



An empirical study of token-based micro commits

Masanari Kondo¹ · Daniel M. German² · Yasutaka Kamei^{1,3} · Naoyasu Ubayashi⁴ · Osamu Mizuno⁵

Accepted: 2 July 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

In software development, developers frequently apply maintenance activities to the source code that change a few lines by a single commit. A good understanding of the characteristics of such small changes can support quality assurance approaches (e.g., automated program repair), as it is likely that small changes are addressing deficiencies in other changes; thus, understanding the reasons for creating small changes can help understand the types of errors introduced. Eventually, these reasons and the types of errors can be used to enhance quality assurance approaches for improving code quality. While prior studies used code churns to characterize and investigate the small changes, such a definition has a critical limitation. Specifically, it loses the information of changed tokens in a line. For example, this definition fails to distinguish the following two one-line changes: (1) changing a string literal to fix a displayed message and (2) changing a function call and adding a new parameter. These are definitely maintenance activities, but we deduce that researchers and practitioners are interested in supporting the latter change. To address this limitation, in this paper, we define *micro commits*, a type of small change based on changed tokens. Our goal is to quantify small changes using changed tokens. Changed tokens allow us to identify small changes more precisely. In fact, this token-level definition can distinguish the above example. We investigate defined micro commits in four OSS projects and understand their characteristics as the first empirical study on token-based micro commits. We find that micro commits mainly replace a single name or literal token, and micro commits are more likely used to fix bugs. Additionally, we propose the use of token-based information to support software engineering approaches in which very small changes significantly affect their effectiveness.

Keywords Empirical study · Micro commits · Tokens · Mining software repositories

1 Introduction

Commits changing a few lines of code are common in software development. Purushothaman et al. defined *small commits* as those modifying less than 10 lines in their study (Purushothaman and Perry 2005). They found that 50% of changes in the examined systems were small commits. They also reported that 10% of all commits were *one-line commits* (modified at

Communicated by: Burak Turhan

Extended author information available on the last page of the article

most one line). In a recent study, Alali et al. reported that in the GCC project, 19.9% of commits were *extra-small*, adding at most 5 lines of code (Alali et al. 2008). Our research found that in the projects we studied, between 6 and 8% of all commits were one-line commits (see Section 2).

Improving code quality is an ultimate goal for software engineering researchers, and several quality assurance approaches have been widely studied so far, such as automated program repair (APR), defect prediction, and fault localization. A good understanding of the characteristics of very small changes can support such approaches, as it is likely that such changes are addressing deficiencies in the system (Purushothaman and Perry 2005); thus, understanding the characteristics of creating small changes can help understand the types of errors that other changes introduce and potentially help with program repair. Eventually, the information can be used to enhance quality assurance approaches for improving code quality.

While prior studies (Purushothaman and Perry 2005; Hindle et al. 2008; Alali et al. 2008) use churn (number of lines added and removed) to identify small changes (e. g., small commits and one-line commits (Purushothaman and Perry 2005), or extra-small commits (Alali et al. 2008)), it has one significant limitation: they consider the line to be the finest-grained entity of changed source code. More specifically, such a definition overlooks the details of what has changed in a line (German et al. 2019; Meng et al. 2013; Servant and Jones 2017). For instance, when several lines have a small change (such as an identifier being renamed in a few places), these modifications might appear as one line added and one line removed for each change, rather than a single identifier change.

Another problem is that splitting or joining a line of code that is being modified can result in noise. For example, splitting a line into two would be reflected as a change to multiple lines in version control systems (e. g., Git), and this type of change can add noise to the analysis of the history of the development process.

These limitations sometimes cause researchers to fail in accurately quantifying small changes. For example, Listing 2 shows a commit in the Linux repository that changes a few lines (i. e., three added and two deleted lines). While this commit corresponds with a multiple-line change and may not correspond to a one-line commit, it only adds a token “static”. This is similar to Listing 1 corresponding with a one-line commit that only adds a token “static”. Studying the actual changed tokens instead of the lines can provide a better understanding of the characteristics of the small changes.

In this paper, we define a new class of commits: *micro commits*. **Micro commits are commits that add at most five tokens and remove at most five tokens of source code.** We aim to quantify small changes using the token-level definition (i. e., micro commits) rather than relying on the line-level definition (i. e., one-line commits). This token-level definition allows us to identify small changes more accurately, and use token information to characterize them. We conducted an empirical study on four large, mature open-source projects to: a) demonstrate that micro commits are common, accounting for between 7.45 and 17.95% of all studied commits in the studied projects, b) understand their qualitative and quantitative characteristics, and c) show our definition of micro commits (a threshold of 5 added and removed tokens) includes approximately 90% of all one-line commits, yet only approximately 40–50% of micro commits are one-line commits.

Specifically, we answer the following research questions (RQs). We have also provided a summary of the key findings for each RQ. The detailed results are described in Sections 4, 5, and 6.

RQ1: What are the characteristics of micro commits?

Motivation: This research question aims to explain their quantitative characteristics:

how frequent they are, and the types of tokens they delete and add.

Results: Most micro commits replace a single token with one of the same types, and this token type is mostly name (e. g., identifier names) or literal (e. g., numbers). Java and C differ on the most frequent tokens in micro commits.

RQ2: **What are the types of changes that micro commits perform?**

Motivation: We intend to understand the purpose of micro commits (e. g., changing control flow, replacing the name of a variable, and modifying an expression) and whether a micro commit performs one or more activities. Specifically, we manually inspected the changes applied to the source code to understand the purpose behind the micro commit and the occurrence of activities.

Results: More than 85% of micro commits apply a single operation to a single target. The four most common types of these micro commits are replacing an existing expression, identifier, constant, or declaration. Multi-operation micro commits usually change the order of statements.

RQ3: **How do micro commits compare to one-line commits?**

Motivation: This research question aims to explore the extent of differences between one-line commits and micro commits. Extracting micro commits requires syntactic parsing of the source code, which is more costly than extracting one-line commits. If they are identical, micro commits may be redundant.

Results: Most one-line commits are micro commits (approximately 89–93%). In contrast, only about 40–50% of micro commits are one-line commits. Indeed, 30–40% of micro commits include two or more hunks (one-line commits only have one hunk).

The main contributions of this paper are as follows:

- We propose the concept of micro commits as commits that add at most five tokens and remove at most five tokens, and demonstrate that these types of changes are common.
- We empirically investigate micro commits and understand their quantitative and qualitative characteristics. We especially shed light on the differences in micro commits between programming languages through our manual inspection.
- We propose the use of token-level information to support software engineering approaches that use extremely small changes (e. g., programing repair).
- We provide the replication package of this study that contains a set of micro commits that have been manually labeled according to their purpose.

The organization of our paper is as follows: Section 2 introduces motivating examples. Section 3 explains our studied dataset. Sections 4, 5, and 6 present the experiments and results based on our RQs. Section 7 proposes the use of token-level information. Section 8 describes the threats to the validity of our case study. Section 9 introduces related work. Section 10 presents the conclusion.

2 Motivating example

In this section, we provide an example of a one-line commit. Also, we demonstrate that they account for a non-negligible proportion of commits. Finally, we highlight the drawback of using lines of code to study extremely small changes, and we discuss how micro commits can address this drawback.

We first show the frequency of one-line commits in four OSS projects used in this study and confirm that it is consistent with Purushothaman and Perry (2005). As in Purushothaman and Perry (2005), we use the diffs generated by Git to identify one-line commits. Table 1 shows

Table 1 The proportion of one-line commits in the studied projects

Project	#total commits	#studied commits	#one-line commits	Proportion(%)
Camel	60,911	38,458	2,405	6.25
Hadoop	69,997	53,796	2,302	4.28
Linux	1,048,688	802,726	65,858	8.20
Zephyr	40,883	25,542	1,979	7.75

Listing 1 An example micro commit in Linux
retrieved from: 092734b4bb227faddf241b116af14357645d963c

```
@@ -385 +385 @@EXPORT_SYMBOL(bt878_device_control);
-struct cards card_list[] __devinitdata = {
+static struct cards card_list[] __devinitdata = {
```

Listing 2 An example micro commit with multiple changed lines
in Linux retrieved from: 0ce6e62bd6591777bd92873e2db93fdbc5228122

```
@@ -1143,2 +1143,3 @@ int path_lookup_open(const char *name, unsigned int lookup_flags,
-int path_lookup_create(const char *name, unsigned int lookup_flags,
- struct nameidata *nd, int open_flags, int create_mode)
+static int path_lookup_create(const char *name, unsigned int lookup_flags,
+ struct nameidata *nd, int open_flags,
+ int create_mode)
```

the proportion of one-line commits. The proportion was computed by using the “#studied commits” column. It only shows the commits that have made changes to the source code. Our analysis is conducted based on these commits. The detailed procedure for extracting commits is explained in Sections 3.2 and 3.3.¹ We observe 4.28–8.20% of one-line commits. Specifically, the proportion is more than 7% in the Linux and Zephyr projects. Hence, one-line commits account for a non-negligible proportion of all commits.

Listing 1 shows an example of a one-line commit, also known as a micro commit. This commit adds a static modifier into a struct definition, and this is not adding a functionality but fixing the code.

However, some extremely small changes are often obscured by splitting or joining lines of code, making them appear more complex than they are. For example, Listing 2 shows an example of a micro commit that is not a one-line commit. This commit semantically adds a static modifier only; however, this commit includes multiple changed lines because of changing the format of the definition of the variable.

Listings 1 and 2 are semantically identical, but one-line commits cannot include Listing 2 because it modifies multiple lines. Because we used Git, we deduced that the diff algorithms could address this limitation. Git has four algorithms to compute diffs, and they exhibit different results (Nugroho et al. 2020). Hence, we investigated four algorithms: *patience*, *minimal*, *histogram*, and *myers* described in the Git manual page.² However, all algorithms generate the same diff. Hence, one-line commits may overlook such commits. If these diffs are analyzed with finer-grained source code entities (e.g., AST), it is easy to realize these

¹ Because of the differences in source code management tools, one-line changes in the prior study (Purushothaman and Perry 2005) and our one-line commits are slightly different.

² <https://git-scm.com/docs/git-diff>

commits have the same intention (i. e. perform the same change). However, AST analysis is expensive, particularly in repositories such as Linux that has more than one million commits and more than 60k source code files.

Therefore, to address this limitation, we define *micro commits* based on tokens. Because tokens are the semantically finest-grained source code entity, micro commits based on tokens can cover ones overlooked by one-line commits. Indeed, Listings 1 and 2 change one token only; thus, they both perform the same change in two different lines of code.

3 Dataset preparation

3.1 Studied datasets

To answer our RQs, we conducted an empirical study on four notable large OSS projects written in Java and C: Camel,³ Hadoop,⁴ Linux,⁵ and Zephyr.⁶ The Camel project is an integration framework that provides a routing engine to integrate systems. The Hadoop project is a distributed computing framework. The Linux project (a.k.a. the Linux Kernel) is one of the most popular open-source operating system kernels. The Zephyr project is a real-time operating system supporting several architectures. Hence, these include four software systems: an integration framework, a distributed computing framework, an operating system kernel, and an operating system. We selected these four projects because of three reasons: (1) they are written in popular programming languages (i.e., Java and C), (2) they are well-known popular OSS projects, and (3) they have a long development history.

3.2 One-line commits and micro commits

Our research aims to accurately quantify small changes using a token-level definition (i. e., micro commits). Additionally, to highlight the differences in accuracy between token-level and line-level definitions, we should compare micro commits with one-line commits. Therefore, we detail the process of extracting one-line commits and micro commits from software development histories below.

Git is language agnostic. The changes performed in a commit are displayed as a diff, comparing the code before and after the commit. These changes are grouped into *hunks*. A hunk is a set of contiguous lines that are added/removed/modified together, along with metadata that indicates its context—where the change occurred. Each hunk can include *context* lines (i. e., lines that were not modified but are used to help interpret the change). The default number of context lines is three, but for the purpose of this paper, we have set it to zero; thus, we ignore context lines in the hunk. Git's diff does not present lines that have been modified. Instead, it simply records lines that have been removed (prefixed with “-”) and lines that have been added (prefixed with “+”); thus, a modified line is represented by a removed line and its corresponding added line. If several continuous lines are modified simultaneously, Git presents first all removed lines, and thereafter the added lines.

³ <https://camel.apache.org/>

⁴ <https://hadoop.apache.org/>

⁵ <https://www.linux.org/>

⁶ <https://www.zephyrproject.org/>

Listing 3 An example commit in a line repository

```
@@ -10 +10 @@ test();  
-int flg = 10;  
+static int flag = 10;
```

Listing 4 An example commit in a token repository

```
@@ -100,2 +100,3 @@ test();  
+specifier|static  
  name|int  
-name|flg  
+name|flag
```

We extracted one-line commits based on the hunks provided by Git. Specifically, one-line commits correspond to commits that have a diff with exactly one removed and one added line **in the same hunk**. Listing 1 is an example of such a commit.

To be able to perform token-level analysis, we processed the repository history using `cregit` (German et al. 2019). This uses `srcML` (Collard et al. 2013) to generate an equivalent commit history where the differences are displayed as changes to sequences of tokens instead of lines (see German et al. (2019) for a detailed description). Effectively, we track tokens removed and/or added during a commit and can easily identify commits that have added and/or removed a certain number of tokens. Similarly to the way we can identify modified lines, we can identify modified tokens if one token is added and another is removed in the same hunk. For example, the commit from Listing 3 is shown in its equivalent token version in Listing 4.

We extracted micro commits based on the hunks provided by Git repositories processed by `cregit`. Micro commits refer to commits that include a maximum of five added tokens and five deleted tokens across all hunks. This number was chosen for the following reasons.

- In the languages being studied (C and Java), it is highly unlikely to add a new statement with only five tokens, suggesting that such commits carry out minor modifications. For example, within five tokens, developers can only add a function call with one parameter and an ending semicolon: `name (parm) ;` includes two identifiers, two parentheses, and one semicolon.
- In the systems we studied, between 7.45 and 17.95% of all studied commits add at most 5 tokens and remove at most 5 tokens.

This number serves as a parameter for micro commits. For example, we use the same number for both added and deleted tokens while different numbers could be used. Its potential threats are discussed in Section 8.2.

Source code comments are important for source code and making changes to comments are also maintenance activities. However, in this paper, we exclude comments and execute our analysis. The reason is to prioritize maintenance activities for code logic. As mentioned in Section 1, our intention is to support various software engineering approaches (e.g., defect prediction), which typically prioritize code logic over comments. Indeed, defect prediction studies typically do not take into account comment issues when identifying target defects (Kondo et al. 2020; Hoang et al. 2019; McIntosh and Kamei 2018). While we

acknowledge the importance of changes made to comments for maintenance purposes, this perspective is beyond the scope of our paper.

3.3 Data collection

We preprocessed the commits in the studied repositories and constructed a database with its diffs (both line-based and token-based) using the following steps. From this database, we extracted one-line commits and micro commits.

Step 1: For each commit, extract the line-based diff of its modified source code ignoring any changes to non-source code:

- Remove changes to non-source files.⁷
- Remove changes to comments and white space using regular expressions (e.g., “//.*”).
- Remove commits that do not have any changes after the aforementioned processes.

Step 2: Using `cregit`, for each commit, extract the token-based diff of its modified source code ignoring any changes to non-source code:

- Remove changes to comments.
- For each source code token, keep its type and its value. `cregit` tokenizes the source code using `srcML`.⁸ Thus, the types of tokens are those created by `srcML`. For example, `int i;` will be converted to the sequence of `type|value: name|int, name|i, decl_stml |;`

Step 3: Create a database in SQLite with these commits (line and token-based) including:

- Identify and store each hunk and its metadata (such as the file where it occurred and the number of lines/tokens added and removed).
- Added and removed lines or tokens in each hunk
- Commit messages
- Metadata (e.g., index)

In summary, we record for each line-based diff: its commit id and its set of hunks (for each hunk, its location, number of lines added, number of lines removed, and its contents as a sequence of added/removed lines). We record the same for token-based diffs (replacing lines with tokens—including their types). Note that when obtaining diffs with Git, we use the *myers* algorithm, which is the default algorithm. Also, we record commit messages. More details can be found in our replication package (see Section 8.1).

Table 2 displays the number of extracted micro commits and one-line commits. We used these micro commits and one-line commits in this study. We found that micro commits can cover approximately 90% of one-line commits. In contrast, only approximately 40% (for Linux and Zephyr) or 50% (for Camel and Hadoop) of micro commits can be covered by one-line commits.

As shown in Table 3, between 7.45% and 17.95% of all studied commits are micro commits, and approximately 1 in 3 or 4 micro commits are one-token commits in all projects. Hence, micro commits constitute a non-negligible portion of all studied commits.

⁷ We extract files with the extension of “java” in Camel and Hadoop and with the extension of “c” and “h” in Linux and Zephyr.

⁸ <https://www.srcml.org/>

Table 2 The number and proportion of the intersection between one-line commits and micro commits in each commit type (i. e., one-line or micro)

Project	#intersects	#one-line	#micro	%one-line	%micro
Camel	2,131	2,405	4,230	88.6	50.4
Hadoop	2,069	2,302	4,010	89.9	51.6
Linux	59,836	65,858	138,142	90.9	43.3
Zephyr	1,849	1,979	4,585	93.4	40.3

The column of “#intersects” indicates the intersection; the columns of “#one-line” and “#micro” indicate the number of one-line commits and micro commits; the column of “%one-line” and “%micro” indicate the proportion of intersection in each commit type (i. e., one-line commits and micro commits)

As expected, most of these commits modify a few lines: between 52.80% and 58.70% modify add or remove at most one line, and between 59.56% and 67.48% add-or-remove two lines.

4 RQ1: What are the characteristics of micro commits?

4.1 Approach

The goal of RQ1 is to understand the characteristics of micro commits. More specifically, we investigated the modified tokens.

In this RQ, we investigated micro commits from two perspectives: (1) *most frequently modified tokens and token types by micro commits* and (2) *modification patterns for each micro commit*. We first count added and removed tokens and their token types from all micro commits and provide researchers with tokens and token types frequently modified by micro commits. Second, we investigate the set of added and removed tokens for every micro commit and show the common *modification patterns* adopted by a single micro commit. Note that we used the set rather than the sequence of tokens. Hence, we characterized modification patterns based on modified tokens in micro commits rather than the sequences of modified tokens.

We use srcML classification for the types of tokens. For example, tokens of type names correspond to names of types and variables (including language predefined ones); literals are constant values; operators are operators to perform mathematical operations; `argument_list` corresponds to either `()` (empty parameter list), or each of the parenthesis around parameters or the comma that separates them; `expr_stmt` is the semicolon at the end of the statement; `block` is a `{ }` or `;`; `file` is a filename; `specifier` is a C storage specifier (e.g., `static`); `directive` a C preprocessor directive; and `annotation` corresponds to Java annotations. The right-hand

Table 3 Number of micro commits and one-token commits and their proportion with respect to all source-code commits

Project	Micro commits	Prop (%)	One-token commits	Prop (%)
Camel	4,230	11.00	1,319	3.43
Hadoop	4,010	7.45	1,288	2.39
Linux	138,142	17.21	32,973	4.11
Zephyr	4,585	17.95	1,247	4.88

side of C macro definitions is not further parsed by srcML and is considered a single token of type value (i. e., the value the macro expands to).

4.2 Results

(1) Most frequently modified tokens and token types by micro commits

The top-3 most frequently touched token types in micro commits are generally the name, literal, and operator token types. Figure 1 shows the frequently added/removed token types by micro commits that account for more than 5% in all projects. We found three token types, which we refer to as the top-3 most frequently touched token types. The token type most frequently included in micro commits is the name token (e.g., the name of a variable or function), the second one is the literal token (e.g., 123, ‘a’, “test”), and the third one is the operator token (such as +). Also, the proportion of the literal token is significantly different between Java and C. Specifically, while the proportion of name tokens is more than three times larger than that of literal tokens in the projects written in C, the difference is less than two times in the projects written in Java. While the proportion of the operator tokens is relatively small, these are also included in the top-3 most frequently touched token types. Hence, micro commits usually modify name, literal, and operator tokens in most cases, but their proportions may differ between programming languages and their token types.

While the tokens corresponding to the top-3 token types differ between Java and C, we observe similar tokens within the same language. Figure 2 shows the top-10 most frequently occurring tokens for the top-3 token types. In Java, boolean literals (e. g., true/false, null), and numeric literals were the most commonly observed, while in C, they were the tokens for 0/1, parentheses and names for types (e. g., int, u32_t, and u8_t).

In conclusion, the types of tokens most frequently changed are the same in both programming languages, but the actual tokens are different.

(2) Modification patterns of micro commits

The single token modification is the most frequently observed pattern in the studied micro commits. Table 4 shows the top-5 most frequently appearing sets of removed and added tokens in micro commits. Each row indicates a set of token types modified by a single micro commit and their frequency and proportion (i. e., # of micro commits). The “n” column indicates the frequency, while “Pro” indicates the proportion. In this paper, we use the same column name in the other tables. In all projects, the most frequently observed micro commits consist of an added and removed token. For example, in the Linux project, micro commits adding and removing a name token are the most frequently observed. This type of single addition and removal usually represents a single token being replaced (e. g., Listing 5) and is the most frequently observed type of micro commit in all projects.

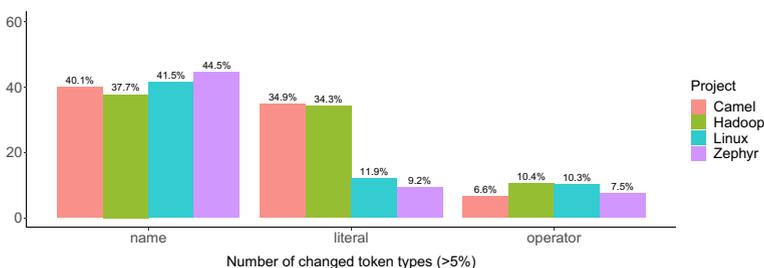
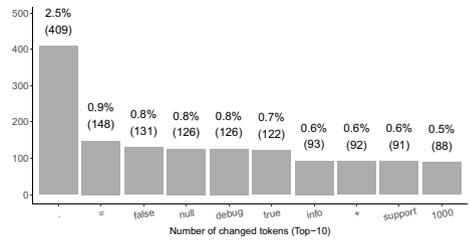
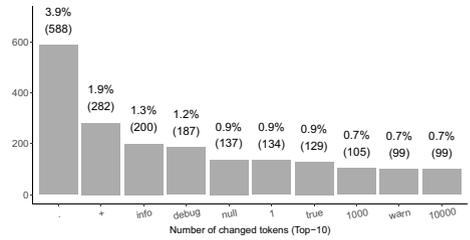


Fig. 1 Proportions of changed token types (> 5%)

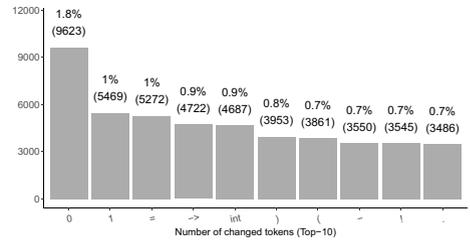
Fig. 2 Numbers of changed tokens



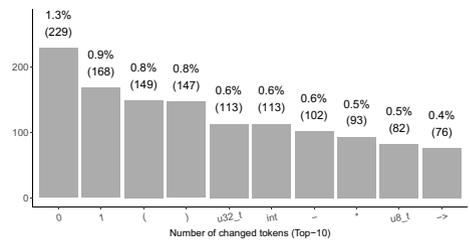
(a) Camel (Java)



(b) Hadoop (Java)



(c) Linux (C)



(d) Zephyr (C)

Listing 5 An example single token modification in a token repository

```
@@ -100,1 +100,1 @@ test();
-name | flg
+name | flag
```

Similar to the results from (1), the modified tokens differ between Java and C. Modifications of literals are the most frequent pattern in Java, accounting for approximately 20% of all micro commits. Modifications of names are the most common pattern in C, accounting for about 10% of all micro commits.

Table 4 Top-5 added and removed token types applied to source code by micro commits

Project	Add	Remove	n	Pro
Camel (Java)	literal	literal	833	19.7
	name	name	435	10.3
	literal,literal	literal,literal	312	7.4
	name,name	name,name	269	6.4
	name,name,name	name,name,name	125	3.0
Hadoop (Java)	literal	literal	761	19.0
	name	name	408	10.2
	literal,literal	literal,literal	266	6.6
	name,name	name,name	229	5.7
	specifier	-	126	3.1
Linux (C)	name	name	13693	9.9
	literal	literal	7350	5.3
	value	value	6835	4.9
	name,name	name,name	6141	4.4
	specifier	-	4171	3.0
Zephyr (C)	name	name	487	10.6
	value	value	352	7.7
	name,name	name,name	260	5.7
	name,name,name	name,name,name	170	3.7
	literal	literal	169	3.7

Summary of RQ1

Most micro commits modify a single token, and this token type is either a name, a literal, or an operator. The distribution of micro commits of each of these types is different in C and Java. The operators being replaced are also significantly different across languages.

5 RQ2: What are the types of changes that micro commits perform?

5.1 Approach

This RQ aims to understand the details of the activities performed by micro commits. Specifically, we manually inspect a large set of micro commits to understand what types of change were performed from a source code perspective, considering removed and/or added tokens. Such an understanding gives us insight into whether understanding micro commits can support several approaches in software engineering (see the details in Section 7).

Our manual inspection consisted of two phases: (1) *constructing a coding guide* and (2) *manual classification*. Constructing a coding guide for manual classification/annotation is a common practice in the field of mining software repositories (Hata et al. 2019; Wang et al. 2021; Chouchen et al. 2021; Hata et al. 2022; Wang et al. 2023; Zanaty et al. 2018). To create the coding guide, we referred to previous studies (Hata et al. 2019; Wang et al. 2021; Chouchen et al. 2021; Hata et al. 2022; Wang et al. 2023; Zanaty et al. 2018) and followed the process detailed below.

The initial coding guide was first discussed by the first and second authors. Since this is the first study to classify micro commits, we examined both micro commits and other types of commits to develop the initial coding guide. After constructing the initial coding guide, we aimed to reach a consensus among the first three authors for this guide and refine the guide. Specifically, we independently annotated 20 micro commits from a subset of all micro commits. This subset consists of micro commits that only change less than or equal to five tokens in main files (.c or .java files) in the Linux, Hadoop, and Zephyr projects to investigate a single operation commit for refining the coding guide. We computed the agreement rate for these 20 micro commits using Fleiss' Kappa (Fleiss 1971) that is used to demonstrate inter-rater agreement when there are more than two raters. It is also frequently applied in the field of mining software repositories (Chouchen et al. 2021; Hata et al. 2022). The Kappa coefficient is commonly interpreted using the following scale (Viera and Garrett 2005): Slight agreement ($0.01 \leq k \leq 0.20$), Fair agreement ($0.21 \leq k \leq 0.40$), Moderate agreement ($0.41 \leq k \leq 0.60$), Substantial agreement ($0.61 \leq k \leq 0.80$), Almost perfect agreement ($0.81 \leq k \leq 0.99$). Then we discussed the coding guide along with any inconsistencies in categorization to reach a consensus. We repeated this process until our categorization substantially matched, indicating that our coding guide was successfully constructed. We, therefore, repeated this process three times (i.e., independently classifying 60 commits). Finally, our agreement rate achieved substantial agreement in two consecutive iterations. Table 5 shows all agreement rates across three authors for each repetition.

Through this process, we identified two perspectives: *operation* and *target*. The operation indicates what kind of operations are applied, such as adding a new statement or changing an expression; the target indicates source code entities where the operation is applied, such as expressions. The details of the coding guide are described below.

We utilized the following coding guide to categorize micro commits in terms of the operations.

- *add*: This refers to operations that add a completely new entity.
- *replace*: This refers to operations that modify an entity.
- *remove*: This refers to operations that completely remove an entity.
- *multi*: This code indicates that multiple operations are applied.
- *no*: This code indicates that no functional change is applied.

We utilized the following coding guide to categorize micro commits in terms of the targets.

- *identifier*: This refers to commits that only modify identifiers, such as variable names. If other entities, such as parentheses, are included, it would not be labelled as an identifier but would be considered an expression.

Table 5 Fleiss' Kappa scores for each repetition

Criteria	First Time	Second Time	Third Time
Operations	0.686	0.669	0.832
Targets	0.425	0.671	0.754

Listing 6 Example “replace identifier” commit diff retrieved from f72e6c3e17be568138d8e4855ac2734d251a6913 in Linux.

```
- strcpy(drvinfo->bus_info, pci_name(mdev->pdev),
+ strcpy(drvinfo->bus_info, dev_name(mdev->device),
```

Listing 7 Example “replace expression” commit diff retrieved from 8b58f261113c442717b9d205ab187e51c3823597 in Linux.

```
- dqm->total_queue_count++;
+ dqm->total_queue_count--;
```

- *statement*: This refers to commits that modify a complete statement, including the semicolon (;), such as an entire function call with its semicolon. C’s #include preprocessor statement is also regarded as a statement.
- *constant*: This refers to commits that only modify literals, such as strings or numbers. If other entities, such as parentheses, are included, they would not be identified as a constant but would be considered an expression.
- *declaration*: This refers to commits that modify declarations, such as variable declarations. However, if the commit can be classified as “identifier” or “constant”, it should be categorized under these two categories rather than “declaration”.
- *control flow*: This refers to commits that modify the control flow of execution, such as adding a new “else” statement.
- *expression*: If a commit does not match other categories and involves modifying a part of a statement, it would be classified into this category. Additionally, this category includes transformations from constants to variables or vice versa, as well as conversions from a variable to a pointer and vice versa.
- *multi*: This code suggests that operations are performed on multiple targets.
- *no*: This code indicates that no functional change is applied.

Table 6 shows the summary and examples. This represents different types of activities performed by micro commits. Let us describe two example commits. Listing 6 shows an example commit. This commit changes a function call and its argument. More specifically, the identifiers of the function call and the argument value are replaced so that we classify this commit as operations=replace, and targets=identifier. The commit of Listing 7 replaces an expression “++” into “--”. Hence, we classify this commit as operations=replace, and targets=expression.

Because our agreement rates for operation and target based on our coding guide achieved almost perfect and substantial agreement respectively (Table 5) and we made an internal consensus of the coding guide, only the first author manually classified the 400 micro commits, similar to previous studies (Wang et al. 2021, 2023; Hata et al. 2019, 2022). All manual categorizations are available in our sheet in our replication package.⁹ The sample size in manual inspection was determined as a statistical representative with a confidence level of 95% and a confidence interval of 5% for 150,967 micro commits from all studied projects.¹⁰ The minimum sample size with this confidence level and this confidence interval is 383. For

⁹ <https://doi.org/10.5281/zenodo.10963270>

¹⁰ <https://www.surveysystem.com/sscalc.htm>

Table 6 The description of each candidate in our manual inspection

Criteria	Candidate	Description	Example Commits and Their Diffs in Linux
Operations	add	Add a new entity	122503683169b21d9cdb90380a20caad7ba994b6 Diff: Listing 11
	replace	Replace an existing entity	b7a90e8043e7ab1922126e1c1c5c004b470f9e2a Diff: Listing 12
	remove	Remove a completely existing entity	b95b4e1ed92a203f4bdfc55f53d6e9c2773e3b6d Diff: Listing 13
	multi	Operations on multiple targets	8df0cfe6c6c4a9355989baa8de9f166b2bc51f76 Diff: Listing 8
	no	Non-functional modification	a092532483e3200a53c8b1170b3988cc668c07ef Diff: Listing 14
Targets	declaration	Change in a type signature	36f062042b0fd9f8e41b97a472f52139886ca26f Diff: Listing 15
	constant	A constant (e. g., literal)	1db76c14d215c8b26024dd532de3daf66ea30f7 Diff: Listing 16
	identifier	An identifier (e. g., function calls)	70e8b40176c75d3544024e7c934720b11a8a11bf Diff: Listing 17
	control flow	Modifies the control flow	415a1975923722f729211a9efca550c60c519bf3 Diff: Listing 18
	statement	A the majority of a statement (delimited by semicolon)	b95b4e1ed92a203f4bdfc55f53d6e9c2773e3b6d Diff: Listing 13
	expression	A part of a statement and not classified into other categories	40cc394be1aa18848b8757e03bd8ed23281f572e Diff: Listing 19
	multi	Multiple targets are altered	8df0cfe6c6c4a9355989baa8de9f166b2bc51f76 Diff: Listing 8
	no	Non-functional modification	a092532483e3200a53c8b1170b3988cc668c07ef Diff: Listing 14

safety, we also inspect 17 additional micro commits. Therefore, we classify a total of 400 micro commits.

5.2 Results

Micro commits usually perform a single operation (85.75%). Micro commits with multiple operations account for the remaining 14.25%, and they usually correspond to two operations. Table 7 lists the proportion of micro commits classified into “multi” in the operation and target. Approximately 86% of micro commits are classified into single operations

Table 7 Proportion of micro commits having multi activities

Single	Multi
85.75%(343)	14.25%(57)

Listing 8 Example “multi” micro commit diffs retrieved from 8df0cfe6c6c4a9355989baa8de9f166b2bc51f76 in Linux.

```

@@ -111,0 +112,5 @@
+ * - EXITCON_PROP_USB_SS (SuperSpeed)
+ * @type: integer (intval)
+ * @value: 0 (USB/USB2) or 1 (USB3)
+ * @default: 0 (USB/USB2)
+ *
@@ -114,0 +120 @@
+#define EXITCON_PROP_USB_SS 2
@@ -117 +123 @@
-#define EXITCON_PROP_USB_MAX 1
+#define EXITCON_PROP_USB_MAX 2
    
```

(i. e., non “multi”); thus, micro commits usually modify an extremely small section. We refer to such commits as *single-operation micro commits* (e. g., Listing 5 is a replace identifier). We surprisingly observe that approximately 14% are classified into *multi-operation micro commits* (i. e., “multi”). Hence, multi-operation micro commits account for a non-negligible portion of the micro commits. Listing 8 shows an example of a multi-operation micro commit. This commit has an add declaration and a replace constant. We disregard the first hunk because it only contains comment lines. It should be noted that such multi-operation micro commits usually have two operations only; we observed that only two commits contain more than two operations. Thus, based on our manual inspection, it is rare to find micro commits containing more than two operations.

Approximately 82.2% of single-operation micro commits replace existing tokens.

Table 8 summarizes the frequency of the combination of operations and targets in single-operation micro commits. The top-4 combinations include the “replace” operation, accounting for approximately 78.7% (270/343). Also, all the “replace” operation commits

Table 8 The frequency of the combination of operations and targets in single-operation micro commits

Operation	Target	n	Pro
replace	expression	85	24.8
replace	identifier	69	20.1
replace	constant	59	17.2
replace	declaration	57	16.6
add	statement	22	6.4
replace	control flow	12	3.5
no	no	8	2.3
remove	statement	8	2.3
add	expression	7	2.0
remove	declaration	7	2.0
add	control flow	4	1.2
add	identifier	2	0.6
remove	control flow	1	0.3
remove	expression	1	0.3
remove	identifier	1	0.3

Listing 9 Example “multi” commit diff retrieved from a71bfb4a6aabfe5e6f145883020153103c7fdfba in Linux.

```
-error_free_data:
-   free(data);
  error_free_buffer_access:
    free(buffer_access);
+error_free_data:
+   free(data);
```

account for 82.2% (282/343). This result suggests that many single-operation micro commits modify the existing source code, but do not add or remove the source code.

Multi-operation micro commits more frequently add-and-remove statements rather than single-operation ones. Table 9 summarizes the frequency of the combination of operations and targets in multi-operation micro commits, and Table 10 summarizes the frequency of the pair of their combination for each commit. The main difference from single-operation micro commits is that the “add statement” and the “remove statement” are top-3 (Table 9). The reason is that multi-operation micro commits that add and remove statements appear most frequently (Table 10). This type of commit is used to move the statements, and therefore, change the order of execution and potentially the control flow of the program. For example, Listing 9 shows an example micro commit. This commit adds and removes a statement and changes the order of execution of the free statement to fix a bug related to freeing data. This swapping activity is frequently observed in multi-operation micro commits.

Summary of RQ2

More than 85% of micro commits apply a single operation to a single target, and they mainly replace existing target. Multi-operation micro commits frequently change the order of statements.

Table 9 The frequency of the combination of operations and targets in multi-operation micro commits

Operation	Target	n	Pro
replace	identifier	22	19.1
replace	expression	19	16.5
add	statement	16	13.9
remove	statement	16	13.9
replace	constant	15	13.0
add	expression	7	6.1
replace	declaration	5	4.3
remove	expression	4	3.5
add	declaration	3	2.6
replace	control flow	3	2.6
add	control flow	2	1.7
remove	declaration	2	1.7
remove	control flow	1	0.9

Table 10 The frequency of the pair of the combination of operations and targets in multi-operation micro commits

Op#1	Tar#1	Op#2	Tar#2	Op#3	Tar#3	n	Pro
remove	statement	add	statement	-	-	11	19.3
replace	identifier	replace	expression	-	-	6	10.5
replace	identifier	replace	constant	-	-	6	10.5
add	expression	replace	identifier	-	-	4	7.0
replace	expression	replace	constant	-	-	4	7.0
remove	statement	replace	expression	-	-	3	5.3
add	expression	replace	constant	-	-	2	3.5
add	control flow	replace	expression	-	-	2	3.5
replace	identifier	replace	declaration	-	-	2	3.5
replace	expression	replace	control flow	-	-	2	3.5
add	declaration	remove	declaration	-	-	2	3.5
add	statement	remove	expression	replace	identifier	2	3.5
replace	constant	remove	expression	-	-	1	1.8
add	statement	replace	identifier	-	-	1	1.8
remove	statement	replace	control flow	-	-	1	1.8
add	statement	replace	declaration	-	-	1	1.8
replace	declaration	replace	constant	-	-	1	1.8
add	statement	replace	expression	-	-	1	1.8
add	statement	remove	control flow	-	-	1	1.8
add	declaration	replace	constant	-	-	1	1.8
remove	statement	replace	identifier	-	-	1	1.8
add	expression	remove	expression	-	-	1	1.8
replace	declaration	replace	expression	-	-	1	1.8

6 RQ3: How do micro commits compare to one-line commits?

6.1 Approach

As discussed in Sections 1 and 2, one-line commits are common in software development and often address deficiencies in the system (Purushothaman and Perry 2005). However, they have a drawback: they overlook changes within a line. Consequently, two commits with the same number of changed tokens could differ; one might be a one-line commit and the other might not. The concept of micro commits, introduced in this paper, address their drawbacks.

However, the extent of the differences between one-line commits and micro commits is unclear. Extracting micro commits is more costly than one-line commits as it requires syntactic parsing of the source code. Hence, this RQ aims to compare the one-line commits and micro commits.

Our methodology can be summarized as follows: we start by analyzing the changed tokens in one-line commits. We then analyze the modified hunks in micro commits. Finally, we discuss the intersection between one-line commits and micro commits.

6.2 Results

Approximately 90% of one-line commits consist of at most five tokens. Table 1 shows the number of one-line commits. As described in Section 2, there are a non-negligible number of these commits in the studied projects (4.28–8.20%). Figure 3 shows the proportion of one-line commits according to the number of tokens that they have added and removed between 0 and 10. As can be seen, there are a significant number of one-line commits that remove and add exactly one token (between approximately 50 and 63% of all one-line commits). Furthermore, except for the case in the Hadoop project where no commits add or delete five tokens, all cells with five or fewer added and deleted tokens have more than one one-line commit across all projects. This implies that there are no empty cells within five added or deleted tokens except for one cell in the Hadoop project. Also, the distribution of one-line commits, with more than five tokens, varies across the projects. For instance, in the Hadoop and Zephyr projects, there are cells with no one-line commits of more than five deleted tokens and less than or equal to one added token. In contrast, every cell in the Camel and Linux projects has at least one one-line commit. Hence, the majority of one-line commits add or remove at most five tokens, and this finding is generally consistent across all projects.

Figure 4 shows the accumulated distribution of one-line commits according to the maximum number of tokens they add or remove. We use the maximum number of tokens added or removed in this figure. This is because our definition of a micro commit applies the same threshold of five tokens to both the number of added and removed tokens. As can be seen, between approximately 57% and 65% add-or-remove at most one token, between 76% and 82% add-or-remove at most three tokens, and between 89% and 93% add-or-remove at most five tokens. Thus, approximately 90% of one-line commits can be covered by our micro commits.

The number of modified hunks is also a crucial characteristic of commits. By our definition, one-line commits only modify one location in the source code (i. e., one hunk). We define micro commits based on the number of tokens, so even if a commit is spread across multiple locations (i. e., multiple hunks), it is still considered a micro commit if the number of modified tokens is below a certain threshold. This is a significant distinction compared to one-line commits. Therefore, we do not impose any limits on the number of modified hunks.

Figure 5 illustrates the accumulated distribution of the number of hunks included in micro commits to investigate their difference from one-line commits. Approximately 70% (Linux

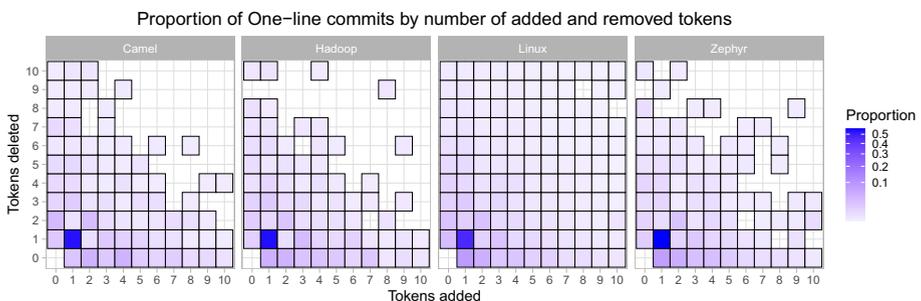


Fig. 3 Proportion of **one-line commits** by the number of tokens added or removed. The x and y-axis show the added and deleted tokens, and each cell indicates the proportion of commits

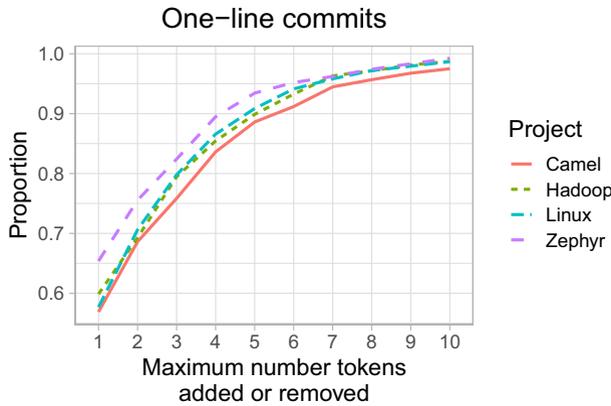


Fig. 4 Accumulated distribution of one-line commits in terms of the maximum number of added or removed tokens

and Hadoop) or 60% (Zephyr and Camel) of micro commits contain a single hunk, while the remaining commits encompass two or more hunks. Hence, while approximately 70% or 60% of micro commits share characteristics with one-line commits, the remaining 30% or 40% represent commits that one-line commits do not detect, even if they modify the same number of tokens.

In conclusion, although micro commits can encompass nearly all one-line commits, the reverse is not typically true: one-line commits do not generally cover micro commits. Indeed, Table 2 in Section 3.3 reveals that around 90% of one-line commits can be encapsulated by micro commits. However, only approximately 40% (for Linux and Zephyr) or 50% (for Camel and Hadoop) of micro commits can be encapsulated by one-line commits. Therefore, micro commits provide new insights compared to one-line commits.

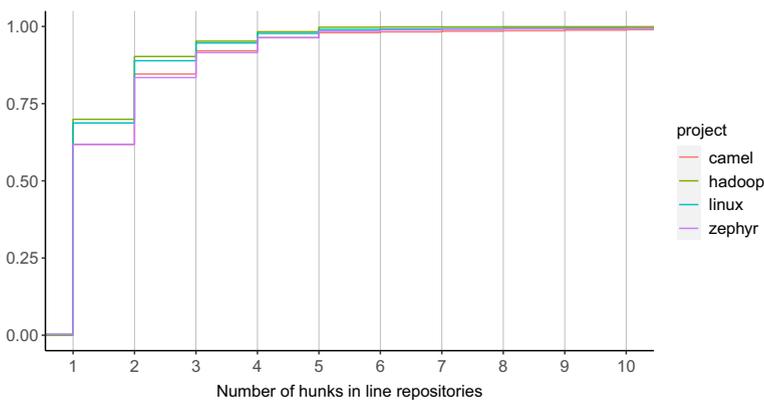


Fig. 5 Accumulated distribution of micro commits ($N = 5$) in terms of the number of hunks included

Listing 10 Example “replace constant” commit diff retrieved from c143708acfb17e91c5e4fc9bd9b496fc7d2db29c in Hadoop.

```
@@-71+71@@ protected void render(Block html) {
-   html.h1()._("Invalid log start value: " + S("end"))._();
+   html.h1()._("Invalid log end value: " + S("end"))._();
```

Summary of RQ3

Approximately 90% of one-line commits add or remove at most five tokens. Therefore, nearly all one-line commits can be covered by micro commits. In contrast, 30 to 40% of micro commits include two or more hunks that are not covered by one-line commits. In fact, only approximately 40% (for Linux and Zephyr) or 50% (for Camel and Hadoop) of micro commits can be encapsulated by one-line commits. Therefore, the characteristics of micro commits can help us understand the attributes of small changes, including those in one-line commits and commits not identified by one-line commits.

7 Discussion

In this paper, we defined micro commits and shed light on their characteristics. The main motivation is that turning our attention to micro commits would benefit software engineering research. In this section, we describe the implications of micro commits on future research.

7.1 Line-based vs token-based complexity metrics

Line-based complexity metrics (e. g., LA, LD, and LT) (Kamei et al. 2013) are one of the most popular source code complexity metrics in software engineering. However, **there are limitations to using line-based complexity metrics, as they may overlook capturing commits with equivalent complexity in terms of changed tokens.** We have shown examples with different numbers of changed lines but the same number of changed tokens (Listings 1 and 2). We have shown that in the projects under analysis, approximately 90% of one-line commits are micro commits, but only approximately 40–50% of micro commits are one-line commits (Table 2).

Thus, **future research should consider tokens (and their types) as an additional metric of the complexity of commits.** Because, for example, prior studies (Kondo et al. 2020; McIntosh and Kamei 2018) in defect prediction reported that current models heavily rely on the added lines, such new metrics would provide new information to identify defective commits accurately.

7.2 Micro commits are non-negligible and should be further studied

In Section 3.3, we showed that micro commits account for between 7.45 and 17.95% of all studied commits, which is quite high. Furthermore, 1 in 3 or 4 these changes (2.39 and 4.88% of all studied commits) simply change one token. Hence, micro commits, including their finest-grained form, the one-token commit, represent a non-negligible development

activity. Thus, **we need to understand how to better support developers, first, by deeply looking at the need for micro commits, and second, by reducing the amount of effort needed to complete these changes.**

7.3 Program repair

The results of RQ1 showed that micro commits frequently modify a single token, and its token type is name, literal, or operator. **Studying micro commits could help understand how software is modified with such a tiny amount of change, and provide datasets that improve methods that attempt to modify software automatically.** For example, datasets based on micro commits might improve data-driven program repair approaches that have been studied so far (Jiang et al. 2018; Liu et al. 2018; Martinez and Monperrus 2015). One potential idea involves utilizing our observations of frequently modified token types and tokens in Java and C. Our observations indicate that while the types of frequently modified tokens are similar, the actual tokens differ across languages. This information is important for developing a program repair approach. When dealing with multiple languages, focusing on token types is crucial. However, when focusing on a specific language, actual tokens can also be beneficial.

Also, **the empirical investigation of micro commits would reveal types of micro commits that are difficult to be generated by program repair approaches.** Listing 10 shows a *replace constant* micro commit example. Specifically, this commit changes a string literal token: “start” into “end”. Such a change might be difficult to be generated automatically, because it is not obvious why a literal token should be replaced by another one; however, other changes (including micro commits) might have performed this specific replacement somewhere else.

Some large commits might actually be composed of several micro commits (i. e., *tangled commits* (Dias et al. 2015; Herzig and Zeller 2013; Kirinuki et al. 2014)). Thus **it is worth also exploring the possibility of untangling micro commits from larger commits.** These untangled micro commits might be very valuable for program repair.

Finally, we present initial analysis results for micro commits regarding their maintenance activities. As stated in Section 1, we hypothesize that small changes are likely intended for maintenance purposes. Therefore, we deduce that exploring micro commits could be beneficial for program repair. To validate this hypothesis, we identify micro commits that fall under the corrective maintenance category as defined by Swanson (1976). Corrective maintenance is performed in response to failures. If corrective maintenance makes up a large proportion of micro commits compared to non-micro commits, it would confirm our hypothesis.

To identify the corrective commits, we followed the methodologies used in prior studies (Mockus and Votta 2000; Purushothaman and Perry 2005), which use keywords in commit messages. More specifically, if at least one of the keywords is included in the commit messages, we classify the commit into the corrective maintenance category. Otherwise, we do not label commits. We used the keyword list defined by Levin and Yehudai (2016) as follows: “fix”, “esolv”, “clos”, “handl”, “issue”, “defect”, “bug”, “problem”, “ticket”.

The detailed procedure is as follows

Step 1: Apply preprocessing to the commit messages using the NLTK package¹¹ in Python by following the steps below:

¹¹ <https://www.nltk.org/>

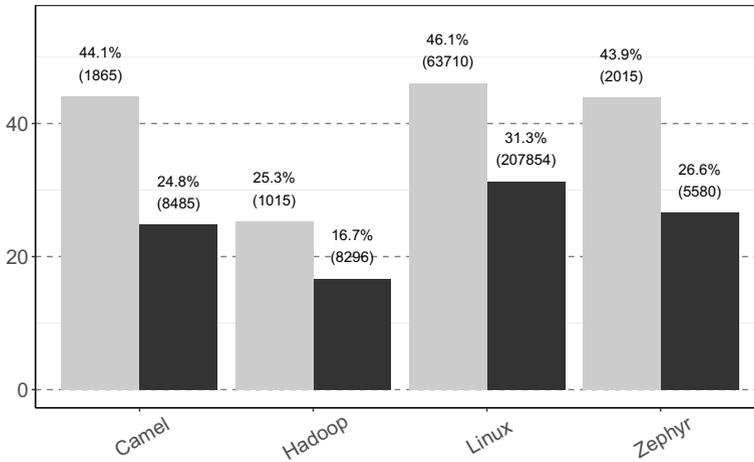


Fig. 6 The proportion of commits for the corrective category. Light gray indicates the proportion in micro commits; dark gray indicates the proportion in non-micro commits.

- Tokenize the text and convert all words to lowercase.
- Remove stopwords and punctuation.
- Perform stemming on all words.

Step 2: Check if the stemmed commit message contains a keyword.

Step 3: Identify commits that fall under the corrective category.

Micro commits are more likely to be failure-fixing activity than other commits.

Figure 6 shows the proportion of corrective micro and non-micro commits. In this figure, we compare the tendency of micro commits (light gray) and non-micro commits (dark gray). Corrective micro commits are larger than non-micro commits. Hence, micro commits distinguishably correspond to the corrective commits. This result shows that micro commits are usually applied to the source code to fix failures.

Also, this finding confirms our initial assumption that micro commits are used more frequently for maintenance purposes than non-micro commits. Interestingly, Hattori and Lanza (2008) found similar results, noting that tiny commits are often associated with corrective activities.

It should be noted that the keyword-based approach generally lacks accuracy (AlOmar et al. 2021; Kondo et al. 2022). To verify the accuracy of the identification, we manually inspect 20 micro commits and 20 non-micro commits identified as corrective, classifying them into three failure types within the corrective category defined by Swanson (1976). If we cannot associate any failure types, those would be considered false positive corrective commits. This allows us to estimate the actual number of corrective micro commits and corrective non-micro commits in the identified commits. We do not examine non-corrective commits to determine the proportion of false-negative corrective commits. This discussion only reports the minimum percentage of corrective commits. Our manual inspection revealed that there were no false-positive corrective micro commits. In contrast, we found 8 out of 20 false positive corrective commits in non-micro commits. This finding suggests that the percentage of corrective micro commits may not change significantly, while the proportion of

corrective non-micro commits could decrease. Therefore, our conclusion remains unchanged. In this manual inspection, we inspect only 20 micro and non-micro commits. Future studies could improve the validity of our findings. Our inspection is available in our sheet in our replication package.¹²

7.4 Size-perspective vs. semantic-perspective for defining micro commits

In this paper, we define micro commits through size metrics (i. e., the number of tokens). This is because we would like to assist with software engineering research, such as program repair. However, micro is a general term, and micro commits can be defined not only by size but also by semantic aspects. For instance, tangled commits (Dias et al. 2015; Herzig and Zeller 2013; Kirinuki et al. 2014) can be considered non-micro, whereas non-tangled commits can be categorized as micro. Additionally, defect-fixing commits can be categorized as micro or non-micro depending on the difficulty of the bug being fixed. We could explore these aspects using non-source code resources, such as source code comments, issue reports, and mailing lists. Exploring these semantic-based micro commits can also contribute to software engineering research.

8 Threats to validity

8.1 External validity

We conducted our empirical study on four OSS projects. To mitigate the threats to generalizability, we selected OSS projects that are active, popular, and well-known OSS projects written in two popular programming languages. However, even if we use these OSS projects, our results may not be generalized to all projects. Indeed, these are system software. To remedy this challenge, replication studies in research or practical scenarios (e.g., actual projects in the industry) are necessary. Hence, we provide a replication package.¹³ Also, the key tool `cregit` is an OSS tool; thus, researchers and practitioners easily convert their Git repositories into token-based ones.

8.2 Construct validity

We define micro commits based on the number of changed tokens. However, micro commits are a general term, and we can make different definitions. The key characteristic of micro commits is that such commits change a small code fragment. Our analysis (RQ1 and 2) shows that our definition is consistent with this characteristic. Hence, we believe our definition can be acceptable. However, our definition may not be the best; thus, future studies are necessary to find a better definition than our first one. For example, future studies can investigate different thresholds for the number of tokens. Also, studying different thresholds for added and deleted tokens (e. g., 3 added tokens and 5 deleted tokens) can be beneficial. Additionally, they can consider changes to source code comments. This definition would encompass not only maintenance activities related to code logic, but also various other maintenance activities.

¹² <https://doi.org/10.5281/zenodo.10963270>

¹³ <https://doi.org/10.5281/zenodo.10963270>

For future studies, researchers can use ASTs to tokenize the source code instead of `cregit`, which we used in this paper. While ASTs are powerful in analyzing token-level information, `cregit` is designed for Git, a de facto standard version control system. Hence, researchers easily analyze the software development history to support its process when using `cregit` instead of ASTs. Hence, we recommend using `cregit` in future studies.

There are several factors that can influence commits. For instance, the way developers write commits can vary depending on the developer and the project. To reduce these influences, we chose and analyzed four projects that involve a large number of developers. As a result, we expect the impact of such influences has been minimized.

8.3 Internal validity

To remove comment lines from source code files, we use regular expressions. This process is not perfect and may overlook comment lines. However, our manual analysis in RQ2 observes that our regular expressions usually work well because we do not find false-positive micro commits. Therefore, we can reduce the risks associated with using regular expressions.

Also, another threat exists in our manual analysis (RQ2). In this analysis, we performed manual labeling to micro commits according to our coding guide. Because this process is performed manually by the first author, the result may have false-positive and false-negative results. Therefore, we have made all labels publicly available to facilitate the validation of future studies. Also, to construct the coding guide, the first three authors independently inspected 20 micro commits three times. This process may also include errors. However, our agreement rate achieved substantial agreement in two consecutive iterations. Hence, we believe the coding guide is reliable. An alternative solution is to use an automatic classification approach rather than manual analysis. We developed a heuristic-based method to classify micro commits into their corresponding targets automatically. Overall, this approach achieved an accuracy of approximately 81.2% for categorizing micro commits into their targets. However, classifying micro commits corresponding to multiple targets can be challenging, with an accuracy of approximately 10.5%. To facilitate replication of this approach, we have included it in our replication package. Finally, we randomly sampled 400 micro commits from all projects. Therefore, our sampled micro commits may be biased by the size of the original projects. To mitigate this risk, we manually inspected additional micro commits from each project. We do not observe significant differences across the projects.

In the discussion, we use keywords to identify the commits related to the corrective maintenance activity as defined by Swanson (1976). While the keyword identification is widely used to categorize commits (Purushothaman and Perry 2005; Mockus and Votta 2000; Levin and Yehudai 2016, 2017; Hindle et al. 2008; Karampatsis and Sutton 2020), it is not perfect (AlOmar et al. 2021; Kondo et al. 2022). To mitigate this threat, we manually review identified commits and estimate their accuracy. Also, there are other sets of maintenance activities that can be used to classify commits, such as the IEEE standard (IEEE standard 2006). While we believe the maintenance activities defined by Swanson are acceptable, future studies are necessary to use other sets. Also, if commit messages do not contain any keywords, we exclude those commits from the analysis. However, it is possible that these commits are related to maintenance activities. Using more precise methods would enhance the validity of this analysis.

The tool “`cregit`” used to tokenize the source code files utilizes `srcML`. Therefore, our analysis can only be applied to specific versions of Java (Java SE8 Edition) and C (up to C11)

that are supported by srcML. We can find the supported versions on the official homepage.¹⁴ To extend our analysis to different versions of Java and C, it is necessary to update srcML and apply our analysis to those versions.

9 Related work

9.1 Challenges of mining Git repositories

Prior studies (German et al. 2019; Spacco and Williams 2009; Asaduzzaman et al. 2013; Bird et al. 2009) investigated and intended to address the challenges of mining Git repositories. For example, as described in Section 2, some non-functional changes update the information for each line and make it difficult to track code changes accurately. `cregit` (German et al. 2019) is proposed to address this problem by improving the blame feature in Git. More specifically, `cregit` tokenizes each line and applies the blame feature to the tokenized files. Spacco and Williams (2009) proposed a technique SDiff to track changes at the statement level instead of the line level. This technique combines previous line- and structural-based approaches. Specifically, SDiff tokenizes each statement and uses diff between revisions. These techniques tokenize the source code to address this problem. Similarly, we define micro commits based on changed tokens in this paper to track code changes accurately.

9.2 Change classification

Classifying changes (e. g., commits) into a certain category is a research topic in mining software repositories so far (Purushothaman and Perry 2005; Mockus and Votta 2000; Levin and Yehudai 2016, 2017; Hindle et al. 2008; Karampatsis and Sutton 2020; Ghadhab et al. 2021; Hindle et al. 2009; Alali et al. 2008; Williams and Carver 2010; German 2006; Mauczka et al. 2015; Yan et al. 2016). For example, many prior studies intend to classify changes in terms of the purpose (Mockus and Votta 2000; Levin and Yehudai 2016, 2017; Hindle et al. 2008; Karampatsis and Sutton 2020; Ghadhab et al. 2021; Hindle et al. 2009; German 2006; Mauczka et al. 2015; Yan et al. 2016). Levin and Yehudai (2016) classified commits into the maintenance activities defined by Swanson (1976). Hindle et al. (2009) used machine learning classifiers to classify changes into the extended Swanson categories. Ghadhab et al. (2021) used a pre-trained deep learning model known as BERT to classify commits into maintenance categories.

On the contrary, in this paper, we classify commits into micro commits based on their size and empirically investigate their characteristics, and there are several similar prior studies (Purushothaman and Perry 2005; Hindle et al. 2008; Alali et al. 2008). Purushothaman and Perry (2005) classified changes into three categories and studied them: one-line changes, small changes, and all. Specifically, this study used the number of changed lines for this classification. Hindle et al. (2008) identified large commits based on the number of changed files and revealed the characteristics of large commits. They also compared their result with the characteristics of the small commits by Purushothaman and Perry (2005). Alali et al. (2008) empirically investigated the characteristics of commits in nine OSS projects. They used three size criteria: the number of files, lines, and hunks. For example, they found that

¹⁴ <https://www.srcml.org/#home>

approximately 19.9% of commits in the GNU gcc system change at most five lines. However, the finest-grained changed source code entity is a line in these papers, and such an entity loses the information of changed tokens in a line. This limitation makes it difficult to define a certain category of commits based on finer-grained source code changes. Hence, our investigation would provide a new research direction in which researchers and practitioners use token-level changes.

9.3 Knowledge gap in previous studies

Compared to these prior studies, this research is the first to define micro commits at a fine granularity, specifically at the token level, through empirical analysis. Small commits defined at the line level, which previous studies often used, may overlook important information for improving existing software engineering research. Our research addresses this knowledge gap by conducting the analysis at the token level.

For instance, as explained in Section 7.3, the findings of RQ1 in this study have implications for research in program repair. These findings indicate the need to explore approaches for fixing bugs caused by a single name or literal token. This is because existing automated program repair approaches (Le Goues et al. 2012; Liu et al. 2019; Jiang et al. 2018) may not be effective in such scenarios due to a lack of information to repair the code. These findings and implications were obtained because the analysis was conducted at the token level. It would have been difficult to obtain such findings and implications using a line-level analysis. The novelty of this study lies in conducting the analysis at the token level and providing these implications. The details of our findings and implications can be found in Sections 5, 6 and 7.

10 Conclusion

In this paper, we defined micro commits (add at most five tokens and remove at most five tokens) and investigated their characteristics. This research is the first to define micro commits at a fine granularity, specifically at the token level. The key novelty of this study lies in conducting the analysis at the token level and providing implications for software engineering research.

Below, we present a summary of the findings from our empirical analysis:

- Our defined micro commits account for between 7.45–17.95% of all studied commits. Approximately 1 in 3 or 4 these changes (2.39–4.88% of all studied commits) involve replacing one token with another. Furthermore, RQ3 demonstrates that approximately 90% of one-line commits are micro commits, but only approximately 40–50% of micro commits are one-line commits. In fact, approximately 30–40% of micro commits include two or more hunks.
- The results of RQ1 show that micro commits primarily affect name token types (37.7–44.5%), literal token types (9.2–34.9%), or operator token types (6.6–10.4%). The most frequently affected tokens vary: the period in Java (2.5% in Camel and 3.9% in Hadoop) and the 0/1 in C (1.8 and 1.0% in Linux and 1.3 and 0.9% in Zephyr). Furthermore, the most frequently observed pattern is the modification of a single token. In Java projects, this modification is typically a single literal token. On the other hand, in C projects, the modification is usually a single name token.

- The results of RQ2 indicate that approximately 86% of micro commits involve a single operation on a single target, with the main focus being the replacement of existing targets. The multi-operation micro commits primarily involve changing the order of statements (19.3%).

In the discussion, we presented the following four implications of micro commits on future research based on the findings:

- Based on RQ3, it is observed that almost all one-line commits are micro commits, whereas only 40–50% of the micro commits are one-line commits. Therefore, token-based complexity metrics offer supplementary information to the commonly used line-based complexity metrics. Designing metrics to measure token-based complexity is a potential area for future research.
- Based on the statistics of micro commits, they account for a non-negligible proportion of all studied commits (7.45–17.95%). Additionally, according to Section 7.3, these commits are more likely used to fix bugs. Therefore, supporting the development of micro commits is an important area for future research.
- Based on RQ1, micro commits frequently modify a single token, with the token type often being either a name or a literal. While these micro commits often address bug fixes, suggesting patches to fix individual name or literal tokens can be challenging with existing program repair approaches. Therefore, it is necessary to investigate these micro commits and propose new program repair approaches for future research.
- We define micro commits based on size metrics. However, micro is a general term, and micro commits can be defined not only by size but also by semantic aspects (e. g., tangled commits or not). Exploring semantic-based micro commits is a potential area for future research.

The key message of this paper is as follows:

The token-level definition could help researchers and practitioners to improve software engineering approaches for software quality assurance activities.

Appendix

A Example micro commits

Listing 11 Example “add statement” commit diff retrieved from 122503683169b21d9cdb90380a20caad7ba994b6 in Linux.

```
@@ -240,0 +241 @@ static int check_partial_mapping(struct drm_i915_gem_object *obj,  
+         cond_resched());
```

Listing 12 Example “replace constant” commit diff retrieved from b7a90e8043e7ab1922126e1c1c5c004b470f9e2a in Linux.

```
@@ -174 +174 @@ static int hfsplus_sync_fs(struct super_block *sb, int wait)  
-     dprint(DBG_SUPER, "hfsplus_write_super\n");  
+     dprint(DBG_SUPER, "hfsplus_sync_fs\n");
```

Listing 13 Example “remove statement” commit diff retrieved from b95b4e1ed92a203f4bdfc55f53d6e9c2773e3b6d in Linux.

```
@@-2973 +2972,0 @@ void * __meminit alloc_pages_exact_nid(int nid, size_t size, gfp_t gfp_mask)
-EXPORT_SYMBOL(alloc_pages_exact_nid);
```

Listing 14 Example “no” commit diff retrieved from a092532483e3200a53c8b1170b3988cc668c07ef in Linux.

```
@@-1035 +1035 @@ static int __dwc3_gadget_ep_queue(struct dwc3_ep *dep, struct dwc3_request *req)
-    };
+    }
```

Listing 15 Example “change declaration” commit diff retrieved from 36f062042b0fd9f8e41b97a472f52139886ca26f in Linux.

```
@@-382 +382 @@ static ssize_t read_vmcore(struct file *file, char __user *buffer,
-static int mmap_vmcore_fault(struct vm_fault *vmf)
+static vm_fault_t mmap_vmcore_fault(struct vm_fault *vmf)
```

Listing 16 Example “change constant” commit diff retrieved from 1db76c14d215c8b26024dd532de3dcaf66ea30f7 in Linux.

```
@@-185 +185 @@ struct mthca_cmd_context {
-static int fw_cmd_doorbell = 1;
+static int fw_cmd_doorbell = 0;
```

Listing 17 Example “change identifier” commit diff retrieved from 70e8b40176c75d3544024e7c934720b11a8a11bf in Linux.

```
@@-564 +564 @@ static irqreturn_t pcie_isr(int irq, void *dev_id)
-    return IRQ_HANDLED;
+    return IRQ_NONE;
```

Listing 18 Example “change control flow” commit diff retrieved from 415a1975923722f729211a9efca550c60c519bf3 in Linux.

```
@@-530,0 +531 @@ static int ir_probe(struct i2c_adapter *adap)
+    break;
```

Listing 19 Example “change expression” commit diff retrieved from 40cc394be1aa18848b8757e03bd8ed23281f572e in Linux.

```
@@-1922 +1922 @@ static int nfs_parse_devname(const char *dev_name,
-    *comma = 0;
+    len = comma - dev_name;
```

Acknowledgements We gratefully acknowledge the financial support of: (1) JSPS for the KAKENHI grants (JP21H04877, JP22K18630, JP22K17874, JP24K02921), Bilateral Program grant JPJSBP120239929; (2) the Inamori Research Institute for Science for supporting Yasutaka Kamei via the InaRIS Fellowship; and (3) NSERC Canada.

Funding All authors certify that they have no affiliations in the subject matter or materials discussed in this manuscript.

Data Availability Statements The replication package that supports the findings of this study is available. (<https://doi.org/10.5281/zenodo.10963270>)

Declarations

Conflicts of interests/Competing interests All authors certify that they have no involvement in any organization or entity with any financial interest or non-financial interest

References

- Iso/iec/ieee international standard for software engineering - software life cycle processes - maintenance. ISO/IEC 14764:2006 (E) IEEE Std 14764-2006 Revision of IEEE Std 1219-1998) pp. 1–58 (2006). <https://doi.org/10.1109/IEEESTD.2006.235774>
- Alali A, Kagdi H, Maletic JI (2008) What's a typical commit? a characterization of open source software repositories. In: Proceedings of the 16th IEEE International Conference on Program Comprehension (ICPC), IEEE, pp 182–191
- AlOmar EA, Mkaouer MW, Ouni A (2021) Toward the automatic classification of self-affirmed refactoring. *J Syst Softw* 171, 110,821
- Asaduzzaman M, Roy CK, Schneider KA, Penta MD (2013) Lhdiff: a language-independent hybrid approach for tracking source code lines. In: Proceedings of the IEEE International Conference on Software Maintenance (ICSM2013), IEEE Computer Society, pp 230–239
- Bird C, Rigby PC, Barr ET, Hamilton DJ, German DM, Devanbu P (2009) The promises and perils of mining git. In: Proceedings of the 6th IEEE International Working Conference on Mining Software Repositories (MSR), pp 1–10
- Chouchen M, Ouni A, Kula RG, Wang D, Thongtanunam P, Mkaouer MW, Matsumoto K (2021) Anti-patterns in modern code review: Symptoms and prevalence. In: Proceedings of the 28th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER2021), IEEE, pp 531–535
- Collard ML, Decker MJ, Maletic JI (2013) srcml: an infrastructure for the exploration, analysis, and manipulation of source code: a tool demonstration. In: Proceedings of the 2013 IEEE international conference on software maintenance, IEEE, pp 516–519
- Dias M, Bacchelli A, Gousios G, Cassou D, Ducasse S (2015) Untangling fine-grained code changes. In: Proceedings of the 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER), IEEE, pp 341–350
- Fleiss JL (1971) Measuring nominal scale agreement among many raters. *Psychol Bull* 76:378–382
- German DM (2006) An empirical study of fine-grained software modifications. *Empir Softw Eng* 11:369–393
- German DM, Adams B, Stewart K (2019) cregit: token-level blame information in git version control repositories. *Empir Softw Eng* 24(4):2725–2763
- Ghadhab L, Jenhani I, Mkaouer MW, Messaoud MB (2021) Augmenting commit classification by using fine-grained source code changes and a pre-trained deep neural language model. *Inform Softw Technol* 135
- Hata H, Novielli N, Baltes S, Kula RG, Treude C (2022) Github discussions: An exploratory study of early adoption. *Empir Softw Eng* 27:1–32
- Hata H, Treude C, Kula RG, Ishio T (2019) 9.6 million links in source code comments: purpose, evolution, and decay. In: Proceedings of the 41st International Conference on Software Engineering (ICSE2019), pp 1211–1221
- Hattori LP, Lanza M (2008) On the nature of commits. In: Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE), pp III–63–III–71

- Herzig K, Zeller A (2013) The impact of tangled code changes. In: Proceedings of the 10th Working Conference on Mining Software Repositories (MSR), IEEE, pp 121–130
- Hindle A, German DM, Godfrey MW, Holt RC (2009) Automatic classification of large changes into maintenance categories. In: Proceedings of the IEEE 17th International Conference on Program Comprehension (ICPC), IEEE, pp 30–39
- Hindle A, German DM, Holt R (2008) What do large commits tell us? a taxonomical study of large commits. In: Proceedings of the 2008 International Working Conference on Mining Software Repositories (MSR), pp 99–108
- Hoang T, Dam HK, Kamei Y, Lo D, Ubayashi N (2019) Deepjit: an end-to-end deep learning framework for just-in-time defect prediction. In: Proceedings of the 16th International Conference on Mining Software Repositories (MSR), IEEE, pp 34–45
- Jiang J, Xiong Y, Zhang H, Gao Q, Chen X (2018) Shaping program repair space with existing patches and similar code. In: Proceedings of the 27th ACM SIGSOFT international symposium on software testing and analysis, pp 298–309
- Kamei Y, Shihab E, Adams B, Hassan AE, Mockus A, Sinha A, Ubayashi N (2013) A large-scale empirical study of just-in-time quality assurance. *IEEE Trans Software Eng* 39(6):757–773
- Karampatsis RM, Sutton C (2020) How often do single-statement bugs occur? the manysstubs4j dataset. In: Proceedings of the 17th International Conference on Mining Software Repositories (MSR), pp 573–577
- Kirinuki H, Higo Y, Hotta K, Kusumoto S (2014) Hey! are you committing tangled changes? In: Proceedings of the 22nd International Conference on Program Comprehension (ICPC), pp 262–265
- Kondo M, German DM, Mizuno O, Choi EH (2020) The impact of context metrics on just-in-time defect prediction. *Empir Softw Eng* 25(1):890–939
- Kondo M, Kashiwa Y, Kamei Y, Mizuno O (2022) An empirical study of issue-link algorithms: which issue-link algorithms should we use? *Emp Softw Eng* 27
- Le Goues C, Nguyen T, Forrest S, Weimer W (2012) Genprog: a generic method for automatic software repair. *IEEE Trans Software Eng* 38(1):54–72
- Levin S, Yehudai A (2016) Using temporal and semantic developer-level information to predict maintenance activity profiles. In: Proceedings of the IEEE International Conference on Software Maintenance and Evolution (ICSME), IEEE, pp 463–467
- Levin S, Yehudai A (2017) Boosting automatic commit classification into maintenance activities by utilizing source code changes. In: Proceedings of the 13th international conference on predictive models and data analytics in software engineering, pp 97–106
- Liu K, Kim D, Koyuncu A, Li L, Bissyandé TF, Le Traon Y (2018) A closer look at real-world patches. In: Proceedings of the International Conference on Software Maintenance and Evolution (ICSME), pp 275–286
- Liu K, Koyuncu A, Kim D, Bissyandé TF (2019) Tbar: revisiting template-based automated program repair. In: Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, pp 31–42
- Martinez M, Monperrus M (2015) Mining software repair models for reasoning on the search space of automated program fixing. *Empir Softw Eng* 20:176–205
- Mauczka A, Brosch F, Schanes C, Grechenig T (2015) Dataset of developer-labeled commit messages. In: Proceedings of the 2015 IEEE/ACM 12th Working Conference on Mining Software Repositories (MSR), IEEE, pp 490–493
- McIntosh S, Kamei Y (2018) Are fix-inducing changes a moving target? A longitudinal case study of just-in-time defect prediction. *IEEE Trans Software Eng* 44(5):412–428
- Meng X, Miller BP, Williams WR, Bernat AR (2013) Mining software repositories for accurate authorship. In: Proceedings of the 2013 IEEE international conference on software maintenance, IEEE, pp 250–259
- Mockus A, Votta LG (2000) Identifying reasons for software changes using historic databases. In: Proceedings of the International Conference on Software Maintenance (ICSM), pp 120–130
- Nugroho YS, Hata H, Matsumoto K (2020) How different are different diff algorithms in git? use-histogram for code changes. *Empir Softw Eng* 25:790–823
- Purushothaman R, Perry DE (2005) Toward understanding the rhetoric of small source code changes. *IEEE Trans Software Eng* 31(6):511–526
- Servant F, Jones JA (2017) Fuzzy fine-grained code-history analysis. In: Proceedings of the IEEE/ACM 39th International Conference on Software Engineering (ICSE), IEEE, pp 746–757
- Spacco J, Williams CC (2009) Lightweight techniques for tracking unique program statements. In: Proceedings of the IEEE International Working Conference on Source Code Analysis and Manipulation (SCAM2009), IEEE Computer Society, pp 99–108
- Swanson EB (1976) The dimensions of maintenance. In: Proceedings of the 2nd International Conference on Software Engineering (ICSE), pp 492–497

- Viera AJ, Garrett JM (2005) Understanding interobserver agreement: the kappa statistic. *Fam Med* 37(5):360–363
- Wang D, Kondo M, Kamei Y, Kula RG, Ubayashi N (2023) When conversations turn into work: a taxonomy of converted discussions and issues in github. *Empir Softw Eng* 28(6):138
- Wang D, Xiao T, Thongtanunam P, Kula RG, Matsumoto K (2021) Understanding shared links and their intentions to meet information needs in modern code review. *Emp Softw Eng* 26
- Williams BJ, Carver JC (2010) Characterizing software architecture changes: a systematic review. *Inf Softw Technol* 52(1):31–51
- Yan M, Fu Y, Zhang X, Yang D, Xu L, Kymer JD (2016) Automatically classifying software changes via discriminative topic model: Supporting multi-category and cross-project. *J Syst Softw* 113:296–308
- Zanaty FE, Hirao T, McIntosh S, Ihara A, Matsumoto K (2018) An empirical study of design discussions in code review. In: *Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement*, pp 1–10

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.



Masanari Kondo currently works as an assistant professor in the Principles of Software engineering and programming Languages Lab. (POSL) at Kyushu University, Kyushu, Japan. He was a Ph.D. student in the Software Engineering Laboratory (SEL) at the Kyoto Institute of Technology, Kyoto, Japan. He was also a Young Scientist of JSPS Research Fellowships (DC1) (2019.4-2021.3) and a visiting researcher in the Software Analysis & Intelligence Lab (SAIL) at Queen's University, Kingston, Canada (2018.3-2019.3). His research interests include supporting software developers by providing tools and methods for software quality assurance activities based on software repository mining, machine learning, and statistical analysis techniques. He received his BSc, MSc, and Ph.D. degrees from the Kyoto Institute of Technology (2017, 2019, and 2021). More about Masanari can be read on his website: <https://mkmknd.github.io/>



Daniel M. German is professor of computer science at the University of Victoria, Canada. His research areas are open source software engineering, mining software repositories and the impact of copyright in software development. He received his PhD in computer science from the University of Waterloo, Canada.



Yasutaka Kamei is a professor at Kyushu University in Japan. He has been a research fellow of the JSPS (PD) from July 2009 to March 2010. From April 2010 to March 2011, he was a postdoctoral fellow at Queen's University in Canada. He received his B.E. degree in Informatics from Kansai University, and the M.E. degree and Ph.D. degree in Information Science from Nara Institute of Science and Technology. His research interests include empirical software engineering, open source software engineering and Mining Software Repositories (MSR). His work has been published at premier venues like ICSE, FSE, ESEM, MSR and ICSM, as well as in major journals like TSE, EMSE, and IST. He served as a program-committee co-chair of the 23rd IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER 2016) and the 15th International Conference on Mining Software Repositories (MSR 2018). More information about him is available online at <http://posl.ait.kyushu-u.ac.jp/~kamei/>



Naoyasu Ubayashi is a professor emeritus at Kyushu University and a professor at Waseda University. He led the POSL (Principles Of Software Engineering and Programming Languages) research group at Kyushu University. Before joining Kyushu University, he worked for Toshiba Corporation and Kyushu Institute of Technology. He received his Ph.D. from the University of Tokyo. He was the general/organizing chair of AOSD 2013, ASPEC 2018, and MODELS 2021. He is a member of ACM SIGSOFT, IEEE Computer Society, and Information Processing Society of Japan (IPSI). He is also a Fellow of the IPSJ.



Osamu Mizuno received M.E. and Ph.D. degrees from Osaka University in 1998 and 2001, respectively. Currently, he is a Professor at the Faculty of Information and Human Sciences, Kyoto Institute of Technology. His current research interests include software repository mining, fault-prone module prediction, software process improvement, and risk evaluation and prediction of software development. He is a member of IEEE.

Authors and Affiliations

Masanari Kondo¹  · Daniel M. German² · Yasutaka Kamei^{1,3} · Naoyasu Ubayashi⁴ · Osamu Mizuno⁵

✉ Masanari Kondo
kondo@ait.kyushu-u.ac.jp

Daniel M. German
dmg@uvic.ca

Yasutaka Kamei
kamei@ait.kyushu-u.ac.jp

Naoyasu Ubayashi
ubayashi@ait.kyushu-u.ac.jp

Osamu Mizuno
o-mizuno@kit.ac.jp

- ¹ Principles of Software engineering and programming Languages (POSL), Kyushu University, Fukuoka, Japan
- ² Department of Computer Science, University of Victoria, Victoria BC, Canada
- ³ Inamori Research Institute for Science, Kyoto, Japan
- ⁴ Waseda University, Tokyo, Japan
- ⁵ Software Engineering Laboratory (SEL), Kyoto Institute of Technology, Kyoto, Japan