset we use to address our research questions. In Figure 1, we provide an overview of our study design, which we detail below.

3.1 Data Preparation

Our study integrates data from various sources (i.e., GitHub API and Gallaba et al. [22]) to acquire the data needed for our analysis. Below, we describe the data preparation steps in detail.

²<https://zenodo.org/records/10578690>

³<https://periodictable.digital.ai/>

⁴<https://www.docker.com/>

¹<https://circleci.com/>

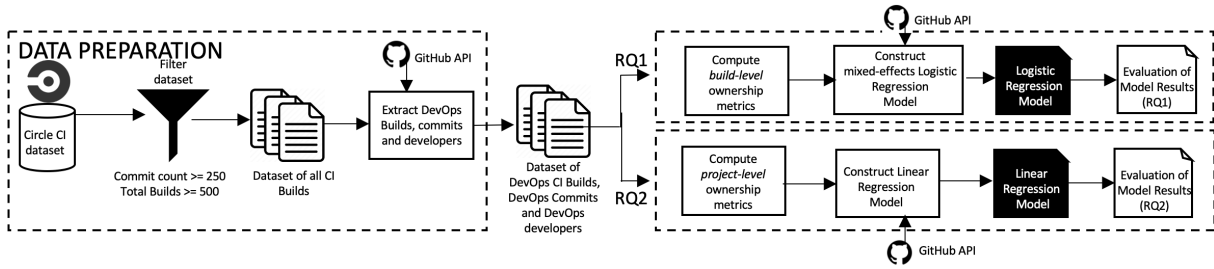


Figure 1: Experimental design overview.




Figure 2: Commit thresholds.




Figure 3: Build thresholds.

3.1.1 Filtering the dataset. We begin with a dataset of 22 million CircleCI builds spanning eight years across 7,795 GitHub repositories, as collected by Gallaba et al. [22]. Note that we choose CircleCI as our CI service due to its widespread adoption, demonstrated by the highest number of installs (748k) on the GitHub marketplace [22]. On this dataset, we first perform project-level filtering to eliminate *toy projects*, inspired by previous work [35, 42], by using thresholds of the number of commits and the number of builds.

Number of commits. Fig. 2 plots candidate threshold values of the number of commits against the number of surviving projects. We select a threshold of 250 commits because it is closer to a “knee” in the curve. This threshold selects 3,247 projects.

Number of builds. In Fig. 3, we illustrate candidate threshold values for the number of builds against the number of surviving projects, with a threshold of 500 builds chosen due to its proximity to a “knee” in the curve, resulting in a dataset of 17,304,451 builds across 1,794 projects.

Table 1: Some regex patterns for detecting DevOps files.

DevOps Category	Regular Expression	DevOps Tool
Continuous Integration	<code>*[.]circleci/config[.]yaml\$</code>	CircleCI
	<code>*[.]github/workflows/*</code>	GitHub Actions
Containers	<code>*[.]dockerfile\$</code>	Docker
	<code>*mesos-master[.]sh\$</code>	Mesos

The dataset provided by Gallaba et al. [22] contains the outcome for every build, which could be one of the following: *success*, *failed*, *timeout*, *infrastructure fail*, etc. We select the builds with either the *success* status or *failed* status.

3.1.2 Retrieving commit data. After filtering the projects and builds of interest, the next step is extracting commit data associated with builds. The original dataset provided by Gallaba et al. [22] contains commit information associated with each build in the dataset. In addition, we need the files associated with the change sets of commits to identify whether a certain commit is a DevOps commit or not, and Gallaba et al.’s dataset does not contain the files changed in the commits. Thus, we use the GitHub API to extract the files in the change set of each commit.

3.1.3 Identifying DevOps-related commits and builds. To filter DevOps commits and builds, we analyze which files are DevOps files in the change sets of commits associated with builds. If there is a DevOps file in a commit, we consider it as a *DevOps commit*, and if there is any DevOps commit associated with a build, we consider that build as a *DevOps CI build*.

To identify DevOps files, we rely on the periodic table of DevOps tools published by Xebia labs⁵ for a comprehensive list of all DevOps tools and the categories of their usage. Since there is a large variety of DevOps tools, we select the tools used for Continuous Integration, Deployment, Containers, and Configuration because such practices are commonly used in projects that use DevOps [58]. These DevOps tools include Kubernetes, Docker, Travis CI, etc.

Next, we review the above tools’ documentation and list the file names or file extensions of their configuration files. Then, we build a DevOps file classifier based on a regular-expression (regex) search to detect DevOps artefacts by examining the file names, file extensions, and directories in which they are stored. Table 1 shows two regex patterns we use and the corresponding tools they

⁵<https://periodictable.digital.ai/>

Table 2: Failed DevOps CI builds.

Build	Build Number	Outcome
DevOps Build 1	t	Failed
DevOps Build 2	t+1	Success

detect. To ensure the validity of this classifier, we manually check a sample of 400 files. We find that the Cohen’s Kappa agreement level [10] between the coder and the classifier is 0.82—a strong agreement. Then, we use this file classifier to analyze all the files in every commit in every build in our filtered project dataset.

We find 444,630 DevOps commits and 1,467,710 DevOps builds. Five projects do not have any DevOps builds; thus, we remove them from the study, reducing the number of projects in our dataset to 1,789. To be certain that build failures in our dataset resulted from changes to DevOps artefacts within the build and not source files, we consider a failed build as a failed “DevOps” build if and only if there is a successful DevOps build observing right after it. Table 2 describes what we define as a failed DevOps build.

We drop all other failed builds with DevOps commits that do not meet this criterion. To mitigate the bias of the previous build’s outcome introduced by this step, we exclude the successful builds immediately after failed DevOps builds. Doing so reduces the number of builds in our dataset to 996,924 across 1,789 projects.

3.1.4 Merging author identities. GitHub allows multiple names and email addresses to be associated with the same GitHub account.^{6,7} As a result, we notice that some commits within a project have very similar names and are most likely the same developer. For example, some commits could be contributed by *Mark Smith* while others would be contributed by *M. Smith*. To account for this, we write a script to merge similar author names. We leverage previous work on string matching author names [13, 19] and computed the jaro-winkler similarity between the author names. If this value was greater than 0.9 [13], we assume the authors are the same and merge their contributions. To ascertain that this script works as it should, we manually unify author names from a random sample of 92 repositories and find a Cohen’s Kappa agreement score [10] of 0.91 with our script, which shows a strong agreement.

3.1.5 Identifying DevOps developers. We extract the DevOps developers who authored the DevOps commits in our dataset. More specifically, we define a DevOps developer as one that makes at least one DevOps commit to the project. We group all DevOps commits in each project by the author names to see all the authors and the corresponding number of DevOps commits they made.

We exclude single-developer repositories in our dataset to avoid bias, as these repositories typically show 100% chronological ownership (see Section 3.2.1) for every build. Doing so also eliminates the inability to compute the *Skewness* of contributions when there is just one author. This filtering reduces the number of repositories from 1,789 to 1,689, and the number of DevOps builds to 978,009.

⁶<https://git-scm.com/book/en/v2/Git-Basics-Git-Aliases>

⁷<https://docs.github.com/en/account-and-profile/setting-up-and-managing-your-personal-account-on-github/managing-email-preferences/setting-your-commit-email-address>

Additionally, we remove the builds that have more than one DevOps contributor, which further reduces the number of DevOps builds to 892,193.

We also exclude commits made by bots from this analysis. To do this, we built a regex-based *bot classifier* that checked if the author of a commit was a bot or not. We evaluate the performance of this bot classifier by manually classifying 400 author names and computing the Cohen’s Kappa agreement with it; the agreement level is 0.88, which shows a strong agreement.

3.2 RQ1 Design of the Build-Level Analysis

In our RQ1, we analyze the build-level impact of the DevOps code ownership on build outcomes. Below, we detail the steps.

3.2.1 Computing build-Level DevOps ownership. Following Bird et al.’s [7] definition of code ownership, we define chronological code ownership of DevOps artefacts as the percentage of the number of commits made by a DevOps developer to the total number of DevOps commits in the project up until when the build is triggered. For instance, the first DevOps build that an author (*author A*) in a project triggers would have a chronological ownership value of 100%, provided all DevOps commits before the time the build was triggered were made by the same author. Similarly, assuming another author (*author B*) makes a DevOps commit that triggers another DevOps build. If there are ten DevOps commits in total at the time of this build, *author B*’s chronological ownership would be 10%. More examples are shown in replication package.²

3.2.2 Constructing the logistic regression model. Based on previous work [34, 44, 72], we collect 16 features to serve as control features to our model. Some of these features include the number of commits in the build, the number of lines of code changed, the time the build was triggered, etc. We first eliminate colinearity from the features by leveraging the `varclus` function in the `Hmisc`⁸ package (in *R*). A total of 15 control features survived the analysis. We build an initial model with these features, and then we analyze the Variance Inflation Factors (VIF) to detect multicollinearity [5]. Following the previous work by Nadri et al. [43], we remove features with VIF values greater than three. Table 3 shows the 13 features that survive. Using these features, we build a mixed-effects logistic regression model to investigate the relationship between the *chronological ownership* of DevOps artefacts and the outcome of DevOps CI builds. Moreover, the use of this “mixed-effects” model accounts for the hierarchical nature of the data by including the name of the repositories as a random effect.

3.3 RQ2 Design of the Project-Level Analysis

Understanding code ownership is not only about who owns what but also unveils insights into how developers contribute to DevOps tasks across the entire project. RQ2 investigates code ownership’s lasting effects during the build stage and its potential cascading impact on DevOps contributions distribution among authors. This analysis aims to uncover how optimizing code ownership can seamlessly enhance collaborative efforts and overall project efficiency.

⁸<https://www.rdocumentation.org/packages/Hmisc/versions/5.1-1>

Table 3: Features used in the mixed-effects logistic regression model.

Category	Feature	Definition	Rationale
Change size (inspired by [38])	number_of_devops_commits	Number of DevOps commits in the build	The number of commits and lines changed in the build may impact the outcome of the build [38].
	devops_change_size	Number of lines changes in DevOps artefacts	
Files Changed [38, 71, 72]	gh_diff_files_removed	Number of files removed from the build	The number and nature of the files changed in the build may impact the outcome of the build [72].
	num_of_devops_artefacts	Number of DevOps artefacts in the build	
Link to Last Build [28, 44]	prev_built_result	Outcome of the last build	The characteristics of the last build may impact the outcome of the current build [28, 44]. For example, if the current build has a recent history of build failures, it is more likely that the build also resulted in a failure.
	same_committer	Boolean value showing if the last build was triggered by the same developer	
Triggering	gh_is_pr	Boolean value showing if the build is part of a Pull Request	The characteristics of the commit that triggered the build may impact the outcome of the build [28, 38, 57].
Commit [28, 38, 57]	day_week	Day of the week the build was triggered	
	time_of_day	Time of the day the build was triggered	
Cooperation [72]	num_of_distinct_authors	Number of distinct authors in the build	The extent of cooperation in the project could impact the outcome of the build [72].
Committer	committer_fail_history	The fail rate of the builds by the current committer in the past	A committer's build history is known to have a significant impact on the outcome of a build [44, 50]. For example, if the committer have a strong history of triggering builds that fail, they may be mostly working on a set of artifacts that are prone to failures.
History [44, 50]	committer_recent_fail_history	The fail rate of the last five builds by the current committer	
	committers_avg_exp	Average number of builds per committer at the time the build was triggered	
Ownership	total_chronological_ownership	Percentage of DevOps commits made by current committer to the total number of DevOps commits at the time of the build	This is our independent variable, and is not a control feature.

Below, we detail our analysis.

3.3.1 Computing project-level DevOps ownership. We compute the *Skewness*⁹ of the distribution of DevOps contributions made by DevOps developers for every project in our dataset. The *Skewness* measures how distorted a particular distribution is from the normal distribution. A positive *Skewness* indicates a right-skewed distribution, negative *Skewness* a left-skewed distribution, and zero *Skewness* a normal distribution.

Fig. 4 shows the helm/helm project's density plot¹⁰ corresponding to the number of contributions made by authors on DevOps commits. The x-axis shows the number of DevOps contributions, and the y-axis shows the *density* of authors. This figure shows an overwhelming majority of developers who made a few DevOps contributions. Thus, the distribution of the author contributions is right-skewed, and the *Skewness* value is 6.49 (a positive value). Such a right-skewed distribution indicates that a majority of the developers in the project made a relatively small number of contributions to DevOps, while a few developers made most of the contributions.




Figure 4: Density plot of DevOps contributions (right-skewed) in helm/helm.

This is particularly relevant to our definition of code ownership of DevOps artefacts because it shows that a few developers in the project have significantly higher ownership of DevOps artefacts than the rest. This also implies that a dedicated group of DevOps developers may maintain DevOps artefacts in the project. Similarly, a left-skewed distribution indicates that the project does not have a potential group of developers working on DevOps artefacts, (i.e.,

⁹<https://www.scribbr.com/statistics/skewness/>

¹⁰A density plot is a representation of the distribution of a numeric variable. It uses a kernel density estimate to show the probability density function of the variable




Figure 5: Boxplot of the Skewness to DevOps contributions.




Figure 6: Boxplot of the rate of successful DevOps builds.

all developers, in general, make both DevOps and non-DevOps contributions. We use this *Skewness* as a proxy to project-level DevOps code ownership, and investigate its impact on the failure rate of DevOps CI builds in the project.

3.3.2 Constructing the linear regression model. We use linear regression to model the impact of project-level code ownership (i.e., *Skewness*) on project-level CI build success rate. To obtain the success rate of each project in our dataset, we express the number of successful DevOps builds over the total number of DevOps builds as a percentage. Since the *Skewness* is not the only feature influencing DevOps build failure rate, we also collect several project-level control features, such as the number of stars, forks, contributors, etc., based on the prior studies [31, 73].

We first remove collinearity and multicollinearity among the extracted features. Table 4 shows the surviving features after removing collinearity and multicollinearity. Next, we pass these features into a linear regression model. We present our results in Section 4.

Challenges of extracting our datasets for the analyses needed for RQ1 and RQ2. Identifying DevOps commits and extracting the features needed to construct the two models are challenging and demanding. Our process involve costly GitHub API calls for every commit in the 15,655,048 CircleCI builds of our initial dataset. To expedite data collection, we use multiprocessing¹² for parallel GitHub API queries, resulting in 7.6 CPU core years of total resource usage. The complete data collection took ten months.

4 RESULTS

4.1 (RQ1) Does the code ownership of DevOps artefacts affect the outcome of DevOps CI builds in OSS?

To analyze the impact of chronological code ownership of DevOps artefact on DevOps build outcomes, we construct a mixed-effects logistic regression model (Section 3.2). The model’s performance is evaluated using conditional R^2 . In particular, the conditional R^2 metric measures the total variance explained by both fixed and random effects. Our model achieves a value of 0.74, indicating that it explains approximately 74% of the variance with both fixed

¹²<https://docs.python.org/3/library/multiprocessing.html>

Table 4: Features used in the linear regression model.

Feature	Definition	Rationale
number_of_repo_builds [59]	Total number of builds in the project	If the number of builds is high, the project is likely to have a mature DevOps pipeline.
number_of_devops_artefacts	Total number of DevOps artefacts in the project	The higher the number of files associated with the build, the more likely the build is to fail [72].
project_maturity [43, 48, 59]	Number of days since the first build was made	The age of a project impacts the outcome of CI builds [59].
recent_failure_rate	Failure rate of last five builds in the project	Having a history of recent may result in more build failures.
commits_per_build	Average number of commits per build in the project	This measures the complexity of builds; more complex builds are more likely to fail.
number_of_devops_authors [3, 72]	Total number of DevOps developers in the project	The extent of cooperation in the project could impact the outcome of the build [72].
forks [3, 43, 48]	Total number of forks on GitHub ¹¹	This measures popularity. A more popular project may have more builds and may be more prone to failures.
repo_skewness	<i>Skewness</i> of developer contributions to DevOps in the project	This is the independent variable.

and random effects considered. Additionally, we report Akaike Information Criterion (AIC) [4] and Bayesian Information Criterion (BIC) [54] scores, which are 303440.5 and 303779.1, respectively. Such high values of AIC and BIC reflect the strong explainability of the dependent variable by the independent variables in our model.

In Table 5, we present the impact of different features (including the chronological code ownership) on the outcomes of DevOps CI builds. The *coefficients’* magnitudes signify their strength, while the *coefficients’* signs indicate the direction of the relationship of the corresponding feature and the DevOps CI build outcome. A negative coefficient shows that as the feature increases, the dependent variable (the DevOps build outcome) becomes more likely to be a failure and vice versa. Conversely, a larger magnitude shows that a feature has a stronger impact on the DevOps CI build outcome and vice versa. Note that we also show the statistical significance levels of all the coefficients in the *p-value* column.

In addition, following the previous work [43, 48], we compute the odds ratio of each feature to better understand the relationship between the dependent and independent variables. We show the odds ratios in the *Odds* column of Table 5. The interpretation of the odds

Table 5: Results of the build-level mixed-effects logistic regression model.

Feature	Coefficient	P-Value	Odds
number_of_devops_commits	0.034758	3.32e-10 ***	1.04
devops_change_size	-0.050963	1.22e-12 ***	0.95
prev_built_resultcanceled	-0.383202	3.54e-09 ***	0.68
prev_built_resultfailed	-2.031670	< 2e-16 ***	0.13
prev_built_resultinfrastructure_fail	0.338537	0.096438 .	1.40
prev_built_resultno_tests	0.281516	0.054136 .	1.33
prev_built_resultsuccess	0.485515	< 2e-16 ***	1.63
prev_built_resulttimeout	-0.035349	0.834786	0.97
num_of_distinct_authors	0.002759	0.612027	1.00
num_of_devops_artefacts	-0.039625	1.94e-08 ***	0.96
gh_diff_files_removed	0.019927	0.000883 ***	1.02
gh_is_prTrue	-0.413991	< 2e-16 ***	0.66
day_weekMonday	-0.050158	0.003287 **	0.95
day_weekSaturday	-0.096326	2.75e-05 ***	0.91
day_weekSunday	-0.022346	0.344655	0.98
day_weekThursday	-0.026590	0.110966	0.97
day_weekTuesday	-0.014659	0.377967	0.98
day_weekWednesday	-0.060893	0.000232 ***	0.94
time_of_day	0.001523	0.758338	1.00
same_committerTrue	-0.461960	< 2e-16 ***	0.63
committer_fail_history	-0.035224	3.92e-06 ***	0.96
committer_recent_fail_history	-1.220983	< 2e-16 ***	0.29
committer_avg_exp	1.837764	< 2e-16 ***	6.28
total_chronological_ownership	0.173893	< 2e-16 ***	1.19

Significance codes: 0: ***, 0.001: **, 0.05: *

ratio depends on the type of variable. In particular, for continuous variables, the odds ratio shows the chances of the dependent variable occurring per unit change in the independent variable, while for categorical variables, the odds ratio shows the chances of the dependent variable occurring compared to the default value of the categorical variable [43]. Below, we describe the main observations from the table and the corresponding discussions.

Observation 1: High chronological code ownership of DevOps artifacts in DevOps CI builds is associated with more successful DevOps CI builds and vice versa.

Chronological code ownership (build-level code ownership) of a DevOps CI build is the percentage of the number of commits made by a DevOps developer to the total number of DevOps commits in the project up until when the build is triggered (Section 3.2.1). Table 5 shows that the chronological ownership of DevOps artefacts has a statistically significant impact ($p < 0.001$) on the outcome of DevOps CI builds. Its positive coefficient indicates that a high code ownership is associated with an increased probability of achieving a successful build outcome. From the odds ratio of the total chronological ownership ($\exp(0.173893) = 1.19$), we can deduce that the chances of a successful build would increase by 19% for every

percentage increase in an author's DevOps commits relative to the total number of DevOps commits in the project.

Furthermore, we use the ANOVA test [53] to determine how much of the variance in build outcomes can be explained by total chronological ownership. We find that code ownership of DevOps artefacts explains 0.69% of the variance of DevOps CI builds' build outcomes. While this percentage of explained variance may be modest, the model coefficients establish a robust and statistically significant correlation between high code ownership in DevOps artefacts and increased success in DevOps CI builds, underscoring the impact of code ownership on the overall outcomes of DevOps CI builds. This observation about code ownership complements Wiedemann et al.'s [69] work, which emphasizes the positive correlation between a concentrated group of developers overseeing DevOps artefacts and enhanced success in DevOps CI builds.

Despite observing a positive relationship between code ownership and DevOps CI build outcomes, a close inspection of our dataset reveals several instances where DevOps builds deviate from this trend. To delve deeper into such deviations, we further analyze builds that failed with high code ownership and those that were successful with low code ownership. We consider builds with high code ownership to be those with code ownership within the upper quartile (top 25%) and builds with low ownership to be those with ownership within the lower quartile (bottom 25%).

Discussion 1: DevOps builds that fail despite having high levels of chronological code ownership are usually early builds and vice versa.

A project's first set of DevOps CI builds typically has high ownership because only a few authors have made contributions at the time. For example, suppose there are ten DevOps commits in a project; six of these commits are made by *author A*, while the remaining four are made by *author B*. If a build is triggered by *author A*, then this build would have chronological ownership of 60%; in the opposite case, where that build is submitted by *author B*, the chronological ownership value would be 40%, which are both substantial levels of chronological ownership.

Our analysis shows 24,971 failing builds with such high ownership levels. These builds are concentrated at the early stages of repository timelines, often associated with the initial configuration of CI pipelines, making them more prone to errors. For builds failing with high code ownership, we find that the mean and median values for the number_of_builds_before_a_build are 511.6 and 113, respectively. In contrast, when considering the entire dataset, these values are notably higher at 5,847 and 1,467, underscoring the unique characteristics of builds with high code ownership. An illustrative case is the #60032dd1e1d87f460fac092f build of the pankona/gomo-simra project, with chronological code ownership of 93.33% at the time but fails; the number of builds before that particular build is 37, which is much smaller than the mean and the median of the entire dataset.

Likewise, the above rationale extends to successful builds despite low ownership levels. We find 204,951 such builds in our dataset, which often appear *later* in the repository timelines, depicting cases where the pipeline must have been set thoroughly

and extensively tested. We find that, for builds in this category, the `number_of_builds_before_a_build` is much higher than that of the entire dataset; the mean and median are 13,853 and 6,467, respectively, clearly surpassing the corresponding values for the whole dataset, which are 5,847 and 1,467.

Observation 2: *As the number of DevOps artefacts in a build increases, the likelihood of a successful build decreases.*

Another interesting observation we made from our model (Table 5) is that the number of DevOps artefacts in a build also has a statistically significant effect ($p < 0.01$) on the outcome of the build. Its coefficient indicates a negative relationship between the number of DevOps artefacts and the dependent variable (the build outcome), i.e., as the number of DevOps artefacts increases, a build's success becomes less likely. In addition, the odds ratio of this feature ($\exp(-0.039625) = 0.96$) shows that for every extra DevOps artefact in the build, there is a 4% decrease in the chances that the build succeeds. A similar trend is also observed in the DevOps change size feature, which measures the number of lines changed in all DevOps artefacts in the build. Table 5 shows that there is a 5% ($\exp(-0.050963) = 0.95$) decrease in the chances of success of a build for every line changed in a DevOps artefact. This observation is in line with the prior work of Luo et al. [38], who modeled CI build outcomes (regardless of being related to DevOps files). They found that as more files are changed in a build, there is less chance for that build to result in a successful one.

Observation 3: *As the committers' average experience increases, the likelihood of the DevOps CI build's outcome being a success also increases.*

Table 5 shows that the average number of builds of committers at the time of the current build (committers' average experience) has a statistically significant impact ($p < 0.001$) on the outcome of DevOps CI builds. Its positive coefficient shows that provided all other features remain constant, as the committers' average experience increases, the chances of that build's success also increase.

Our results further show that the committers' average experience explains about 44% of the variance of the DevOps CI build outcome by conducting the ANOVA test. We believe this is because the committers' average experience is pivotal in capturing the typical instability observed during the initial stages of build creation, where inexperienced contributors may encounter challenges. It is not uncommon for early builds to exhibit failures as they are being set up. Notably, committers' average experience, a historical statistic as reported by Ni et al. [44], is crucial for predicting build outcomes.

Summary RQ1: *On a build level, code ownership of DevOps artefacts does matter. The results of our experiments distinctly reveal a positive relationship between the code ownership of DevOps artefacts and outcomes of DevOps CI builds.*

4.2 (RQ2) Does the Skewness of DevOps contributions affect the success rate of DevOps CI builds in OSS?

We construct a linear regression model (Section 3.3) to investigate the project-level impact of code ownership (i.e., the *Skewness* of DevOps contributions) on the DevOps CI build success rate of the 1,689 projects in our dataset. Table 6 shows the model results. As similar to Table 5 in RQ1, the magnitude and sign of the *Coefficient* column in Table 6 represents the strength and direction of the relationship between the feature and the dependent variable (build success rate), respectively. From the table, we make the below observations.

Observation 4: *A high Skewness of DevOps contributions within a project is related to elevated success rates in DevOps builds in that project.*

From Table 6, we observe that the *Skewness* of DevOps contributions has a statistically significant effect ($p\text{-value} < 0.05$) on the success rate of DevOps CI builds. We also use the ANOVA test [53] to determine how much of the variance in the success rate of DevOps builds is explained by the *Skewness* of DevOps contributions, and we find this value to be about 2.8%. The model coefficient of the *Skewness* is positive, indicating that as the *Skewness* of DevOps contributions in a project increases, the rate of successful DevOps builds increases. For example, the `helm/helm` project shown in Figure 4 (Section 3) has a *Skewness* of 6.49; the success rate of DevOps CI builds in this project is 91.38%. Furthermore, from figures 5 and 6 (in Section 3), we see that the median *Skewness* and success rate of DevOps builds are 3.05 and 87.62 respectively. Thus, `helm/helm` has a high *Skewness* and a high success rate of DevOps CI builds.

Although highly skewed DevOps contributions elevate the success rate of DevOps builds in projects, we still find some projects in our dataset that deviate from this trend. We perform an extended analysis, investigating the characteristics of the deviated projects to understand the underlying reasons for such deviations better. In particular, we investigate projects with high *Skewness* values and low build success rates, as well as projects with low *Skewness* values and high success rates. We consider projects with high *Skewness* values to be the ones with *Skewness* values in the upper quartile, i.e., the top 25%; on the other hand, we consider the projects with low *Skewness* values to be the ones with *Skewness* values in the

Table 6: Results of the project-level linear regression model.

Feature	Coefficient	P-Value
<code>number_of_repo_builds</code>	5.148e-05	0.000244 ***
<code>number_of_devops_files</code>	-2.745e-04	0.931195
<code>project_maturity</code>	-7.436e-04	0.177558
<code>recent_failure_rate</code>	-1.344e-01	<2e-16 ***
<code>commits_per_build</code>	-2.092e-02	0.521927
<code>number_of_devops_authors</code>	-5.267e-04	0.904448
<code>forks</code>	-6.500e-05	0.711234
<code>repo_skewness</code>	2.286e-01	0.019384 *

Significance codes: 0: ***, 0.001: **, 0.05: *

lower quartile (bottom 25%). Similarly, we obtain projects with high *Skewness* and low build success rates. This analysis reveals that 78 projects have high *Skewness* values and low build success rates.

Discussion 2: *Projects with high Skewness of DevOps contributions and low DevOps CI build success rates have more complex builds.*

By delving deeper into the attributes of these projects, we find that they have more intricate builds, i.e., builds with a high average number of lines changed per build. The mean and median of the average number of lines changed per build in a project with high *Skewness* values with low build success rates are 36,015 and 1,944, respectively, and are much higher than those for the projects in our entire dataset (mean and median of the whole dataset are 12,167 and median 1,247, respectively).

On the other hand, we find 97 projects with low *Skewness* and high build success rates, all having fewer DevOps developers.

Discussion 3: *Projects with low Skewness values and high build success rates have a small number of developers.*

For example, `hanami/validations` project has a *Skewness* of 0.707, a build success rate of 95.41%, and has just three DevOps developers. In general, among the projects with low *Skewness* values and high build success rates, we find that the mean and median number of developers are four and six, respectively, which are much less than the mean (61) and median (20) values of the entire dataset. The fact that these projects have high build success rates despite having a small number of developers implies that when a project has a small number of developers, a more balanced allocation of contributions may improve the DevOps build success rate.

Because the number of developers in a project impacts the DevOps build success rate, we study the distribution of the number of developers in the projects that fully satisfy our Observation 4, i.e., the projects with high *Skewness* and build success rates.

Discussion 4: *Although the total number of developers that contribute to projects with high Skewness values and high build success rates is large, the number of DevOps developers that make most of the DevOps commits is relatively small.*

This analysis shows that projects with high *Skewness* and high success rates have more DevOps developers relative to the entire dataset and projects with low *Skewness* and high success rates, with mean and median developer counts being 222 and 113, respectively.

To examine how many developers are accountable for the majority of the DevOps commits in the projects with high *Skewness* values and high build success rates, we retrieve the number of developers in each project who are responsible for 80% of the DevOps commits in the project. We find that the mean and median numbers of DevOps developers contributing to 80% of the DevOps commits made in these projects are seven and three, respectively. This reinforces the idea that, despite a sizable developer pool, a select few drive

most changes in DevOps artefacts, leading to greater ownership and success in DevOps CI builds (Observation 1 in RQ1).

Summary RQ2: *On a project level, the Skewness of DevOps contributions does affect the success rate of DevOps CI builds. Our results show a statistically significant positive correlation between the two.*

5 THREATS TO VALIDITY

Construct Validity—We use a regular expression-based classifier to decide whether a code artefact is a DevOps artefact. This also cascades into our definitions for DevOps commits and DevOps builds. We rely solely on the project's file names, extensions, and directories to make our classification. Relying on a regular expression-based classifier could lead to misclassifying artefacts, when distinguishing between DevOps and other file types. To mitigate this, we randomly select a sample of 400 files from our dataset and manually classify these files. We then compute the Cohen's Kappa agreement score [10] between the file classifier and the coder. We find an agreement level of 0.82, which indicates a near-perfect agreement.

Due to GitHub allowing aliases, authors can commit under different names,^{6,7} complicating the quantification of their DevOps contributions. To mitigate this threat, we leverage previous work [13, 19] and compute the `jaro-winkler` similarity between author names. If the calculated similarity exceeds 0.9 [13], we assume they are the same author and sum up the contributions. We validate this alias-merging method by manually unifying author names in a random sample of projects and comparing it with the Jaro-Winkler similarity-based alias-merging method. We find that the Cohen's Kappa agreement to be near-perfect (0.91).

Internal Validity—We may have missed confounding factors that could impact the interpretations of our results, e.g., we observe that the previous build outcome has a significant impact on the current build outcome (Table 5 in Section 4.2). A possible root cause for repeated failed builds might be developers retrying builds due to perceived flakiness [6]. We call them *failure streaks*. While distinguishing whether our features have causal or merely correlational links with DevOps build outcomes is a future research avenue, our models demonstrate that several features share useful correlational relationships with DevOps build outcomes, which can explain the studied cases of DevOps build failures effectively.

We define DevOps commits as commits that add and/or modify DevOps artefacts, and we consider the builds triggered by DevOps commits as DevOps builds. This definition of DevOps builds may be biased. This is because a build may be associated with DevOps commits and other commits, and thus, the build's outcome could result from a commit that is not a DevOps commit. Consequently, the failure of a build that we consider a DevOps build may not be failed because of a DevOps commit but because of another commit. To ensure that we only consider DevOps builds that failed only because of a DevOps commit (and not because of other commits), we consider failed DevOps builds as the DevOps builds preceded by a successful DevOps build, as depicted in Table 2 in Section 3.1.3. Doing so also mitigates any biases in our results, for example, due to *failure streaks*.

External Validity—Our findings based on CircleCI may not fully apply to other CI services but could be adaptable due to common features like YAML and pipeline structures. Future research should explore different CI services for more general insights.

6 RELATED WORK

DevOps Technologies — In contemporary software engineering, the emergence of DevOps represents a relatively recent but pivotal concept. Several studies [11, 16, 32, 58] defined the term DevOps. For example, Jabbari et al. [32] defined DevOps as “a development methodology aimed at bridging the gap between Development and Operations.” Among several benefits of adopting DevOps, the major one is that it forces the development and operations teams to interact with each other more than before, leading to enhanced collaboration and communication [49].

However, adopting DevOps is not without its limitations; several studies [12, 26, 55, 58, 67], thus, focused on the problems in adopting DevOps practices in organizations. Grande et al. [26] performed a systematic review on DevOps challenges in globally distributed teams, identifying key issues: the need for complex skillsets, communication barriers between developers and operations, and employee resistance to change. Diel et al. [12] specifically examined communication challenges, noting that geographic distance and interaction frequency among team members can impact the communication gap. Other studies [29, 62] discussed the challenges of configuring complex DevOps files. In particular, Tamraw et al. [62] and Henkel et al. [29] discussed the challenges of configuring build files (e.g., `Makefiles`) and writing `Dockerfiles`, respectively. These studies further suggested tools to help developers work on such files.

Other work studied the tasks of DevOps developers [36, 70]. For example, Kerzazi et al. [36] conducted an empirical analysis of online job postings, aiming to identify and compare the primary responsibilities of DevOps engineers. This analysis revealed that *automation* is the essential activity expected by DevOps developers, which is further explored in several other studies [41, 68].

While many of the above studies explore various facets of DevOps, our work takes a distinct approach by examining DevOps from a new perspective; we shed light on the impact of code ownership of DevOps artefacts and the *Skewness* of DevOps contributions in a project, areas that have not been extensively studied.

Continuous Integration (CI) — CI is a DevOps practice that frequently integrates code changes into a shared code base [15]. CI provides several advantages to the software teams that adopt it [30, 64]. For example, Vasilescu et al. [64] studied a dataset of 246 OSS projects that use CI. They found that CI improves the productivity of project teams. While CI is advantageous for projects, other studies revealed the difficulties of adopting CI. Among the several challenges that come with CI, build failures are a challenge that is unavoidable for many projects [23, 37, 40, 45, 46, 60, 65, 66]. For example, Kerzazi et al. [37] analyzed 3,214 builds produced in a large software company over a period of six months, and found that a substantial proportion (17.9%) of builds are failing. In addition, Vassallo et al. [65] analyzed build failures in 349 OSS projects and 418 proprietary projects. This study revealed, for both the OSS and proprietary projects, the overall percentage of failing builds during the period of observation is 26%, which is much higher compared

to Kerzazi et al.’s [37] study. Vassallo et al. [65] further showed that failures related to deployments are a key category of build failures.

Due to challenges of observing build failures, several studies [9, 28, 38, 44, 50, 51, 71, 72] focused on predicting build outcomes in advance. For example, Luo et al. [38] constructed four build outcome prediction models for the TravisTorrent dataset. They found that the *number of commits in a build* is the most critical factor determining the probability of build failures. Hassan et al. [28] conducted a similar study on the same dataset but restricted their investigation to the projects using Ant, Gradle, and Maven. They utilized a random forest model, and found the outcome of the previous build to be the most crucial factor in predicting build outcomes. Furthermore, Saidani et al. [50] used deep learning to predict CI build outcomes. They also leveraged the TravisTorrent dataset but restricted the study to the top ten projects with the highest number of builds.

In addition to the issue of build failures, there are several other challenges. The need to restart builds has been highlighted as a substantial drain on CI build time [14, 39, 56], and longer build durations [20, 22, 24, 25] are known to not only overuse CI resources but also increase the time to feedback. To that end, previous work also focused on reducing the time-to-feedback of CI builds [1, 2, 21, 52]. For example, Abdalkareem et al. [1, 2] proposed a method to detect when to skip CI builds based on commits that do not affect the source code (e.g., changes to `Readme` files). This method could reduce the number of commits triggering CI by 18%.

The above studies focused on CI builds in general. Our study, however, thoroughly analyzes builds associated with commits on DevOps artefacts (DevOps CI builds), and discusses the impact of code ownership of these artefacts on the outcomes of such builds.

Code Ownership — Several studies have studied the relationship between code ownership and software quality [7, 17, 18, 27, 47, 63]. Bird et al. [7] conducted an empirical study on two large proprietary software projects, and investigated the relationship between different code ownership metrics and the prevalence of software failures. They find that in all cases, metrics like the number of low-expertise developers and the proportion of ownership of the top owner have a relationship with the occurrence of pre-release faults and post-release failures. Foucault et al. [18] replicated Bird et al.’s [7] study in the context of OSS projects. They studied seven OSS projects and found that the results in the OSS context were inconsistent with those of Bird et al., i.e., code ownership does not impact software quality in OSS projects. Thongtanunam et al. [63] extended the work done by Bird et al. [7] and Foucault et al. [7] by using new code ownership metrics that consider code reviews. They performed empirical analyses on six releases of two large OSS projects, and found that code ownership metrics that consider code review are related to software quality.

Several prior studies have delved into the dynamics of having dedicated DevOps developers versus employing full-stack developers. For example, Wiedemann et al. [69] underscore the significance of establishing a dedicated DevOps team within an organization. The study emphasizes that the expertise needed to oversee and automate IT infrastructure operations differs substantially from that of a software developer. In contrast, major tech companies,

such as Amazon, advocate for a model where all software developers responsible for product development also run and operate the product. The article “Why You Should Run What You Build” by Stephen [61] highlights that developers embracing this approach apply the same creativity used in building applications to DevOps, ultimately enhancing the overall user experience.

While code ownership has been studied in the context of open source and code quality, the existing studies do not investigate code ownership in the context of CI. Our study extends the existing studies by focusing solely on the influence of code ownership (i.e., chronological code ownership and *Skewness* of DevOps contributions) on the DevOps CI builds.

7 CONCLUSION

By studying the impact of code ownership of DevOps artefacts on the outcome and success rate of 892,193 DevOps CI builds across 1,689 OSS GitHub projects, we make the following two conclusions:

Large software organizations are better off having dedicated DevOps developers handle all DevOps-related commits. In Section 4, both RQs point to the same conclusion: higher build-level (Observation 1) and project-level (Observation 4) code ownership are associated with more successful outcomes for DevOps CI builds. Thus, we advise software practitioners to take advantage of this relationship by ensuring dedicated developers handle DevOps-related tasks in the project. Doing so allows developers with high ownership to focus on DevOps files while gaining project-specific DevOps knowledge that can benefit the project in the long run.

Small software organizations would benefit from all their developers working on DevOps code. Our findings also show that while having a dedicated team is paramount for large organizations, relatively smaller organizations (e.g., projects with six developers, as we observe in Discussion 3 in Section 4.2) could also benefit by letting all the developers work on DevOps artefacts. This collaborative approach may foster a more integrated and agile development process in the early stages of large projects as well.

While our results show that DevOps code ownership impacts DevOps CI builds, we also propose a broader hypothesis—code ownership, in a general sense, could impact CI builds. For example, developers are often rotated among different project components, which may change their ownership of other components and ultimately affect the CI build outcomes. To the best of our knowledge, no previous studies on build prediction models have accounted for code ownership in CI build outcome prediction. We encourage future researchers to consider the chronological code ownership and the *Skewness* of contributions (to specific files and/or code components) when curating features for their prediction models.

DATA AVAILABILITY

Our replication package includes all the scripts associated with our analysis to foster future studies: <https://zenodo.org/records/10570324>.

REFERENCES

- [1] Rabe Abdalkareem, Suhaib Mujahid, and Emad Shihab. 2021. A Machine Learning Approach to Improve the Detection of CI Skip Commits. *IEEE Transactions on*

- Software Engineering* 47, 12 (12 2021), 2740–2754. <https://doi.org/10.1109/TSE.2020.2967380>
- [2] Rabe Abdalkareem, Suhaib Mujahid, Emad Shihab, and Juergen Rilling. 2021. Which Commits Can Be CI Skipped? *IEEE Transactions on Software Engineering* 47, 3 (3 2021), 448–463. <https://doi.org/10.1109/TSE.2019.2897300>
- [3] Roozbeh Aghili, Heng Li, and Foutse Khomh. 2023. Studying the characteristics of AIOps projects on GitHub. *Empirical Software Engineering* 28, 6 (11 2023), 143. <https://doi.org/10.1007/s10664-023-10382-z>
- [4] H. Akaike. 1974. A new look at the statistical model identification. *IEEE Trans. Automat. Control* 19, 6 (12 1974), 716–723. <https://doi.org/10.1109/TAC.1974.1100705>
- [5] Olusegun Akinwande, H G Dikko, and Samson Agboola. 2015. Variance Inflation Factor: As a Condition for the Inclusion of Suppressor Variable(s) in Regression Analysis. *Open Journal of Statistics* 05 (1 2015), 754–767. <https://doi.org/10.4236/ojs.2015.57075>
- [6] Eman Abdullah AlOmar, Moataz Chouchen, Mohamed Wiem Mkaouer, and Ali Ouni. 2022. Code Review Practices for Refactoring Changes: An Empirical Study on OpenStack. (3 2022).
- [7] Christian Bird, Nachiappan Nagappan, Brendan Murphy, Harald Gall, and Premkumar Devanbu. 2011. Don’t touch my code!. In *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*. ACM, New York, NY, USA, 4–14. <https://doi.org/10.1145/2025113.2025119>
- [8] Benjamin M Bolker, Mollie E Brooks, Connie J Clark, Shane W Geange, John R Poulsen, M Henry H Stevens, and Jada-Simone S White. 2009. Generalized linear mixed models: a practical guide for ecology and evolution. *Trends in Ecology & Evolution* 24, 3 (2009), 127–135. <https://doi.org/10.1016/j.tree.2008.10.008>
- [9] Bihuan Chen, Linlin Chen, Chen Zhang, and Xin Peng. 2020. Buildfast: History-aware build outcome prediction for fast feedback and reduced cost in continuous integration. In *Proceedings of the 35th International Conference on Automated Software Engineering*.
- [10] Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin* 70, 4 (1968), 213–220. <https://doi.org/10.1037/h0026256>
- [11] Breno B. Nicolau de França, Helvio Jeronimo, and Guilherme Horta Travassos. 2016. Characterizing DevOps by Hearing Multiple Voices. In *Proceedings of the XXX Brazilian Symposium on Software Engineering*. ACM, New York, NY, USA, 53–62. <https://doi.org/10.1145/2973839.2973845>
- [12] Elisa Diel, Sabrina Marczak, and Daniela S. Cruzes. 2016. Communication Challenges and Strategies in Distributed DevOps. In *2016 IEEE 11th International Conference on Global Software Engineering (ICGSE)*. IEEE, 24–28. <https://doi.org/10.1109/ICGSE.2016.28>
- [13] Paul Donner. 2016. Enhanced self-citation detection by fuzzy author name matching and complementary error estimates. *Journal of the American Society for Information Science and Technology* 67, 3 (2016), 662.
- [14] Thomas Durieux, Claire Le Goues, Michael Hilton, and Rui Abreu. 2020. Empirical study of restarted and flaky builds on Travis CI. In *Proceedings of the 17th International Conference on Mining Software Repositories*.
- [15] Paul M Duvall, Steve Matyas, and Andrew Glover. 2007. *Continuous integration: improving software quality and reducing risk*.
- [16] Floris Erich. 2019. DevOps is Simply Interaction Between Development and Operations. 89–99. https://doi.org/10.1007/978-3-030-06019-0_7
- [17] Csaba Faragó, Péter Hegedűs, and Rudolf Ferenc. 2015. Code ownership: Impact on maintainability. In *Computational Science and Its Applications—ICCSA 2015: 15th International Conference, Banff, AB, Canada, June 22–25, 2015, Proceedings, Part V* 15. Springer, 3–19.
- [18] Matthieu Foucault, Jean-Rémy Falleri, and Xavier Blanc. 2014. Code ownership in open-source software. In *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*. ACM, New York, NY, USA, 1–9. <https://doi.org/10.1145/2601248.2601283>
- [19] Tanner Fry, Tapajit Dey, Andrey Karanach, and Audris Mockus. 2020. A Dataset and an Approach for Identity Resolution of 38 Million Author IDs extracted from 2B Git Commits. In *Proceedings of the 17th International Conference on Mining Software Repositories*. ACM, New York, NY, USA, 518–522. <https://doi.org/10.1145/3379597.3387500>
- [20] Keheliya Gallaba, John Ewart, Yves Junqueira, and Shane McIntosh. 2020. Accelerating continuous integration by caching environments and inferring dependencies. *Transactions on Software Engineering* (2020).
- [21] Keheliya Gallaba, John Ewart, Yves Junqueira, and Shane McIntosh. 2022. Accelerating Continuous Integration by Caching Environments and Inferring Dependencies. *IEEE Transactions on Software Engineering* 48, 6 (6 2022), 2040–2052. <https://doi.org/10.1109/TSE.2020.3048335>
- [22] Keheliya Gallaba, Maxime Lamothe, and Shane McIntosh. 2022. Lessons from eight years of operational data from a continuous integration service. In *Proceedings of the 44th International Conference on Software Engineering*. ACM, New York, NY, USA, 1330–1342. <https://doi.org/10.1145/3510003.3510211>
- [23] Keheliya Gallaba, Christian Macho, Martin Pinzger, and Shane McIntosh. 2018. Noise and heterogeneity in historical build data: an empirical study of travis

- ci. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 87–97.
- [24] Taher Ahmed Ghaleb, Daniel Alencar Da Costa, and Ying Zou. 2019. An empirical study of the long duration of continuous integration builds. *Empirical Software Engineering* 24 (2019).
- [25] Taher A Ghaleb, Safwat Hassan, and Ying Zou. 2022. Studying the Interplay between the Durations and Breakages of Continuous Integration Builds. *Transactions on Software Engineering* 49 (2022).
- [26] Rubén Grande, Aurora Vizcaino, and Félix O. García. 2024. Is it worth adopting DevOps practices in Global Software Engineering? Possible challenges and benefits. *Computer Standards & Interfaces* 87 (1 2024), 103767. <https://doi.org/10.1016/j.csi.2023.103767>
- [27] Michaela Greiler, Kim Herzig, and Jacek Czerwonka. 2015. Code Ownership and Software Quality: A Replication Study. In *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*. IEEE, 2–12. <https://doi.org/10.1109/MSR.2015.8>
- [28] Foyzul Hassan and Xiaoyin Wang. 2017. Change-Aware Build Prediction Model for Stall Avoidance in Continuous Integration. In *2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. IEEE, 157–162. <https://doi.org/10.1109/ESEM.2017.23>
- [29] Jordan Henkel, Christian Bird, Shuwendu K Lahiri, and Thomas Reps. 2020. Learning from, Understanding, and Supporting DevOps Artifacts for Docker. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 38–49. <https://doi.org/10.1145/3377811.3380406>
- [30] 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 international conference on automated software engineering*.
- [31] Dongyang Hu, Tao Wang, Junsheng Chang, Gang Yin, and Yang Zhang. 2018. Multi-Discussing across Issues in GitHub: A Preliminary Study. In *2018 25th Asia-Pacific Software Engineering Conference (APSEC)*. 406–415. <https://doi.org/10.1109/APSEC.2018.00055>
- [32] Ramtin Jabbari, Nauman bin Ali, Kai Petersen, and Binish Tanveer. 2016. What is DevOps?. In *Proceedings of the Scientific Workshop Proceedings of XP2016*. ACM, New York, NY, USA, 1–11. <https://doi.org/10.1145/2962695.2962707>
- [33] Ramtin Jabbari, Nauman bin Ali, Kai Petersen, and Binish Tanveer. 2016. What is DevOps? A Systematic Mapping Study on Definitions and Practices. In *Proceedings of the Scientific Workshop Proceedings of XP2016 (XP '16 Workshops)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/2962695.2962707>
- [34] Xianhao Jin and Francisco Servant. 2021. What Helped, and what did not? An Evaluation of the Strategies to Improve Continuous Integration. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 213–225. <https://doi.org/10.1109/ICSE43902.2021.00031>
- [35] 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*. ACM, New York, NY, USA, 92–101. <https://doi.org/10.1145/2597073.2597074>
- [36] Noureddine Kerzazi and Bram Adams. 2016. Who needs release and devops engineers, and why?. In *Proceedings of the international workshop on continuous software evolution and delivery*, 77–83.
- [37] Noureddine Kerzazi, Foutse Khomh, and Bram Adams. 2014. Why do automated builds break? an empirical study. In *2014 IEEE International Conference on Software Maintenance and Evolution*. IEEE, 41–50.
- [38] Yang Luo, Yangyang Zhao, Wanwangying Ma, and Lin Chen. 2017. What are the Factors Impacting Build Breakage?. In *2017 14th Web Information Systems and Applications Conference (WISA)*. IEEE, 139–142. <https://doi.org/10.1109/WISA.2017.17>
- [39] Rungraj Maipradit, Dong Wang, Patanamon Thongtanunam, Raula Gaikovina Kula, Yasutaka Kamei, and Shane McIntosh. 2023. Repeated Builds During Code Review: An Empirical Study of the OpenStack Community. In *Proc. of the International Conference on Automated Software Engineering*.
- [40] Ade Miller. 2008. A hundred days of continuous integration. In *Agile 2008 conference*. IEEE, 289–293.
- [41] Sikender Mohsienuddin Mohammad. 2017. DevOps automation and Agile methodology. *International Journal of Creative Research Thoughts (IJCRT)*, ISSN (2017), 2320–2882.
- [42] Nuthan Munaiah, Steven Kroh, Craig Cabrey, and Meiyappan Nagappan. 2017. Curating GitHub for engineered software projects. *Empirical Software Engineering* 22, 6 (12 2017), 3219–3253. <https://doi.org/10.1007/s10664-017-9512-6>
- [43] Reza Nadri, Gema Rodriguez-Perez, and Meiyappan Nagappan. 2022. On the Relationship Between the Developer's Perceptible Race and Ethnicity and the Evaluation of Contributions in OSS. *IEEE Transactions on Software Engineering* 48, 8 (8 2022), 2955–2968. <https://doi.org/10.1109/TSE.2021.3073773>
- [44] Ansong Ni and Ming Li. 2017. Cost-Effective Build Outcome Prediction Using Cascaded Classifiers. In *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*. IEEE, 455–458. <https://doi.org/10.1109/MSR.2017.26>
- [45] Zi Peng, Tse-Hsun Chen, and Jinqiu Yang. 2020. Revisiting test impact analysis in continuous testing from the perspective of code dependencies. *IEEE Transactions on Software Engineering* 48, 6 (2020), 1979–1993.
- [46] Jackson A Prado Lima, Willian DF Mendonça, Silvia R Vergilio, and Wesley KG Assunção. 2022. Cost-effective learning-based strategies for test case prioritization in continuous integration of highly-configurable software. *Empirical Software Engineering* 27, 6 (2022), 133.
- [47] Foyzur Rahman and Premkumar Devanbu. 2011. Ownership, experience and defects: a fine-grained study of authorship. In *Proceedings of the 33rd International Conference on Software Engineering*. 491–500.
- [48] Ayushi Rastogi, Nachiappan Nagappan, Georgios Gousios, and André van der Hoek. 2018. Relationship between geographical location and evaluation of developer contributions in github. In *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/3239235.3240504>
- [49] Leah Riungu-Kalliosaari, Simo Mäkinen, Lucy Ellen Lwakatara, Juha Tiihonen, and Tomi Männistö. 2016. DevOps adoption benefits and challenges in practice: A case study. In *Product-Focused Software Process Improvement: 17th International Conference, PROFES 2016, Trondheim, Norway, November 22-24, 2016, Proceedings* 17. Springer, 590–597.
- [50] Islem Saidani, Ali Ouni, Moataz Chouchen, and Mohamed Wiem Mkaouer. 2020. Predicting continuous integration build failures using evolutionary search. *Information and Software Technology* 128 (2020), 106392. <https://doi.org/10.1016/j.infsof.2020.106392>
- [51] Islem Saidani, Ali Ouni, and Mohamed Wiem Mkaouer. 2022. Improving the prediction of continuous integration build failures using deep learning. *Automated Software Engineering* 29 (2022).
- [52] Mark Santolucito, Jialu Zhang, Ennan Zhai, Jurgen Cito, and Ruzica Piskac. 2022. Learning CI Configuration Correctness for Early Build Feedback. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 1006–1017. <https://doi.org/10.1109/SANER53432.2022.00118>
- [53] Steven Sawyer. 2009. Analysis of Variance: The Fundamental Concepts. *Journal of Manual & Manipulative Therapy* 17 (4 2009), 27E–38E. <https://doi.org/10.1179/jmt.2009.17.2.27E>
- [54] Gideon Schwarz. 1978. Estimating the Dimension of a Model. *The Annals of Statistics* 6, 2 (1978), 461–464. <http://www.jstor.org/stable/2958889>
- [55] Tushar Sharma, Marios Fragkoulis, and Diomidis Spinellis. 2016. Does your configuration code smell?. In *Proceedings - 13th Working Conference on Mining Software Repositories, MSR 2016*. Association for Computing Machinery, Inc, 189–200. <https://doi.org/10.1145/2901739.2901761>
- [56] August Shi, Wing Lam, Reed Oei, Tao Xie, and Darko Marinov. 2019. iFixFlakies: A framework for automatically fixing order-dependent flaky tests. In *Proceedings of the 27th Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.
- [57] Jacek Śliwerski, Thomas Zimmermann, and Andreas Zeller. 2005. When do changes induce fixes? *ACM SIGSOFT Software Engineering Notes* 30, 4 (7 2005), 1–5. <https://doi.org/10.1145/1082983.1083147>
- [58] Jens Smeds, Kristian Nybom, and Ivan Porres. 2015. DevOps: A Definition and Perceived Adoption Impediments. 166–177. https://doi.org/10.1007/978-3-319-18612-2_14
- [59] Eliezio Soares, Daniel Alencar da Costa, and Uirá Kulesza. 2023. Continuous Integration and Software Quality: A Causal Explanatory Study. (9 2023).
- [60] Eliezio Soares, Gustavo Sizilio, Jadsou Santos, Daniel Alencar da Costa, and Uirá Kulesza. 2022. The effects of continuous integration on software development: a systematic literature review. *Empirical Software Engineering* 27, 3 (2022), 78.
- [61] Stephen Orban. 2015. Enterprise DevOps: Why You Should Run What You Build.
- [62] Ahmed Tamrawi, Hoan Anh Nguyen, Hung Viet Nguyen, and Tien N Nguyen. 2012. Build code analysis with symbolic evaluation. In *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 650–660.
- [63] Patanamon Thongtanunam, Shane McIntosh, Ahmed E Hassan, and Hajimu Iida. 2016. Revisiting Code Ownership and Its Relationship with Software Quality in the Scope of Modern Code Review. In *Proceedings of the 38th International Conference on Software Engineering (ICSE '16)*. Association for Computing Machinery, New York, NY, USA, 1039–1050. <https://doi.org/10.1145/2884781.2884852>
- [64] Bogdan Vasilescu, Yue Yu, Huaimin Wang, Premkumar Devanbu, and Vladimir Filkov. 2015. Quality and productivity outcomes relating to continuous integration in GitHub. In *Proceedings of the 2015 10th joint meeting on foundations of software engineering*. 805–816.
- [65] Carmine Vassallo, Gerald Schermann, Fiorella Zampetti, Daniele Romano, Philipp Leitner, Andy Zaidman, Massimiliano Di Penta, and Sebastiano Panichella. 2017. A tale of CI build failures: An open source and a financial organization perspective. In *2017 IEEE international conference on software maintenance and evolution (ICSME)*. IEEE, 183–193.
- [66] Carmine Vassallo, Fiorella Zampetti, Daniele Romano, Moritz Beller, Annibale Panichella, Massimiliano Di Penta, and Andy Zaidman. 2016. Continuous delivery practices in a large financial organization. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 519–528.

- [67] Mario Villamizar, Oscar Garcés, Lina Ochoa, Harold Castro, Lorena Salamanca, Mauricio Verano, Rubby Casallas, Santiago Gil, Carlos Valencia, Angee Zambrano, et al. 2017. Cost comparison of running web applications in the cloud using monolithic, microservice, and AWS Lambda architectures. *Service Oriented Computing and Applications* 11 (2017), 233–247.
- [68] Yuqing Wang, Maaret Pyhäjärvi, and Mika V Mäntylä. 2020. Test automation process improvement in a DevOps team: experience report. In *2020 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 314–321.
- [69] Anna Wiedemann, Manuel Wiesche, Heiko Gewalt, and Helmut Kremer. 2023. Integrating development and operations teams: A control approach for DevOps. *Information and Organization* 33, 3 (9 2023), 100474. <https://doi.org/10.1016/j.infoandorg.2023.100474>
- [70] Pavlina Wurzelová, Fabio Palomba, and Alberto Bacchelli. 2019. Characterizing women (not) contributing to open-source. In *2019 IEEE/ACM 2nd International Workshop on Gender Equality in Software Engineering (GE)*. IEEE, 5–8.
- [71] Jing Xia and Yanhui Li. 2017. Could We Predict the Result of a Continuous Integration Build? An Empirical Study. In *2017 IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C)*. IEEE, 311–315. <https://doi.org/10.1109/QRS-C.2017.59>
- [72] Jing Xia, Yanhui Li, and Chuanqi Wang. 2017. An Empirical Study on the Cross-Project Predictability of Continuous Integration Outcomes. In *2017 14th Web Information Systems and Applications Conference (WISA)*. IEEE, 234–239. <https://doi.org/10.1109/WISA.2017.53>
- [73] Jiaxin Zhu, Minghui Zhou, and Audris Mockus. 2014. Patterns of Folder Use and Project Popularity: A Case Study of Github Repositories. In *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '14)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/2652524.2652564>

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009