

Why Developers Refactor Source Code: A Mining-based Study

JEVGENIJA PANTIUCHINA, Università della Svizzera italiana

IORELLA ZAMPETTI, University of Sannio

SIMONE SCALABRINO, University of Molise

VALENTINA PIANTADOSI, University of Molise

ROCCO OLIVETO, University of Molise

GABRIELE BAVOTA, Università della Svizzera italiana

MASSIMILIANO DI PENTA, University of Sannio

Refactoring aims at improving code non-functional attributes without modifying its external behavior. Previous studies investigated the motivations behind refactoring by surveying developers. With the aim of generalizing and complementing their findings, we present a large-scale study quantitatively and qualitatively investigating *why* developers perform refactoring in open source projects. First, we mine 287,813 refactoring operations performed in the history of 150 systems. Using this dataset, we investigate the interplay between refactoring operations and process (*e.g.*, previous changes/fixes) and product (*e.g.*, quality metrics) metrics. Then, we manually analyze 551 merged pull requests implementing refactoring operations, and classify the motivations behind the implemented refactorings (*e.g.*, removal of code duplication). Our results led to (i) quantitative evidence of the relationship existing between certain process/product metrics and refactoring operations; and (ii) a detailed taxonomy, generalizing and complementing the ones existing in the literature, of motivations pushing developers to refactor source code.

CCS Concepts: • **Software and its engineering** → **Maintaining software**.

ACM Reference Format:

Jevgenija Pantiuchina, Fiorella Zampetti, Simone Scalabrino, Valentina Piantadosi, Rocco Oliveto, Gabriele Bavota, and Massimiliano Di Penta. 2020. Why Developers Refactor Source Code: A Mining-based Study. *ACM Trans. Softw. Eng. Methodol.* 37, 4, Article 1 (August 2020), 32 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Software refactoring has been widely studied in the research community, with most of the works falling into three main research threads: (i) approaches aimed at identifying refactoring opportunities [84]; (ii) techniques to recommend refactoring solutions for a

Authors' addresses: Jevgenija Pantiuchina, jevgenija.pantiuchina@usi.ch, Università della Svizzera italiana, Lugano, Switzerland; Fiorella Zampetti, fiorellazampetti@gmail.com, University of Sannio, Benevento, Italy; Simone Scalabrino, simone.scalabrino@unimol.it, University of Molise, Pesche, Italy; Valentina Piantadosi, valentina.piantadosi@unimol.it, University of Molise, Pesche, Italy; Rocco Oliveto, rocco.oliveto@unimol.it, University of Molise, Pesche, Italy; Gabriele Bavota, gabriele.bavota@usi.ch, Università della Svizzera italiana, Lugano, Switzerland; Massimiliano Di Penta, dipenta@unisannio.it, University of Sannio, Benevento, Italy.

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 ACM 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.

© 2020 Association for Computing Machinery.

1049-331X/2020/8-ART1 \$15.00

<https://doi.org/10.1145/1122445.1122456>

given design flaw [37]; and (iii) empirical studies looking at software refactoring from many different perspectives [44, 81, 92, 105, 106]. The knowledge of motivations pushing developers to perform refactoring [92] can help in building recommender systems able to propose suitable solutions for that. For this reason, understanding when and why developers perform refactoring has been the goal of many previous studies [40, 61, 92, 105, 106].

Some of these studies tried to answer this question by looking at specific factors that might correlate with refactoring operations, such as code quality proxies (*i.e.*, quantitative measures providing indications about the internal quality of code components, such as quality metrics or code smells) [40, 105]. While valuable, these studies provide limited insights into the reasons behind the performed refactorings, since their analysis is mostly quantitative and limited to a small number of factors. Other studies opted for a more qualitative approach by interviewing developers [61, 106] to identify the major factors that motivate their refactorings. Although these studies have pioneered the investigation of the reasons pushing developers to refactor their code, as observed by Silva *et al.* [92], the previously mentioned surveys are general purpose, meaning that they do not ask developers to *justify* specific refactorings they performed, but rather study refactoring habits in general. To address this limitation, Silva *et al.* [92] interviewed developers who authored 222 refactoring-related commits to understand the reasons behind these specific operations.

Stemming from the studies discussed above and to generalize their findings [61, 92, 106], this paper describes large-scale mining study combining quantitative and qualitative analyses to investigate the motivations behind refactoring operations, by observing code and discussions rather than interviewing developers. From a quantitative point of view, we mine the change history of 150 Java repositories hosted on GitHub to extract 287,813 refactoring operations of 25 different types performed by developers through the RMiner tool [100]. Then, we analyze product- (*e.g.*, slopes indicating whether the quality of code components as assessed by quality metrics is decreasing over time) and process-related (*e.g.*, source code change- and fault-proneness) factors that contribute to trigger refactoring actions. As compared to previous work [40, 105], we consider a more comprehensive set of factors and, more importantly, analyze them in a single model rather than in isolation, showing which ones are related to refactoring operations. From a qualitative point of view, we use the same set of systems to manually analyze a statistically significant sample of 551 pull requests (PRs) in which (i) developers discuss refactoring **and** (ii) RMiner identifies at least one refactoring operation. Through a manual analysis, we identify the rationale of the refactoring change, and whether it is the main intent of the change or, rather, they are triggered by the code review process of the PR. As main contribution of this analysis, we defined an extensive taxonomy of 67 motivations pushing developers to implement refactoring operations. Our qualitative analysis complements and generalizes the findings in previous survey-based studies [61, 92, 106] by investigating the same research question with a completely different experimental design.

As compared to the most similar work (*i.e.*, Silva *et al.* [92]), the following notable differences can be highlighted for what concerns the study design and findings:

- *Study Design: surveying developers vs analyzing their activities.* While Silva *et al.* contacted the developers authoring the refactorings asking their motivations for the implemented changes, we manually inspect pull requests implementing refactorings by analyzing their discussion and related commits in order to create our taxonomy of motivations. Investigating the same research question with two different experimental designs can lead to additional insights, and helps in generalizing previous findings.

- *Study Design: complementing qualitative and quantitative analysis.* In our work, we analyze the motivations behind refactoring operations not only from a qualitative perspective (as done by Silva *et al.* [92]), but also by quantitatively studying the influence of product and process metrics on the triggering of refactoring operations. Also, to the best of our knowledge, our study is the first one analyzing these metrics in a single model rather than in isolation (as done in [40], for example).
- *Findings: complementing and generalizing Silva et al. [92].* As output of their study, Silva *et al.* defined a list of 44 motivations for 12 frequently applied refactoring operations. Our taxonomy, besides confirming 41 of their motivations, thus improving the generalizability of their findings, includes 26 additional ones that are not covered in the previous study.

Our quantitative analysis indicates that code readability and process-related factors correlate with the changes a commit containing refactoring operations has. As the main result of the qualitative analysis, we provide a comprehensive taxonomy of 67 categories of motivations leading developers to refactoring operations. We describe and exemplify each category, and discuss its implications in refactoring research and practice.

2 STUDY DESIGN

The *goal* of this study is to quantitatively and qualitatively analyze the context in which refactoring operations occur in open source projects, with the aim of identifying the circumstances that may make a refactoring happen. The *quality focus* relates not only to code quality, but also to the improvement of the software development process. The *context* consists of 287,813 refactoring actions automatically identified in 150 open-source projects and, for the qualitative analysis, of 551 manually-analyzed PRs mentioning refactoring operations and linked onto refactoring-related commits.

We address the following two research questions (RQs):

RQ₁: *Which product and process-related factors relate with an increase of refactoring operation chances?* We are interested in studying if various source code features or process features correlate with the presence of refactoring operations in a commit.

RQ₂: *What are the reasons for performing a refactoring operation?* We investigate the rationale behind refactoring opportunities. We consider refactorings occurred in PRs, and perform a qualitative analysis of developers' discussions over the PR. Also, since previous work found that most refactoring operations occur with other changes [81], by analyzing PRs we give a closer look at this phenomenon, investigating if the refactoring was tangled with other changes, and looking at whether the refactoring was the primary purpose of the PR. We decided to answer this RQ by looking at PRs rather than at commits implementing refactorings since PRs offer a richer set of information to analyze to derive the rationale behind refactoring operations. Indeed, they often feature a discussion among developers that can help in better understanding what the goal of the implemented change was.

The formulated RQs investigate the same phenomenon (*i.e.*, what the motivations for refactoring operations are) from two different perspectives (quantitative—RQ₁ vs qualitative—RQ₂). The catalog of motivations identified in the two RQs can complement and support each other.

2.1 Study Context

We identified the projects to be studied among repositories hosted on GitHub. Since the infrastructure used in our study (*e.g.*, the refactoring detection tool) only supports Java,

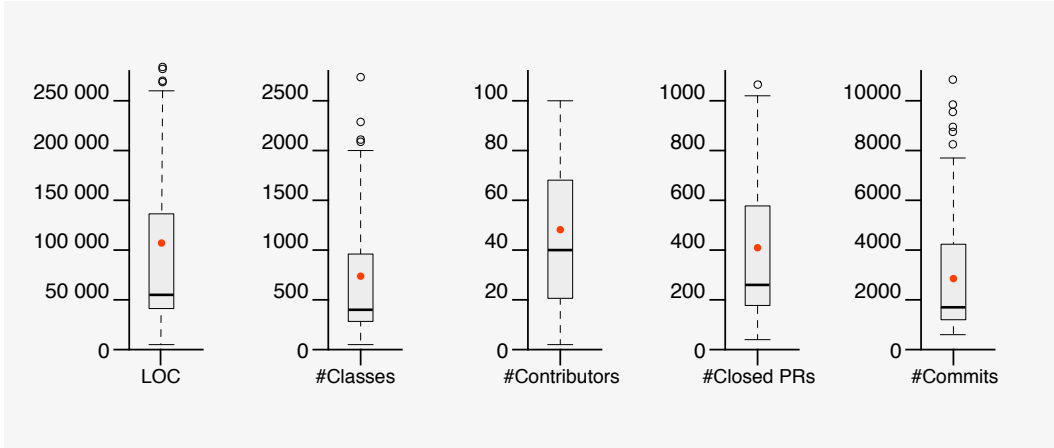


Fig. 1. Characteristics of the 150 projects used in our study

we focus on Java projects. Among all Java projects on GitHub, we aim at studying active projects having a non-trivial change history to study (needed to mine the PRs needed for our study) and not representing personal and/or toy projects (*e.g.*, a project created by a student during an assignment). To identify these projects we applied a number of selection criteria, only retaining projects having:

- *At least 5 contributors and 1 fork*, to exclude personal/toy projects.
- *At least 500 commits and 100 PRs*, to exclude projects having a short change history and unlikely to provide useful PRs for our study.
- *Modified at least once in the period Jan-May 2019*, to exclude inactive projects at the time in which this study has been run.

From the set of 303 remaining projects, we randomly selected 150 of them for our study (list available in [28]). The choice of selecting a subset of the 303 projects was dictated by the computationally expensive data extraction process adopted in our study. Indeed, as detailed in the following, besides detecting refactoring operations, we computed 42 product- and process-metrics (*e.g.*, code quality metrics, change-proneness of classes) for each of the 213,102 commits in the studied projects. This process took three months on a 56-core server. Figure 1 reports boxplots depicting the distribution of Lines of Code (LOC), number of classes (#Classes), number of contributors (#Contributors), number of closed PRs (#Closed PRs), and number of commits (#Commits) for the analyzed 150 systems. The raw data from which this figure has been created is available in our replication package [28].

We used the RMiner tool [100] to detect the refactorings implemented by developers in the studied projects. We focus on commits performed in the master/default branch of each project. We have chosen RMiner due to its high reported precision (98%) and recall (87%) [100]. RMiner takes as input two consecutive commits and provides as output the set of detected refactorings (see [100] for the supported refactorings).

2.2 Quantitative Analysis (RQ₁)

The occurrences of the detected refactorings constitute the *dependent variable* for RQ₁. As *independent variables*, we consider process-/product-related factors for each snapshot s_i (commit) of the master branch.

2.2.1 Identification of product and process metrics. The considered metrics are summarized in Tables 1, 2 and 3 and described in the following. The selection of these metrics (detailed in the following) is based on the will to include in our study:

- (1) *Metrics capturing code quality from different perspectives (Table 1).* We included both structural and semantic (*i.e.*, textual) metrics that have been shown to capture orthogonal code quality aspects [75]. Also, we considered the recent readability metrics proposed in the literature [42, 91], that have been shown to highly correlate with the developers' assessment of code readability.
- (2) *Code smells and quality issues widely studied in the literature (Table 2).* The presence of code smells has been correlated with higher change- and fault-proneness of code [83] and, thus, they could also be responsible for the triggering of refactoring actions. Also, static analysis tools are more and more used in the context of continuous integration to perform basic code quality checks at commit time. Thus, we decided to include the warnings raised by one of these state-of-the-art tools, *i.e.*, PMD [13].
- (3) *Process-related factors (Table 3).* These metrics are meant to provide a view on the development process, the developers involved in it, and historical information about the code components. We conjecture that these factors can play an important role in taking refactoring decisions, as also partially confirmed by previous work in the literature [105].

As detailed in Section 2.2.2, to avoid multicollinearity, we performed a variable selection. Near each metric, we indicate the cluster it belongs to and whether (✓) it was retained.

Source Code Quality Metrics. We consider, for each class C changed in each snapshot s_i , its quality *trend* as assessed by the 18 metrics in Table 1. These metrics capture different aspects of code quality, including size (*e.g.*, ELOC), coupling (*e.g.*, CBO), inheritance (*e.g.*, DIT), complexity (WMC), encapsulation (*e.g.*, NOPM), and readability (*e.g.*, StrRead). The first 13 metrics in Table 1 (*i.e.*, until NOSI included) have been computed by using the CK tool [4]. For the HsLCOM and C3, we used our implementation, while for the readability metrics we relied on the original implementations of the tools computing these metrics kindly made available by the original authors of the papers that introduced them [42, 91]. We start by measuring these 18 metrics on each class in each mined snapshot. Then, based on this information, we compute, for each snapshot, the slope of each metric over a window of N preceding commits (we set $N=10$ according to a previous work recommending just-in-time refactoring [85]). The slope of a line describes its steepness and in our case can highlight, for example, continuing degradation of some quality aspects (*e.g.*, a high positive slope for the WMC metrics indicates a steep increase in complexity for a class over time). Thus, using slopes we capture the improvement or degradation of quality factors, where the latter may trigger a refactoring. Clearly, slopes were considered unavailable for the first ten commits of a class.

Code Design Flaws and Quality Warnings. We consider code design flaws related to the lack of adoption of good Object-Oriented coding practices (*i.e.*, Spaghetti Code, Excessive Coupling), to complex/large code components (*i.e.*, Blob Class, Complex Class) as well as other design flaws and warnings (*i.e.*, Excessive Imports, Too Many Methods) raised by

Table 1. Quality Metrics (Product-related factors). Near each factor, we indicate whether (✓) it was retained. The factors retained in the model are also highlighted in boldface.

Metric	Description
CBO ✓	Coupling Between Object classes: measures the dependencies a class has [47]
WMC ✓	Weighted Methods per Class: sums the cyclomatic complexity of the methods in a class [47]
RFC ✓	Response For a Class: the number of methods in a class plus the number of remote methods that are called recursively through the entire call tree [47]
ELOC ✓	Effective Lines Of Code: the lines of code excluding blank lines and comments
NOM ✓	Number Of Methods in a class
NOPM ✓	Number Of Public Methods in a class
DIT ✓	Depth of Inheritance Tree: the length of the path from a class to its farthest ancestor [47]
NOC ✓	Number Of Children (direct subclasses) of a class
NOF ✓	Number Of Fields declared in a class
NOSF ✓	Number Of Static Fields declared in a class
NOPF	Number Of Public Fields declared in a class
NOSM ✓	Number Of Static Methods in a class
NOSI ✓	Number Of Static Invocations of a class
HsLCOM ✓	Henderson-Sellers revised Lack of Cohesion Of Methods (LCOM): a class cohesion metric based on the sharing of local instance variables by the methods of the class [47]. HsLCOM dresses limitations of the original LCOM [59]
C3 ✓	Conceptual Cohesion of Classes: avg. textual similarity between all pairs of methods in a class [75]
StrRead ✓	Structural readability: uses structural aspects (<i>e.g.</i> , line length) to model code readability [42]
ComRead ✓	Comprehensive readability model: combines structural, visual (<i>e.g.</i> , alignment) and textual features (<i>e.g.</i> , comments readability) [91]

a static analysis tool. We detect five types of code smells using an implementation of the DECOR smell detector based on the original rules defined by Moha *et al.* [78]. The choice of using DECOR is driven by the fact that (i) it is a state-of-the-art smell detector having high accuracy in detecting smells [78]; and (ii) it applies simple detection rules that allow it to be very efficient. The latter was a strict requirement for our analysis since we detected smells in all classes and for all studied systems' snapshots. In addition, we also consider 13 flaws from a widely-used static analysis tool that does not require code compilation, *i.e.*, PMD [13]. The set of detected design flaws and code quality warnings is described in Table 2.

Process-related factors. Besides the product-related factors previously described, we also study how process-related factors correlate with refactoring. In this case, we extract for each analyzed snapshot the factors summarized in Table 3.

Given a snapshot s_i , we compute its distance (in commits) from the previous and next release (first two rows in Table 3). This to verify the conjecture of Vassallo *et al.* [105] that refactoring does not occur immediately before/after a release. This information was retrieved using the GitHub API, through which it is possible to access all the tags related to a project. Then, we manually looked at the tags assigned to each project to isolate the ones referring to a new release.

We also consider the change- and fault-proneness of classes. The change-proneness is computed as the ratio between the total number of lines changed in the class C from the date of its addition to the project and the total number of commits in which C was changed, until each snapshot s_i .

Table 2. Code Design Flaws (Product-related factors). Near each factor, we indicate whether (✓) it was retained. The factors retained in the model are also highlighted in boldface.

Design Flaw	Description
DECOR Code Smells	
Blob Class	A large class that monopolizes most of the application logic [41]
Complex Class ✓	A class characterized by a high cyclomatic complexity [41]
Spaghetti Code ✓	A class declaring long methods without parameters [41]
CDSBP ✓	Class Data Should be Private: violation of information hiding principle [55]
Functional Decomposition ✓	Scarcely used object-oriented principles, such as inheritance and polymorphism; few methods and many private fields [41]
PMD code quality warnings	
Excessive Coupling ✓	A highly coupled class hindering reuse and maintainability [55]
Too Many Nested If Statements ✓	Makes the code harder to understand and increase error-proneness
Excessive Imports ✓	It might indicate too high coupling
Too High NPath Complexity	NPath is the number of acyclic execution paths throughout a method
Excessive Method Length ✓	It might indicate too many functionalities in a single method
Excessive Class Length	It might indicate too many responsibilities implemented in a class
Too Many Fields ✓	It can make the code hard to understand
Too Many Methods ✓	It might indicate too many responsibilities in a class
Cyclomatic Complexity	An excessive degree of decisional logic in a class
Excessive Parameter List ✓	It might indicate the need for a new object to wrap them
NCSS Type Count ✓	Similar to excessive class length, but it only considers actual statements
NCSS Method Count ✓	Similar to excessive method length, but it only considers actual statements
NCSS Constructor Count ✓	Equivalent of NCSS Method Count for constructors

The fault-proneness for C is computed as the number of bug-fixing commits it has been subject to in the past (*i.e.*, before s_i). For each project, we firstly identified all bug-fixing commits by matching patterns [53]: “fix” or “solve” or “close” **and** “bug” or “defect” or “crash” or “fail” or “error”. Then, for a given class C and for each snapshot s_i , we compute the number of bug-fixing commits preceding s_i and impacting C . Section 4 discusses the extent to which this simple heuristic for identifying bug fixes leads towards imprecisions.

Finally, we consider two metrics capturing the experience of the developers who worked on the system’s classes. The first metric, named *Developer Overall Experience*, assesses the experience of each developer as the number of commits she performed in the past. For each snapshot s_i and for each of its classes C , we extract the list of developers who modified C in the past (*i.e.*, before s_i). For each commit c_j (with $j < i$) in which C has been modified, we compute the experience of the developer authoring c_j (*i.e.*, the number of commits she performed *before* c_j). This gives us a distribution of developers’ experiences, for which we compute the minimum. Indeed, the minimum represents the lowest experience of a developer

Table 3. Process-related factors. Near each factor, we indicate whether (✓) it was retained. The factors retained in the model are also highlighted in boldface.

Metric	Description
Closeness to a previous release ✓	The number of commits until the previous minor/major release
Closeness to a next release ✓	The number of commits until the next minor/major release
Fault-Proneness ✓	Number of bugs fixed in the project history on a given class
Change-Proneness ✓	The average number of lines impacted in commits related to a class
Developer Overall Experience ✓	The number of past commits a developer performed
Developer Class Experience ✓	The number of past commits on a class performed by a developer

who worked on C , and we assume it might be correlated with future refactoring actions taken on C .

The *Developer Class Experience* computes a class-related experience: for each snapshot s_i and for each of its classes C , this form of experience is computed for a given developer as the number of commits impacting C she performed in the past. Thus, it is a more specific version of the overall experience. We compute this metric for each s_i and C under study in the same way explained for the overall experience.

2.2.2 Metrics Aggregation and Preprocessing. Since in RQ₁ we are interested to build an *explanatory model* explaining which factors correlate with the presence of refactoring actions in a snapshot, we had to aggregate metrics for all classes involved in each snapshot. For the product metrics, we compute the maximum slope among all classes involved in the snapshot, except for the conceptual cohesion (*C3*) and readability (*StrRead* and *ComRead*) metrics, which go in the opposite directions than other metrics (higher values are better). In such cases, we consider the minimum. In both cases, the rationale is to identify the “worst case” in a snapshot, which could ideally trigger a refactoring. As for the DECOR smells, we count the number of classes exhibiting a smell in each snapshot, while for PMD we sum the number of warnings of each type among changed classes. Similarly to what done for product metrics, for process metrics, we compute the maximum (*e.g.*, maximum number of bugs), except for the experience-related metrics, where we consider the minimum, again to consider the worst-case scenario. Finally, release-related metrics do not need to be aggregated, since they are already at commit granularity.

After that, to avoid multi-collinearity, we use the *R redun* function of the *Hmisc* package [58] for removing redundant variables. The *redun* function stepwise removes variable, starting from the most predictable one, until no variable can be predicted with an adjusted R^2 greater than a given threshold (0.8 in our study). Once again, we use the whole dataset to perform correlation analysis, because we intend to build an explanatory model and not a predictive model.

Since the value of our independent variables can depend on projects’ characteristics, and to properly interpret the importance of each variable in the model, we normalize variable values, within each project, in the interval $[0, 1]$. This is done by subtracting the minimum and dividing by the difference between the maximum and minimum. Finally, to build a

model easy to be interpreted, we invert (*i.e.*, compute 1-x) the values of variables going towards a different direction than the others (*i.e.*, those for which the higher the better).

2.2.3 Mixed-model Building. Once variables have been preprocessed, we address **RQ₁** by building mixed-effect generalized linear models. The model, built using the *glmer* function of the *lme4* [33] R package, is a logistic regression mixed-effect model where: (i) the dependent variable is a dichotomous variable indicating whether at least a refactoring was performed in a given commit; (ii) the independent variables (fixed effects) are all the aforementioned ones, after having pruned out those highly correlating with others; (iii) the random effect is the project in which the change occurred. The latter aims at controlling within-project effects, *e.g.*, a project following a specific development process had better code quality assurance policies than others. To simplify, our model reports whether the status of the system (as assessed by the used independent variables) in the snapshot S_{i-1} triggered a refactoring in the subsequent commit C_i .

To answer **RQ₁**, we report the details of the model, among others the coefficient of each factor in the model, and the p -value indicating whether the factor is statistically significant or not (for a significance level of 95%). We also report the odds ratio (OR) which, for a logistic regression model, is given by e^{c_i} where c_i is the coefficient of the i -th factor. An $OR > 1$ indicates that a unity increase of a variable increases of OR times the chances of a refactoring to occur.

2.3 Qualitative Analysis of Refactoring Discussions in Pull Requests (RQ₂)

For the qualitative analysis, we identified PRs likely discussing refactorings using two criteria to be satisfied: (i) whether a commit is a part of PR or made during its review contains a refactoring identified by RMiner, and (ii) whether the PR title or comments contain refactoring-related keywords. We used a list of refactoring keywords defined in a previous work [43] (available in our replication package [28]) and augmented it with all names of refactorings identified by RMiner [100]. Note that, while this selection process can generate false positives (*i.e.*, PRs unrelated to refactoring operations), these will be discarded during the manual analysis and, thus, do not represent a source of noise for our study.

Once the candidate set of 2,400 PRs has been identified, we created a randomly-stratified sample of 551 PRs. The strata here were represented by the projects, *i.e.*, PRs were sampled across projects proportionally based on the number of candidate PRs found in the previous step. The total number of PRs sampled allows us to ensure a significance interval (margin of error) of $\pm 5\%$ with a confidence level of 99%, and feature a total of 8,108 refactoring operations identified by RMiner. This estimation has been performed using a sample size (SS) calculation formula for an unknown population [89]:

$$SS = p \cdot (1 - p) \frac{Z_\alpha^2}{E^2}$$

and SS_{adj} for a known population pop :

$$SS_{adj} = \frac{SS}{1 - \frac{SS-1}{pop}}$$

where p is the estimated probability of the observation event to occur (we assume it being 0.5 if we don't know it a priori), Z_α is the value of the Z distribution for a given confidence level, and E is the estimated margin of error (5%).

We then uploaded the sample of PRs on a tagging webapp we used to perform a manual coding of PRs. The webapp presented to the annotator the following information: (i) the

PR title and hyperlink to the discussion; (ii) the refactoring-related keyword(s) matched in the PR text; and (iii) the list of refactorings detected by RMiner in commits linked to the PR, as well as the links to the GitHub diff pages of the commits themselves.

Through the coding app, each annotator could add one or more tagging items, containing the following information: (i) the type of refactoring action performed and discussed in the PR, or whether the change discussed was related to a combination of refactorings; (ii) whether the refactoring was the original intent of the PR, whether it happened as a consequence of the PR discussion, or whether it happened accidentally because of another change; (iii) whether the refactoring was tangled with other changes, or if it was the only purpose of the PR; (iv) finally, a tag indicating the motivation behind the refactoring, as it could be inferred from the inspection of the PR title/description, from its discussion, and from the commits related to it, looking at commit messages and, when needed, code diff. Note that each annotator could add more than one motivation for each PR (*e.g.*, one for each refactoring operation, or even more than one for the same refactoring). To assign the tag describing the motivation, the annotator could choose an available tag in a drop-down menu (from those previously created by other annotators or by herself), or add a new one if no tag was fitting the specific case. If an annotator realized that the PR discussion was not related to refactoring, the PR was tagged as “false positive”.

Six of the seven authors took part in the annotation process. The webapp we developed took care of automatically assigning each PRs to at least two of the involved annotators. We collected a total of 1,223 tags each one reporting a motivation for a refactoring (or combination of refactorings) performed in a PR. After each PR was tagged by two annotators, three of the authors jointly worked on the available tags to perform a card sorting activity [95] aimed at merging duplicates (*i.e.*, similar tags having the same meaning), and start grouping tags into categories. Then, they created a first taxonomy describing the different purposes of refactorings by only using the 699 tags for which there was no conflict (*i.e.*, the same tag was used by the two annotators for motivating the refactoring observed in a PR). After a first draft of the taxonomy was produced, two different authors refined it, by renaming some categories and moving sub-categories through the taxonomy. Once the final taxonomy was produced, three authors jointly discussed the conflicting cases in the categorization (524 out of 1,223 tags) and assigned them to suitable taxonomy categories, creating new ones when needed, and ensuring a consistency of category naming.

To address **RQ₂**, we report and discuss the taxonomy of refactoring motivations inferred as previously explained. In particular, we discuss the various categories, highlighting the percentages of PRs belonging to the category, reporting some examples, and highlighting the implications resulting from our empirical findings.

3 RESULTS

In the following we report and discuss the results addressing our RQs (Section 2).

3.1 Which product and process-related factors relate with an increase of refactoring operation chances?

Over the 213,102 snapshots analyzed, RMiner identified a total of 287,813 refactoring operations. More in details, our dataset contains 35,560 commits ($\simeq 17\%$) with at least one refactoring operation. If we exclude renaming operations (*Rename Method* and *Rename Class*), RMiner found a total of 209,385 refactorings in 28,716 different snapshots (14%).

Table 4 and Table 5 reports the results of the logistic regression mixed-effect model. More specifically, Table 4 reports the model diagnostics (Akaike Information Criterion — AIC

Table 4. Generalized mixed effect logistic regression model: diagnostics, residuals, and random effect

Diagnostics				
AIC	BIC	logLik	deviance	df resid.
21,071.5	21,362.4	-10,499.7	20,999.5	23,860
Scaled residuals				
Min	1Q	Median	3Q	Max
-2.0565	-0.5256	-0.3867	-0.1094	11.3848
Random effects				
Groups Name	Variance	Std.Dev.		
ProjectName (Intercept)	1.549	1.245		

[24], Bayesian Information Criterion — BIC, log likelihood, deviance, and degree of freedom residuals), the scaled residuals, and the random effect (project estimate). Concerning the model fitting (Table 4), we tried different models, namely logistic (*i.e.*, the one reported, AIC=21,071), linear (AIC=22,565), and Poisson (AIC=22,560). Also, although the analysis performed using the *redun* function already used a goodness-of-fit to iteratively remove variables, we experimented logistic models using structural metrics only (AIC=89,751), conceptual metrics only (AIC=89,475), code design flaws only (AIC=179,528), and process metrics only (AIC=23,583). Ultimately, the comprehensive logistic regression model we report is the one with the smallest AIC among those considered.

Table 5 reports the OR, estimate, standard error, *z*-value and *p*-value for the various factors we considered. We report in bold face the coefficient for which there is a statistically significant correlation. Metrics that have been inverted (*e.g.*, C3) are named with the prefix “Lack”.

Looking at code quality metrics, we found that the lack of structural readability plays a significant role: lack of structural readability [42] increases the odds of refactoring (OR=3.14). At the same time, *ComRead* readability metric and *LackC3* show a marginal significance (*p*-value = 0.02 and *p*-value = 0.04), respectively. In particular, looking at the *ComRead* readability metric combining structural and textual features [91] the OR is 2.68, while classes showing a decrease in their conceptual cohesion (*LackC3*) have 1.87 times higher odds of being refactored.

Among the structural metrics, we found that *NOM*, *NOPM* and *HsLCOM* have a statistically significant effect (marginally significant for *HsLCOM*), although the OR for *NOM* and *NOPM* is close to one. Instead, the OR for *HsLCOM* is 1.92, indicating that, as expected, a lack of cohesion increases the odds of inducing a refactoring operation.

In conclusion, from our analysis it results that conceptual and readability metrics play a more important role in the model than structural metrics. This finding is aligned with previous work aimed at applying conceptual metrics to suggest software refactoring [34] and modularization [39], and with findings of the seminal work about C3, indicating that such a metric is complementary to structural metrics [76].

None of the design flaws plays a statistically significant role. Although we expect that developers take care of removing smells, or try to “make static analysis tools happy”, and although previous work has pointed the role of refactoring for improving code having a poor quality, *e.g.*, overly complex code [40, 61, 92, 106], an evolutionary study on code smells

Table 5. Generalized mixed effect logistic regression model: effect of considered factors

	Metric	OR	Estimate	Std. error	z-value	p-value
	(Intercept)	0.00	-7.52	0.52	-14.50	< 0.01
Quality	LackStructRead	3.14	1.14	0.38	3.05	< 0.01
	LackComRead	2.68	0.99	0.41	2.42	0.02
	LackC3	1.87	0.63	0.30	2.06	0.04
	CBO	1.02	0.02	0.03	0.66	0.51
	WMC	0.98	-0.02	0.01	-1.57	0.12
	DIT	1.17	0.16	0.12	1.39	0.17
	NOC	0.95	-0.06	0.28	-0.20	0.84
	RFC	0.98	-0.02	0.02	-1.05	0.29
	NOM	0.88	-0.12	0.05	-2.54	0.01
	NOPM	1.22	0.20	0.06	3.19	< 0.01
	NOSM	0.91	-0.09	0.12	-0.76	0.45
	NOF	1.12	0.11	0.08	1.34	0.18
	NOSF	0.89	-0.11	0.13	-0.88	0.38
	NOSI	1.01	0.01	0.07	0.12	0.90
	LOC	1.00	0.00	0.00	0.91	0.36
	HsLCOM	1.94	0.66	0.29	2.28	0.02
Code Design Flaws	IsGodDecor	1.12	0.12	0.14	0.81	0.42
	IsCDSBPDecor	0.90	-0.11	0.15	-0.71	0.48
	IsComplexDecor	1.03	0.03	0.14	0.23	0.82
	IsFuncDecDecor	0.76	-0.28	0.25	-1.11	0.27
	IsSpaghCodeDecor	1.10	0.09	0.13	0.69	0.49
	AvoidDeeplyNestedIfStmts	0.92	-0.09	0.20	-0.43	0.67
	CouplingBtwObjects	0.73	-0.31	0.22	-1.42	0.16
	ExcessiveImport	1.04	0.04	0.20	0.18	0.86
	ExcessiveMethodLength	0.96	-0.04	0.23	-0.18	0.85
	ExcessiveParameterList	1.28	0.25	0.20	1.22	0.22
	TooManyFields	1.08	0.08	0.20	0.40	0.69
TooManyMethods	0.66	-0.42	0.29	-1.42	0.16	
Process	LackGeneralExp	1.26	0.23	0.10	2.30	0.02
	LackFileExp	8.93	2.19	0.13	16.45	< 0.01
	FilesRelatedToIssueFix	2.09	0.74	0.07	10.12	< 0.01
	AvgLinesImpactedInCommit	1.93	0.66	0.17	3.96	< 0.01
	DistancePreviousRelease	1.13	0.12	0.08	1.59	0.11
	DistanceNextRelease	1.43	0.36	0.09	4.12	< 0.01

indicate that smells mostly disappear when the source code is being rewritten, and only in less than 10% of the cases because of a refactoring action [102].

Interestingly, process-related metrics are highly representative if compared to product-related metrics: five of the six considered process-related metrics are statistically significant. Previous bug fixes play a role: a unit increase of the *FileRelatedToIssueFix* factor results in 2.09 higher odds of applying a refactoring in the system. Not only classes subject to

bug fixes are likely to be fault-prone in future [64, 80] but, since they are subject to (often quick-and-dirty) patches, they may necessitate refactoring actions. For related reasons, classes changing a lot (*AvgLinesImpactedInCommit*) also need to be refactored, although the OR is smaller (1.93).

Moving the attention to the metrics capturing the developers' experience, the *Developer Class Experience* (*LackFileExp*) has the highest OR. A unit increase of this factor related to the lack of specific experience (in terms of past commits) of developers that have recently modified a class, and therefore a decrease of experience increases the odds of refactoring by 8.93 times. In other words, changes applied by developers with little knowledge about a code component increase the need for restructuring it in the future. The general experience also plays statistically significant role, although the OR is relatively small (1.26).

Finally, looking at the proximity to a release (*DistanceNextRelease* and *DistancePreviousRelease*), the metrics indicate that refactoring operations are applied to the system far from a release of the system. More specifically, increasing the number of commits to a subsequent release, there are 1.43 higher odds of applying a refactoring. However, results are not significant while looking at the number of commits from a previous release (*i.e.*, p -value = 0.11). Our findings confirm previous literature [62, 105] since developers are aware that some kind of refactorings may result in the introduction of new faults [35] and in any case, refactoring represents a costly and risky operation [62]. For this reason, it is not very common to apply refactoring close to a new release of the software product. Furthermore, once released a new version of the software, developers likely tend to focus on bug-fixing activities instead of applying refactoring operations. In summary, based on our observations, refactorings are less likely to occur either immediately before major releases (developers focus on new features and, for what possible, on reliability of what they release), and immediately after (developers work on bug fixes). Instead, refactorings are more likely to happen in-between. Once again, note that this conclusion is based on purely observational data, and distance from the next release is unlikely to be used for prediction purposes (developers do not know when the next release is, unless a project or organization adopts very rigid release timelines).

We conclude RQ_1 stating that, on the one hand, by observing product metrics, only code readability plays a significant role. On the other hand, process-related metrics play a significant role. These are metrics related to previous changes and bug fixes, and to the experience of recent change authors.

3.2 What are the reasons for performing a refactoring operation?

Before digging into the results Table 6 reports statistics about the refactoring operations we found in the analyzed PRs. Note that: (i) several refactorings can be applied in one PR, therefore the number of refactorings is higher than that of PRs; (ii) we only list refactoring operations we observed at least 10 times. However, the overall number of refactorings (1,117) also includes the instances related to refactoring types we do not show in Table 6 (since having less than 9 occurrences). In most cases, developers do not discuss a specific refactoring operation. Instead, they rather provide a rationale for a combination of refactorings ($\sim 52\%$ of the cases). Then, *extract operation* ($\sim 10\%$) and renaming refactorings in general ($\sim 17\%$, in total) are the ones more discussed by developers. Surprisingly, we found that only a small percentage (5%) of refactorings were done collaterally, *i.e.*, without mentioning them at all. Instead, many of them were done as the original intent of the PR ($\sim 60\%$) or after discussing with other developers ($\sim 35\%$).

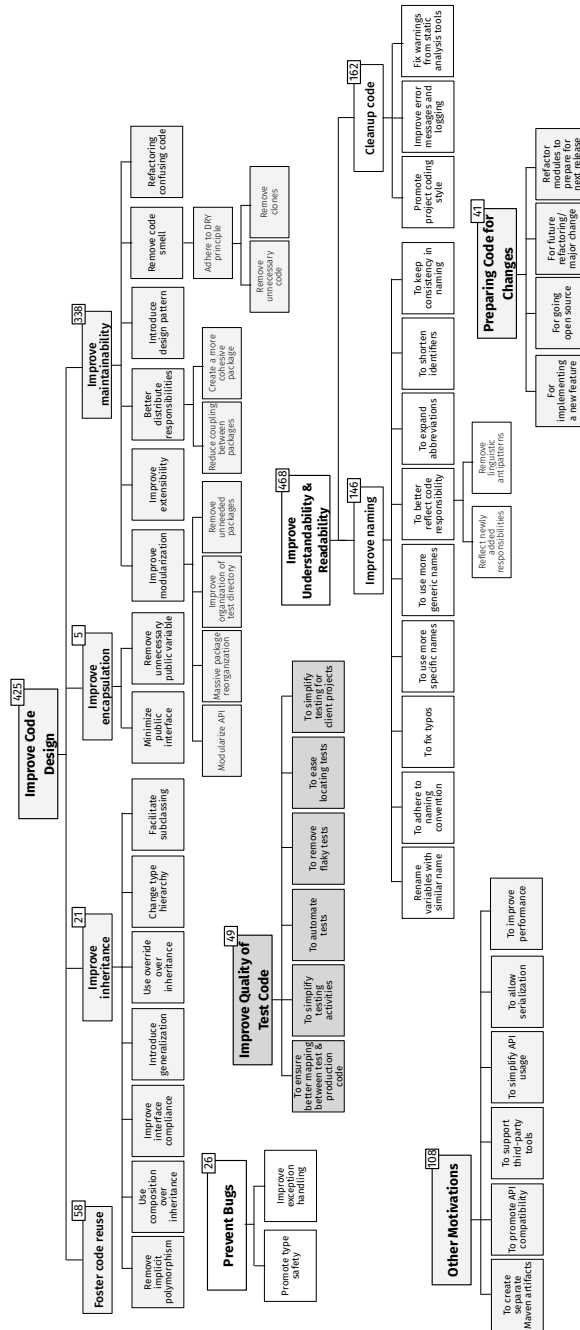


Fig. 2. Motivations behind refactoring operations

Table 6. Statistics of refactoring operations labeled in the 551 analyzed PRs. The “Freq.” column reports the number of times that the annotators defined a tag explaining the rationale behind each specific type of refactoring operation. The total number of 1,117 tags is the result of the 1,223 tags we defined, excluding the 94 “unclear” (*i.e.*, cases in which the annotators did not manage to identify the rationale for the refactoring) and 12 “false positives” (*i.e.*, PRs that were unrelated to refactoring).

Refactoring operation	Freq.	When?			Tangled	
		Orig. intent	Collateral	After discuss.	Yes	No
Combination of refactor.	578	65%	3%	32%	86%	14%
Extract operation	112	62%	3%	35%	62%	38%
Rename method	69	47%	4%	49%	59%	41%
Rename class	53	39%	6%	55%	56%	44%
Move class	39	64%	8%	28%	49%	51%
Extract variable	29	52%	35%	14%	72%	28%
Rename variable	29	35%	14%	52%	83%	17%
Extract interface	27	59%	11%	30%	59%	41%
Rename attribute	22	35%	8%	57%	44%	56%
Extract and move	22	68%	0%	32%	55%	45%
Rename parameter	19	69%	5%	26%	37%	63%
Move attribute	15	53%	14%	33%	73%	27%
Move operation	13	69%	0%	31%	39%	61%
Extract superclass	10	70%	10%	20%	80%	20%
Overall	1117	60%	5%	35%	73%	27%

Some refactoring operations, such as *extract variable* and *rename variable*, were performed collaterally more often, given their “local” nature: a variable name only matters in the methods in which it is declared, while a class name can possibly impact the whole system. It is worth noting that developers perform most of the renaming operations after they receive feedback from their peers. This shows that names are often discussed in PR reviewing activities. Finally, in line with Murphy-Hill *et al.* [81], we found that about a fourth of the refactorings are tangled with other changes.

Figure 2 depicts the taxonomy of refactoring motivations we have identified. It comprises six root categories: (i) *Improve Code Design* groups refactoring operations targeting an improvement of the system design, *e.g.*, by fostering the reusability of code; (ii) *Improve Understandability & Readability* includes refactorings aimed at reducing the effort to read and understand code, *e.g.*, by renaming identifiers; (iii) *Improve Quality of Test Code* groups all refactorings performed to improve the quality of the test code or ease the testing process; (iv) *Prevent Bugs* identifies refactorings performed to prevent the future introduction of bugs; (v) *Preparing Code for Changes* includes refactorings performed in preparation of other changes, *e.g.*, refactoring the code before implementing a new feature; (vi) finally, *Other Motivations* groups those motivations that cannot be classified into one of the previous categories.

It is important to point out that some of the categories of the taxonomy are not mutually exclusive. For example, a refactoring aimed at improving code readability, is also likely to improve maintainability. However, readability can be improved for multiple purposes (*e.g.*, simplify testing), and, for this reason, we separated these categories. We acknowledge

that other choices in terms of categories and assignment of instances to these categories are possible. Also, the hierarchical organization of the categories only indicates that child categories are specialization of their parent categories, while it does not imply that two categories at the same hierarchy level represent motivations at the same level of abstraction (e.g., *Improve Inheritance* and *Foster Code Reuse* appear at the same level, but the former is a more concrete motivation as compared to the latter).

Figure 2 also reports for each category of “motivations” the number of PRs in which we found related refactoring operations. For readability purposes, we only report these numbers for the main categories. Note that the number for a parent category does not correspond to the sum of the children, because some PRs were only assigned to the parent category, as the motivation was not specific enough. Also, the sum of refactoring instances in all root nodes does not correspond to the total number of 551 manually analyzed PRs because some PRs comprise refactorings falling into multiple categories, and we labeled some refactorings as *Unclear* (94) and discarded 12 PRs as *False Positive*.

We compared our taxonomy with the list of 44 motivations derived by Silva *et al.* [92] for 12 frequently applied refactoring operations (see Tables 3 and 4 in [92]). In particular, two of the authors tried to map Silva *et al.*'s motivations into our taxonomy, to see whether they were covered or not. Note that the mapping is not one-to-one since one motivation identified by Silva *et al.* [92] may be mapped to more than one category in our taxonomy, as well as one of our categories can group more than one of their motivations. This is expected since their motivations and the categories in our taxonomy have been derived by using two different methodologies. Indeed, while we have categorized the possible motivations behind the application of refactoring operations by looking at the discussions in PRs, Silva *et al.* [92] have asked the reasons behind specific instances of refactoring operations to the original developer who has applied it.

Only 3 out of the 44 motivations from Silva *et al.* [92] cannot be mapped in our taxonomy. The main reason is that for these three instances (i.e., *Enable recursion*, *Convert to top-level container*, and *Convert to inner class*) it was unclear to us the actual motivation behind the refactoring. For example, enabling recursion could be done to improve performance as well as to improve code readability. However, this high overlap between the two sets of motivations (i) validates and generalizes the work done by Silva *et al.* and (ii) supports the comprehensiveness of our taxonomy. As reported in Table 7, our taxonomy features 16 inner categories that are not covered in [92] (e.g., *To Ensure Better Mapping Between Test And Production Code*), 3 inner categories that are only partially covered (e.g., *Preparing Code for Changes*), and 6 inner categories (e.g., *Forster Code Reuse*) that are completely covered. We provide in our replication package [28] a spreadsheet reporting the mapping between the two taxonomies.

In the following, we discuss each root category, reporting interesting examples and outlining implications for researchers and practitioners (indicated with the ♡ icon), as well as highlighting the differences with the taxonomy provided by Silva *et al.* [92]. The complete list of manually analyzed PRs together with their refactorings/assigned tags is publicly available [28].

Improve Code Design (425 instances). Unsurprisingly, a large proportion of the analyzed refactorings are aimed at improving code design from several perspectives [56]. In 58 of these, the refactorings are aimed at making source code easier to be reused (see *Foster code reuse* in Figure 2). In 42% of these cases, this was accomplished through a combination of several refactoring operations, while in the remaining 58% specific refactorings were applied in isolation. When this happened, almost always (91%) an operation aimed at extracting a code

Table 7. Comparison with refactoring motivations found by Silva *et al.* [92]: ↑↑ highlights a perfect match (the motivation also emerges from their study); ↑ highlights a partial match (we found additional, specific motivations); ↓↓ stands for a mismatch (the motivation was not found in their study).

Root Category	Inner Category	Match
Improve Code Design	Foster code reuse	↑↑
	Improve inheritance	↑
	Improve encapsulation	↓↓
	Improve maintainability	↑↑
Improve Understandability & Readability	Improve naming	↑
	Cleanup code	↓↓
Improve Quality of Test Code	To ensure better mapping between test and production code	↓↓
	To simplify testing activities	↑↑
	To automate tests	↓↓
	To remove flaky tests	↓↓
	To ease locating tests	↓↓
	To simplify testing for client projects	↓↓
Preparing Code for Changes	For implementing a new feature	↑↑
	For going open source	↓↓
	For future refactoring/major change	↓↓
	Refactor modules to prepare for next release	↓↓
Prevent Bugs	Promote type safety	↓↓
	Improve exception handling	↓↓
Other Motivations	To create separate Maven artifacts	↓↓
	To promote API compatibility	↑↑
	To support third-party tools	↓↓
	To simplify API usage	↓↓
	To allow serialization	↓↓
	To improve performance	↑↑

component from an existing one was applied. In particular, *extract method* operations were performed in 70% of cases, to extract a small piece of functionality from an existing method thus avoiding code duplications and fostering the reuse of the extracted code. For example, during the code review of the PR #626 in the *nakadi* project [20], the reviewer observed that two of the implemented methods were “*almost the same except the very last line*” and suggested to “*extract a helper method*” in such a way to reduce code duplication and also allow other methods, in future, to reuse the same functionality. This was accomplished through an extract method refactoring. In other cases, the refactoring was more substantial and directly justified by the need for reusing specific pieces of functionality, as discussed in the PR #488 of the *dropwizard* project [6]. Here the contributor explains, when submitting the PR: “*I was looking at starting/stopping a Dropwizard app in Cucumber tests and DropwizardAppRule has all the functionality I need but obviously it doesn’t expose startIfRequired and stop methods. I’d happy to extract a DropWizardAppTestSupport class from DropwizardAppRule*”. After approval, this triggered an *extract class* refactoring. This last example is interesting for several reasons. 🕒 First, extract class is a non-trivial refactoring possibly having substantial ripple effects in the system, with the obvious possibility of introducing bugs. For example, the discussed commit impacted a total of 499 lines of code, thus showing that code reuse is a strong motivation for triggering refactoring operations. Second, while many approaches to identify extract class refactoring opportunities have been proposed (see *e.g.*, [36, 54]), they focus on the identification of complex classes implementing several responsibilities (*i.e.*, *God* or *Blob* classes [41]) that could be split into several classes. The class subject of the *extract class* refactoring (*i.e.*, *DropwizardAppRule*) is a fairly simple class composed by 156 effective LOC (excluding comments and blank lines) that is unlikely to be reported by refactoring

recommenders as a candidate for *extract class* refactoring. To cope with these cases, these recommenders could be combined with clone detectors [90] to factor out a class to be used by multiple other ones. Note that these “special” cases should complement the more standard refactoring recommendations done for complex and low-cohesive classes. Indeed, as shown in our RQ₁, classes characterized by a low cohesion as assessed by the C3 and HsLCOM metrics are more likely to be subject to refactoring operations.

Many (338) of the refactorings performed in the code design taxonomy aimed at *Improve Maintainability* (see Figure 2). In this category, refactorings aimed at improving the modularization were often implemented through simple move class refactorings, while we rarely observed massive package reorganizations (7 cases). ♡ This is in line with recommendations from previous literature, suggesting that approaches performing big-bang remodularization through clustering algorithms have limited applicability, and techniques suggesting fine-grained and incremental adjustments to software modularization should be preferred [57, 82]. Also, it was interesting why developers decided to perform remodularization. For example, in some cases move class refactorings are performed to group, in specific “API-related” packages, utility classes potentially useful in different parts of the system and/or to third-party components, *e.g.*, PR #324 from DSpace “*I suggest to move this class in dspace-api as it will be useful to port this feature to JSP UI as well*” [8]. While these changes might look suboptimal from the cohesion-coupling point of view (*i.e.*, they could generate a low-cohesive package) they are justified by a clear rationale. ♡ As also observed for the approaches automating *extract class* refactoring, tools recommending modularization solutions (see *e.g.*, [29, 66, 74, 87]) just strive to maximize the cohesion-coupling trade-off. Given the availability of historical data, they could also learn from previous changes what a meaningful modularization is from the developers’ perspective. While learning code changes is already an active research field [103], no previous work has attempted to design refactoring recommenders learning from developers’ activities what a meaningful refactoring is in a given context.

Removal of code clones is one of the two motivations behind the refactorings in the *Adhere to DRY principle* (Don’t Repeat Yourself) subcategory (child of *Remove Code Smell*), together with the removal of unnecessary code. Note that this category is strictly related to the *Foster code reuse* one. Indeed, some of the analyzed PRs fall into both these categories because factoring out duplicates also creates a more generic code element (*e.g.*, a class or method) that can be further reused. For example, PR #366 in the *fineract* project [2] can be seen as an example of improving reusability by adhering to the DRY principle, since it features an *extract method* refactoring suggested by the reviewer and avoiding code duplication while allowing the reuse of a piece of functionality now embedded in the extracted method. ♡ This confirms once more the relevance for practitioners of clone detectors [90] as well as of refactoring tools aimed at removing clones [101] and encourages their use in the Continuous Integration (CI) pipeline, as advocated by Duvall *et al.* [51].

Concerning the removal of unnecessary code, besides cases simply related to removing unused `imports`, we found refactorings performed to remove redundant code (*e.g.*, PR #1481 from `testng` [3]). ♡ While some work has investigated the automatic identification of redundant code in software systems [60, 71], the provided support is still very limited to specific redundancy cases (*e.g.*, those related to API usages [60]) or programming languages (*e.g.*, LISP [71]). The 55 cases related to refactorings motivated by the removal of unnecessary code suggest room for more research in this field.

Concerning the operations targeting a better distribution of the responsibilities across code components, one very interesting example comes from the DSpace project (PR #1083 [7]).

This PR implements a massive refactoring aimed at ensuring a better “separation of concerns/responsibilities” for an API module, and has been subject to votes by the community, because the refactored API was not backward compatible. Despite this issue, the merging has been approved thanks to the numerous positive advantages brought by the refactored API: “*makes it much easier to achieve future goals on our Roadmap, especially, moving us towards potentially better support of third-party modules*”, “*it cleans up one of the messiest areas of our existing API [...]*”. ♡ This case shows the non-trivial trade-offs that developers should consider in case of massive refactoring: for example, smartphones have limited battery life and they require software optimized to reduce the energy consumption. State-of-the-art refactoring recommenders [36, 99] ignore the heterogeneity of modern software, and the different priorities that non-functional requirements, possibly more important (*e.g.*, maintainability, performance, backward API compatibility) may have in different contexts. Future work should consider integrating into these recommender systems the possibility to define a priority list of non-functional properties that developers are or are not willing to sacrifice when applying refactoring. This would allow generating more meaningful and sensible refactoring recommendations.

The two sub-categories described above (*Foster code reuse* and *Improve maintainability*), *i.e.*, the ones highly represented in our study, have a complete matching with the motivations identified by Silva *et al.* [92]. This confirms their findings, and stresses once more the importance from the developers’ perspective of improving both the reusability and maintainability of code, especially when discussing whether to accept or not a PR.

Other less represented subcategories in the *Improve code design* taxonomy include refactorings aimed at improving the usage of inheritance (21 instances) and the ones working on the encapsulation (5). In these cases, considering the low number of instances belonging to each category, it is quite obvious that while comparing with the motivations by Silva *et al.* we found that these reasons did not emerge from their study. The only exception is the one related to inheritance that is only partially covered, as shown in Table 7.

Improve Understandability & Readability (468 instances). The majority of refactorings we found in the manually-analyzed PRs aim at improving understandability and readability of source code. This supports the findings of RQ₁, which indicate a significant correlation (and high OR) of readability metrics with refactoring operations. In this category, 146 refactorings were done to improve naming. The observed renamings had a variety of motivations (see Figure 2), ranging from fixing typos to keeping naming consistency throughout the project. Naming decisions were often carefully discussed, showing their importance for developers. An interesting example of discussion about naming is the PR #150 of the `optaplanner` project [11]. The original intent of the PR was to add a new feature, but the author explicitly asked for feedback about the naming of a new interface he extracted: “*we should discuss the naming and the usage of the SolverProblemBenchmarkResult interface*”. Such a name was changed after the discussion: “*renamed to BenchmarkResult as agreed in a meeting*”. In the same PR the developers also discussed several other names in the contributed code. For example, the PR author introduced a boolean field named `hasNonDefaultSubSingleCount`; another developer asked why such a value was introduced, since the name was not clear enough. The discussion triggered not only a renaming operation, but also a type change (from boolean to integer) to represent additional information that could be useful in future: “*maximumSubSingleCount is the best name here, as it gives us more potential information for the future at no cost*”.

We also found cases of renaming aimed at better reflecting the code responsibility (38). A representative example is in PR #251 of the `kafka-connect-elasticsearch` project [5], in

which one of the reviewers suggested to rename a test method to something very specific and clearly depicting the responsibility of the test case: “*you could change the name of the test method to something like `testCreateAndWriteToIndexForTopicWithUppercaseCharacters`. I like test names that read like the condition they are testing*”. Other cases aimed at removing linguistic antipatterns defined in the literature by Arnaoudova *et al.* [32] and known to have negative effects on the understandability of code [52]. ♡ This is only one of the studies linking the poor quality of identifiers to difficulties experienced by developers in code comprehension [49, 68–70, 98]. Our findings show that developers care about the quality of identifiers and carefully discuss their choice.

♡ Most of the rename refactoring recommendation approaches aim at fostering the usage of consistent naming [25, 72], while only a single attempt has been done, to the best of our knowledge, to recommend rename method refactorings with the goal of better reflecting the responsibilities implemented by the code [26]. Our results show that more effort in this direction is needed since this is the scenario in which developers more frequently perform rename refactorings. Also, ♡ the high number of rename refactorings implemented as a consequence of code review indicates the possibility to mine this data to evaluate automated rename refactoring techniques [25, 72]: the originally submitted identifier represents an opportunity for rename refactoring while the one adopted after the code review process can be used as reference of a good refactoring. This would avoid the evaluation of the rename refactoring techniques in artificial scenarios.

Looking at Table 7, and considering that developers can modify the names of packages, classes, variables or methods for different reasons, we can state that our *Improve Naming* category is only partially covered by the motivations reported in the previous study by Silva *et al.* [92]. For instance, while both studies identify the need for adhering to naming conventions, for keeping consistency in naming, or for better representing code responsibilities, in our taxonomy we also found cases where the renaming occurs to fix typos, shorten identifiers and expand abbreviations. ♡ The latter challenge (*i.e.*, expansion of abbreviations in code identifiers) has been vastly investigated in the software engineering research literature (see, for example, the works by Lawrie *et al.* [67]), with approaches proposed and empirically evaluated. However, to the best of our knowledge, there are no ready-to-use tools that, for example, can be integrated in a CI pipeline and can recommend to developers identifiers to expand at commit time. The implementation of such a tool is a clear next step to perform in this research field.

We also found several refactorings implemented to make the source code less confusing. Such changes involved improvements to both names and structural aspects. For example, we found an interesting example in PR #599 of the `htsjdk` project [15]. Note that the refactoring performed in this PR is an *extract interface* rather than a rename, that resulted in the interface `CRAMReferenceSource` implemented by the class `ReferenceSource`. The main goal of the refactoring was to improve code reusability: for this reason, we include such a case in our taxonomy under the *Foster code reuse* category. However, it is interesting to discuss this refactoring in the context of renaming: during the code review process, one of the reviewers argued that the chosen names were confusing because he expected an inheritance relationship in the opposite direction (*i.e.*, `CRAMReferenceSource` implements `ReferenceSource`) by reading the names alone. As a consequence, `ReferenceSource` was renamed to `CRAMReferenceSourceImpl`, making the relationship between the two classes more evident. ♡ This naming issue could be characterized as a sort of linguistic antipattern, and shows that the original catalog of these antipatterns defined by Arnaoudova *et al.* [32]

could be expanded by analyzing recommendations provided by reviewers in a code review process.

Finally, we found many cases in which the developers made more generic clean-ups in the code (162 instances), to improve the coding style and, in some cases, the quality of error messages and logging. Some of these changes were performed as a result of tools' recommendations. For example, the PR #346 of the `spring-amqp` project [18] fixes warnings raised by SonarQube. ♡ This suggests that developers are willing to fix issues identified by automatic tools. However, in our sample of PRs, SonarQube was the only tool mentioned in many discussions. Note that a specific analysis of the extent to which static analysis tool warnings are removed was not in scope of our work (but rather addressed in related literature [48, 63, 93]); this is the reason why, in RQ₁, we only considered two tools — DECOR and PMD — that could raise warnings that triggered refactoring operations.

More than 60% of the clean-up operations were done as part of the original intent of the PRs. Differently from other categories, we found very little discussion among developers regarding clean-ups. This is likely due to (i) the limited impact that these clean-ups generally have; and (ii) a general agreement on the need for improving code quality.

Prevent Bugs (26 instances). These refactorings are motivated by the will to prevent bugs, for example through a better exception handling or by promoting type safety. Note that these are changes that preserve the program's behavior. For this reason, we contemplated them in our taxonomy, even if they do not belong to the canonical cases of refactorings, such as those defined by Fowler [56]). This is why those categories are completely uncovered in the 44 motivations provided by Silva *et al.* [92]. Indeed, they studied the reasons behind specific refactoring operations that are detected by RMiner and inline with those defined by Fowler [56].

Some PRs are explicitly motivated by the will of improving the exception handling mechanism. This is the case for PR#933 of the `nakadi` project [21], in which the developer implements several different refactorings (*e.g.*, *rename class*, *move class*) to improve the overall handling of the exceptions in the project. Other PRs, instead, simplify the handling of exceptional conditions. For example, in PR#1067 from the `htsjdk` project [14], the developer fixes a possible issue caused by the invocation of the method `mFile.getSource()` within several exception messages. Indeed, in specific cases the `mFile` object could be `null`, leading to the throwing of a `NullPointerException`. For this reason, the developer implemented an *extract method* refactoring, creating the method `getSource()` which returns the value of `mFile.getSource()` when `mFile` is not `null`, and a constant string otherwise. This allowed to easily prevent `NullPointerException` by replacing the many usages of `mFile.getSource()` with an invocation to the newly created `getSource()` method. ♡ This is an interesting application of *extract method* refactoring, since it aims at refactoring a very small clone, *i.e.*, a method invocation reused, in the same way, in different parts of the code. Extract method is widely applied in the refactoring of clones [65]. However, the focus is usually on more complex clones sharing several statements, rather than on the identification of refactoring opportunities that, as in the discussed case, involve few code tokens but can have a positive impact on the reliability and maintainability of the system.

Another very interesting example is the PR#238 from the `minio-java` project [12]. Here the developer replaced general, unchecked exceptions such as `NullPointerException`, with more specific and checked ones. With “unchecked” we refer to those exceptions that in Java can be thrown without declaring them in the method signature. For example, in the specific case of PR#238, the method `getClient` of the class `Client` was throwing a `NullPointerException` in specific situations: note that this was an intended behavior of

the method, *i.e.*, there was an explicit `throw new NullPointerException()` in the code. However, in the method signature the only visible exception was `MalformedURLException`. Through the implemented changes, including *class rename* refactoring, the more general exceptions have been specialized (*e.g.*, to `ClientException` in the case of the `getClient` method), forcing the exposure in the signature of the thrown exception. This has the double effect of (i) giving developers compilation errors if they do not catch the thrown exception, thus preventing bugs; and (ii) using more expressive exception names. ♡ Note that the usage of unchecked exceptions in Java code should not be considered as a “bad smell” since, in general, unchecked exceptions should be used to reveal bugs, while checked exceptions to throw errors that the program should handle [46]. However, the misuse of unchecked exceptions where checked ones are needed can lead to higher chances of introducing bugs (as in the case of PR#238). The automatic identification of these situations is, to the best of our knowledge, a problem still not faced in the research literature.

Preparing Code for Changes (41 instances). This category includes refactorings facilitating the implementation of new features or of other planned activities. Refactorings preparing for future changes are usually implemented in dedicated PRs including major refactorings.

Looking at the comparison highlighted in Table 7, it is possible to state that our taxonomy provides more insights compared to the list of motivations in the previous study. Indeed, the category *Preparing Code For Changes* contains some motivations already highlighted in [92] while others missed such as the need for refactoring operations aimed at moving to open-source. An interesting case in this category is represented by PR #317 of the `sagan` project [17]. As mentioned in the title, the goal of the PR is to “*refactor in preparation for open source*”. Note that this is kind of an exception in our taxonomy, and we decided to put it into the *Preparing Code for Changes* category just because open sourcing a project is, in some way, a decision taken to foster the future evolution of the project. Code refactoring is one of the action items in a checklist defined in issue #179 for preparing the project to be open-sourced [16]; other action items included, for example, the introduction of installation and configuration instructions. ♡ This suggests that an appropriate code cleanup, including refactoring, should be part of packaging checklists when putting a project in the open-source.

Another example of refactoring performed to accommodate other changes is in PR#136 of the `zhcet-web` project [23], where the developer extracted the class `CryptoUtils` from `SecurityUtils` to accommodate the implementation of new functionalities (*e.g.*, the `decrypt` method) in a suitable class (*i.e.*, `CryptoUtils`).

♡ Tools supporting preemptive refactoring are lacking in the literature. Indeed, the only effort in this direction is the work by Pantiuchina *et al.* [85] in which, however, the focus is on identifying classes that will be affected by code smells in the future, thus recommending them for a preemptive refactoring action. ♡ Our manual analysis indicates that a novel family of recommender systems able to suggest developers how to refactor the code in order to “accommodate” the implementation of a given change request could be valuable.

Improve Quality of Test Code (49 instances). As the production code might be in need for refactoring, this also holds for test code [104]. The quality of the test code is also assessed in the context of PR discussion and code reviews, as observed by Spadini *et al.* [94]. We found 49 test code refactoring cases, 37% of which performed with a specific type of refactoring operation, and 63% with multiple operations. The observed refactorings include changes similar to those performed on production code, *e.g.*, better distribution of responsibilities to have a better mapping between test and production code, see *e.g.*, PR #1071 in the `error-prone` project [10] in which the author comments “[...] *separate the tests into logical classes*”.

Other cases we found concern the removal of flaky tests which introduce non-determinism in test outcome [77]. In the `microprofile-fault-tolerance` project, the PR #363 [9] aims at removing flaky test: “*The tests `testCircuitInitialSuccessDefaultSuccessThreshold` and `testCircuitLateSuccessDefaultSuccessThreshold` were moved to an independent test to avoid dependencies between tests that use the same bean [...] that can generate possible failures when the circuit breaker leaves open [...]*”.

More interesting and specific for test code are the refactorings performed to improve testability. In PR #73 of the WPS project [1] a developer performed an *extract class* refactoring in production code motivated by the will of simplifying integration testing: “*The functionality to create a `GTVectorDatabinding` out of `shapefiles` was removed from the `GenericFileData` class and moved to a new `GenericFileDataWithGT` class. Due to this change, the processes used for the integration tests do not depend on `GeoTools` anymore [...] also the tests now use only local resources*”. ♡ This example shows how refactoring performed on production code can have an impact on many software quality aspects (in this case, testability). Cases like this one suggest to always ponder the positive and negative impacts of refactoring beyond maintainability *e.g.*, some refactoring actions can aid testability, while others might improve maintainability at the cost of testability. ♡ Clearly, considering this aspect in the context of a refactoring recommender is far from trivial, given the need for automatically assess the testability of a given component. Besides, the presence of refactorings specifically aimed at improving testability shows room for approaches aimed at recommending such kind of operations.

While the motivations of refactorings for simplifying testing activities were already presented in the work by Silva *et al.* [92], our taxonomy provides other new categories such as those aimed at removing flakiness, at automating testing activities, or at improving the overall testability of the project under development.

Other Motivations (108 instances). In this category, we put all motivations that did not find their place among other root categories.

Thirty-seven PRs were performed to improve software performance. This is not surprising and in line with Fowler [56], who observed that the internal program structure is closely related to its performance due to better optimization opportunities. Moreover, the latter also emerges from the motivation provided in the previous study by Silva *et al.* [92]. Also, our quantitative results of RQ₁ indicate that refactorings have a high chance to occur on classes frequently subject to bug fixes, which may have affected performance, especially in the case of quick patches. A concrete example is the PR #1577 of `AmazeFileManager` [19] Android application. In a linked issue, a user reports that when she is “*copying a large file using SFTP, the process can take more than 1 minute, so the phone goes in stand by mode*”. PR #1577 improves the I/O performance of this feature through refactoring.

♡ This example confirms the importance of specific non-functional attributes (in this case, performance) for different types of software (in this case, a mobile app). Also, once again, it points to the need for developing refactoring techniques able to consider this heterogeneity of non-functional requirements rather than mainly focusing on maintainability as done in state-of-the-art refactoring tools [36, 99]. Indeed, to the best of our knowledge, only a few authors have developed refactoring techniques having the improvement of performance as the main objective [31, 50, 107]; however, these approaches either target very specific performance issues [50, 107] or are designed to work on models rather than on source code [31].

4 THREATS TO VALIDITY

Construct validity. A source of inaccuracy is represented by the automated refactoring detection. However, RMiner has been reported to exhibit a very high precision (98%) and recall (87%) [100]. This threat is mitigated at least in RQ₂, where refactorings have been manually reviewed. To identify bug fixes, we used an approach matching regular expressions onto commit messages [53], as also done in previous work [88]. To limit threats due to this heuristic [30], two authors independently analyzed, for each project we considered, a random sample of commits classified as a bug fix to mark true and false positives. After discussing disagreements, only 8% of the analyzed commits resulted to be false positive bug fixes (mostly related to CheckStyle fixes).

In RQ₁, we only analyzed the correlation between the presence of any refactoring with various metrics. While it may be interesting correlating specific types of refactorings with metrics, our qualitative analysis showed that refactoring goals are often achieved through a combination of refactorings. To build the explanatory model of RQ₁, we have selected a broad set of metrics capturing different aspects of software product and process. It is important to note that the aim was to correlate such metrics with the presence of at least one refactoring action of any kind. Building models for specific refactoring types is out of scope of this paper and could, possibly, require to identify further specific indicators.

In RQ₂ we identified refactoring-related PRs as those having (i) one of their commits containing a refactoring identified by RMiner, and (ii) a refactoring-related keyword in their title. Such selection criteria can result in false negatives (*i.e.*, missing some refactoring-related PRs) and, in turn, this may have resulted in missing categories in our taxonomy. Indeed, it is possible that our taxonomy is only representative of the motivations behind PRs that can be captured through the adopted selection criteria.

As context for our study we targeted non-personal/toy projects having a substantial change history to study and being active. For this reason, we defined a number of selection criteria (*i.e.*, at least 5 contributors, 1 fork, 500 commits, 100 PRs, and one recent commit) that, however, may fail in capturing the type of systems we were interested in.

Internal validity. In the quantitative analysis (RQ₁), although we tried to capture factors from different dimensions (*i.e.*, different kinds product and process metrics), there could be many other factors that could have influenced the need for refactoring. We mitigated this threat through (i) the use of a mixed-model considering project as random effect, and (ii) the qualitative analysis of RQ₂. It is important to note that the aim of RQ₁ is to mainly identify correlations between metrics and refactoring activities, and not about claiming any causation. Only the qualitative analysis of RQ₂, taking into account developers' discussion, can highlight the rationale for refactoring actions.

Conclusion validity. In RQ₁ we performed a careful preprocessing of data and variable selection to avoid multi-collinearity, and normalized metrics to allow properly interpreting the ORs.

External validity. Our analysis is limited to a sample of 150 Java open source projects hosted on GitHub, and the qualitative analysis to 551 PRs. We do not claim the generalizability of our findings to other programming languages or to industrial systems. For this reason, a further investigation on a more diverse set of projects, developed with different programming languages and belonging to both open and closed-source is highly desirable. Also, it is worth mentioning that in our manual analysis we only considered PRs for which RMiner identified at least one refactoring operation. This means that we did not consider

PRs that, for example, targeted a complete remodularization of the system that involved refactoring operations not captured by RMiner.

5 RELATED WORK

Different studies have empirically analyzed refactoring operations from different perspectives, including how developers perform refactoring [81]; the relationship between refactoring and other software-related activities (*e.g.*, merge conflicts [73]); the impact of refactoring operations on the likelihood of introducing bugs [35]; the impact of refactoring on specific quality indicators (*e.g.*, quality metrics) [27, 45, 96, 97], or on developers' productivity [79]. While these studies mine and analyze refactoring for different purposes, they address different research questions as compared to the ones subject of our work. For this reason, we focus the discussion on studies analyzing why developers perform refactoring.

Wang *et al.* [106] interviewed 10 developers from four software companies to reveal the major factors motivating refactoring operations. Results highlight external motivators *e.g.*, *Recognitions from Others*, and intrinsic motivators, *i.e.*, when refactoring is initiated without any obvious external reward (*e.g.*, *Self Esteem*) behind refactorings.

Kim *et al.* [61] present a field study surveying and interviewing 328 Microsoft engineers to investigate when and how engineers refactor code. They identify the low readability of source code as the most important symptom that pushes developers to perform refactoring (mentioned by 21% of developers). Our quantitative analysis confirms the central role of low code readability in triggering refactoring. Other frequently mentioned symptoms were code duplication (11%), and fostering code reuse (10%).

Silva *et al.* [92] observe that the previously discussed works [61, 106] report findings from surveys asking developers about their general refactoring habits, without focusing on real refactorings they performed. Thus, they use RMiner [100] to monitor refactorings performed in open source repositories and contacted the developers authoring the refactorings asking to motivate the performed changes. Then, they grouped the responses and defined a catalog of 44 motivations for 12 refactoring operations. For example, they find that three main motivations exist for *Rename Package* refactoring: improve the package name, enforce naming consistency, and move package to appropriate container [92].

While we share with previous work [61, 92, 106] the goal of identifying factors motivating developers to perform refactoring operations, we adopt a different and complementary experimental design. First, we quantitatively investigate process- and product-related factors that may correlate with refactoring actions (*e.g.*, does low code quality as assessed by quality metrics trigger refactoring operations?). Second, we qualitatively analyze refactorings performed by developers in the context of 551 merged PRs, to complement and validate the refactoring motivators already identified in the literature [61, 92, 106]. Note that, as for the work by Silva *et al.* [92] and differently from the ones by Wang *et al.* [106] and Kim *et al.* [61], we look at motivations related to specific refactoring operations (*i.e.*, those performed in the analyzed PRs). Differently from Silva *et al.* [92], we do not rely on answers collected through a survey, but we inspected the code changes performed in the commits related to the PRs, and read the reviewing process carried out before merging the PR and the related discussion. As previously discussed, we validated and complemented the taxonomy of motivations they defined.

Peruma *et al.* [86] also look at the motivations pushing developers to refactor their code, but they focused only on rename refactoring operations, finding that, in most cases, renaming is applied to narrow the identifier meaning. Our qualitative study confirmed this finding.

Bavota *et al.* [40] study the relationship between metrics, code smells and refactoring. Their findings indicate that there is no clear causation between code having a smell or exhibiting elevated complexity metric values and subsequent refactoring of this code. Our study, performed at the commit level, confirms their results but also points out that readability metrics and process-related factors play a significant role.

Finally, very related work in this line of research is the one by Vassallo *et al.* [105], in which the authors mine 200 systems to quantitatively investigate factors correlating with refactoring. They consider, in isolation, factors related to when, why, and by whom refactoring is performed. They found that (i) refactorings are mostly performed after one year from the project startup and rarely close to a new release; (ii) most of refactorings are performed while enhancing existing features; and (iii) developers that refactor code are often the owners of the impacted files. Our work, besides quantitatively investigating a more comprehensive set of process- and product-related factors (considered altogether in a mixed model), also complements this analysis with a qualitative investigation on the motivations behind refactorings.

6 CONCLUSION

In this paper we quantitatively and qualitatively analyzed the reasons behind refactoring operations performed by developers. Our quantitative analysis highlighted that (i) code readability is the product-related factor mostly correlated with refactoring operations, and (ii) process-related factors such as source code change- and fault-proneness and, especially, the experience of developers changing a code component, play a significant role in triggering refactoring operations. Our qualitative analysis resulted in an extensive taxonomy of 67 motivations behind refactorings, relating to quantitative results where possible. We have made the study material and data available in our replication package [28].

The implications of our study trigger several directions for future work.

Identifying “when” to trigger refactoring recommendations. Our quantitative study (RQ₁) shed some light on the product and process-related factors that contribute to trigger refactoring operations. Such an empirical evidence represents the basis for future approaches able to predict “when”, during a project’s evolution to trigger refactoring recommendations. Indeed, such an aspect is currently ignored in the refactoring recommenders literature, with researchers focusing their attention on the core problem of generating meaningful recommendations. However, recommending refactorings in a “context” in which developers do not feel in need to refactoring their code is unlikely to provide benefits. An interesting research direction in this field is to build on top of our RQ₁’s findings to devise models able to predict when refactoring recommendations would be welcome by software developers.

On the relevance of “semantic metrics” for code smell detectors. As shown in our RQ₁, the product metrics exhibiting the stronger relationship with refactoring operations are those exploiting textual information extracted from the code, usually referred to as semantic metrics. This was the case for the readability and for the C3 (Conceptual Cohesion of Classes) metrics. This result somehow supports the previous findings reported in the literature indicating a stronger alignment between the developers’ perception of code quality and semantic metrics as compared to structural ones [38]. These findings suggest the adoption of these metrics in code smell detectors that, at the end, are used to identify refactoring opportunities.

On learning refactoring operations. While discussing the results of our qualitative study (RQ₂), a bold message repeatedly emerged from the analyzed cases: the motivations behind a refactoring are variegated and heterogeneous and, probably, cannot be captured by any combination of metrics. This implies that refactoring recommenders using product-

and/or process-metrics as fitness functions to recommend refactorings will always miss many refactoring opportunities. Indeed, these techniques mostly target complex code components for refactoring (*e.g.*, God Classes in the case of extract class refactoring), ignoring all the other scenarios we documented (*e.g.*, applying extract class to remove a code clone). One possibility worth exploiting in the future is the application of deep learning techniques to refactoring recommenders. Indeed, recent work already shown the possibility to learn from code changes [103]. However, no previous work has attempted to design refactoring recommenders learning from developers' activities what a meaningful refactoring is in a given context.

Automatic support for relevant refactorings is missing. In RQ₂ we found many cases of manually performed refactoring operations that could benefit from the development of techniques to automate such code transformations. This includes the identification and automatic refactoring of redundant code and misuse of checked/unchecked exceptions in Java as well as the lack of support for preemptive refactoring (*i.e.*, refactoring operations aimed at accommodating future changes). We see these areas as possible directions for future work in software refactoring.

Lack of production-ready tools. Some of the cases discussed in RQ₂ highlighted that, while the research community has developed many good approaches for supporting specific refactorings, these approaches find little application in practice. One clear example of this are the rename refactoring operations suggested by reviewers to expand abbreviations used in the identifiers of the code contributed in PRs. These recommendations could be easily generated at commit time by one of the many approaches proposed in the literature for the automatic expansion of abbreviations (see *e.g.*, [67]). However, the lack of production-ready tools could be the reason for the lack of adoption of such techniques. This is an opportunity not only for researchers, but also for developers interested in building tools useful for the partial automation of code review activities.

On the need for pursuing "trade-offs" when refactoring. Finally, our RQ₂ also highlighted the need for refactoring techniques able to consider the many contrasting objectives that a code transformation brings with it. Indeed, while most of the refactoring recommenders strive to maximize maintainability, our findings show that maintainability is only one of the many aspects considered by developers while refactoring, accompanied by performance, testability, *etc.* Future approaches to support refactoring should consider the pros and cons of the recommended solutions from different perspectives, suggesting operations that are sensible to different quality criteria.

Based on the above findings, our future research agenda will focus on the two points described above: (i) predicting *when* to recommend refactorings, and (ii) developing refactoring tools sensible to different non-functional requirements.

REFERENCES

- [1] 52North/WPS Pull Request #73. <https://github.com/52North/WPS/pull/73>.
- [2] apache/fineract Pull Request #366. <https://github.com/apache/fineract/pull/366>.
- [3] cbeust/testng Pull Request #1481. <https://github.com/cbeust/testng/pull/1481>.
- [4] *Code metrics for Java code by means of static analysis*. <https://github.com/mauricioaniche/ck>
- [5] confluentinc/kafka-connect-elasticsearch Pull Request #251. <https://github.com/confluentinc/kafka-connect-elasticsearch/pull/251>.
- [6] dropwizard/dropwizard Pull Request #488. <https://github.com/dropwizard/dropwizard/pull/488>.
- [7] DSpace/DSpace Pull Request #1083. <https://github.com/DSpace/DSpace/pull/1083>.
- [8] DSpace/DSpace Pull Request #324. <https://github.com/DSpace/DSpace/pull/324>.

- [9] eclipse/microprofile-fault-tolerance Pull Request #363. <https://github.com/eclipse/microprofile-fault-tolerance/pull/363>.
- [10] google/error-prone Pull Request #1071. <https://github.com/google/error-prone/pull/1071>.
- [11] kiegroup/optaplanner Pull Request #150. <https://github.com/kiegroup/optaplanner/pull/150>.
- [12] minio/minio-java Pull Request #238. <https://github.com/minio/minio-java/pull/238>.
- [13] *PMD Source Code Analyzer*. <http://pmd.sourceforge.net>
- [14] samtools/htsjdk Pull Request #1067. <https://github.com/samtools/htsjdk/pull/1067>.
- [15] samtools/htsjdk Pull Request #599. <https://github.com/samtools/htsjdk/pull/599>.
- [16] spring-io/sagan Pull Request #179. <https://github.com/spring-io/sagan/issues/179>.
- [17] spring-io/sagan Pull Request #317. <https://github.com/spring-io/sagan/pull/317>.
- [18] spring-projects/spring-amqp Pull Request #346. <https://github.com/spring-projects/spring-amqp/pull/346>.
- [19] TeamAmaze/AmazeFileManager Pull Request #1577. <https://github.com/TeamAmaze/AmazeFileManager/pull/1577>.
- [20] zalando/nakadi Pull Request #626. <https://github.com/zalando/nakadi/pull/626>.
- [21] zalando/nakadi/ Pull Request #933. <https://github.com/zalando/nakadi/pull/933>.
- [22] zendesk/maxwell Pull Request #392. <https://github.com/zendesk/maxwell/pull/392>.
- [23] zhcet-amu/zhcet-web Pull Request #136. <https://github.com/zhcet-amu/zhcet-web/pull/136>.
- [24] H. Akaike. Information theory and an extension of the maximum likelihood principle. In *2nd International Symposium on Information Theory*.
- [25] Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. Learning Natural Coding Conventions. In *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE 2014)*. 281–293.
- [26] Uri Alon, Meital Zilberstein, Omer Levy, and Eran Yahav. Code2Vec: Learning Distributed Representations of Code. *Proc. ACM Program. Lang.* 3, POPL (Jan. 2019), 40:1–40:29.
- [27] Mohammad Alshayeb. Empirical investigation of refactoring effect on software quality. *Information and Software Technology* 51, 9 (2009), 1319 – 1326. <http://www.sciencedirect.com/science/article/pii/S095058490900038X>
- [28] Anonymous. Replication Package. <https://github.com/replication-package/why-refactoring>.
- [29] Nicolas Anquetil and Timothy Lethbridge. Experiments with Clustering as a Software Remodularization Method. In *WCRE*. 235–255.
- [30] Giuliano Antoniol, Kamel Ayari, Massimiliano Di Penta, Foutse Khomh, and Yann-Gaël Guéhéneuc. Is it a bug or an enhancement?: a text-based approach to classify change requests. In *Proceedings of the 2008 conference of the Centre for Advanced Studies on Collaborative Research*. IBM, 23.
- [31] Davide Arcelli, Vittorio Cortellessa, and Catia Trubiani. Antipattern-based Model Refactoring for Software Performance Improvement. In *Proceedings of the 8th International ACM SIGSOFT Conference on Quality of Software Architectures (QoSA '12)*. 33–42.
- [32] Venera Arnaoudova, Massimiliano Di Penta, Giuliano Antoniol, and Yann-Gaël Guéhéneuc. A New Family of Software Anti-patterns: Linguistic Anti-patterns. In *17th European Conference on Software Maintenance and Reengineering, CSMR 2013*. 187–196.
- [33] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48.
- [34] Gabriele Bavota. Using structural and semantic information to support software refactoring. In *34th International Conference on Software Engineering, ICSE 2012*. IEEE Computer Society, 1479–1482.
- [35] Gabriele Bavota, Bernardino De Carluccio, Andrea De Lucia, Massimiliano Di Penta, Rocco Oliveto, and Orazio Strollo. When Does a Refactoring Induce Bugs? An Empirical Study. In *Proceedings of the IEEE International Working Conference on Source Code Analysis and Manipulation*. 104–113.
- [36] Gabriele Bavota, Andrea De Lucia, Andrian Marcus, and Rocco Oliveto. Automating extract class refactoring: an improved method and its evaluation. *Empirical Software Engineering* (2013), 1–48.
- [37] Gabriele Bavota, Andrea De Lucia, Andrian Marcus, and Rocco Oliveto. Recommending Refactoring Operations in Large Software Systems. In *Recommendation Systems in Software Engineering*. Springer Berlin Heidelberg, 387–419.
- [38] Gabriele Bavota, Bogdan Dit, Rocco Oliveto, Massimiliano Di Penta, Denys Poshyvanyk, and Andrea De Lucia. An empirical study on the developers' perception of software coupling. In *35th International Conference on Software Engineering, ICSE '13*. IEEE Computer Society, 692–701.
- [39] Gabriele Bavota, Andrea De Lucia, Andrian Marcus, and Rocco Oliveto. Using structural and semantic measures to improve software modularization. *Empirical Software Engineering* 18, 5 (2013), 901–932.

- [40] Gabriele Bavota, Andrea De Lucia, Massimiliano Di Penta, Rocco Oliveto, and Fabio Palomba. An experimental investigation on the innate relationship between quality and refactoring. *Journal of Systems and Software* 107 (2015), 1 – 14.
- [41] William J. Brown, Raphael C. Malveau, William H. Brown, Hays W. McCormick III, and Thomas J. Mowbray. *Anti Patterns: Refactoring Software, Architectures, and Projects in Crisis* (1st ed.). John Wiley and Sons.
- [42] Raymond P. L. Buse and Westley Weimer. Learning a Metric for Code Readability. *IEEE Transactions on Software Engineering* 36, 4 (2010), 546–558.
- [43] Gerardo Canfora, Luigi Cerulo, Marta Cimitile, and Massimiliano Di Penta. How changes affect software entropy: an empirical study. *Empirical Software Engineering* 19, 1 (2014), 1–38.
- [44] Diego Cedrim, Alessandro Garcia, Melina Mongiovi, Rohit Gheyi, Leonardo Sousa, Rafael de Mello, Balduino Fonseca, Márcio Ribeiro, and Alexander Chávez. Understanding the Impact of Refactoring on Smells: A Longitudinal Study of 23 Software Projects. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2017)*. 465–475.
- [45] Alexander Chávez, Isabella Ferreira, Eduardo Fernandes, Diego Cedrim, and Alessandro Garcia. How Does Refactoring Affect Internal Quality Attributes?: A Multi-project Study. In *Proceedings of the 31st Brazilian Symposium on Software Engineering (SBES'17)*. 74–83.
- [46] Chien-Tsun Chen, Yu Chin Cheng, Chin-Yun Hsieh, and I-Lang Wu. Exception handling refactorings: Directed by goals and driven by bug fixing. *Journal of Systems and Software* 82, 2 (2009), 333 – 345.
- [47] Shyam R. Chidamber and Chris F. Kemerer. A Metrics Suite for Object Oriented Design. *IEEE Transactions on Software Engineering (TSE)* 20, 6 (June 1994), 476–493.
- [48] Cesar Couto, João Eduardo Montandon, Christofer Silva, and Marco Tulio Valente. Static correspondence and correlation between field defects and warnings reported by a bug finding tool. *Software Quality Journal* 21 (2011), 241–257.
- [49] Florian Deissenbock and Markus Pizka. Concise and Consistent Naming. In *IWPC*.
- [50] Danny Dig. A Refactoring Approach to Parallelism. *IEEE Software* 28, 1 (2011), 17–22.
- [51] Paul Duvall, Stephen M. Matyas, and Andrew Glover. *Continuous Integration: Improving Software Quality and Reducing Risk*. Addison-Wesley Professional.
- [52] Sarah Fakhoury, Yuzhan Ma, Venera Arnaoudova, and Olusola O. Adesope. The effect of poor source code lexicon and readability on developers' cognitive load. In *Proceedings of the 26th Conference on Program Comprehension, ICPC*. 286–296.
- [53] Michael Fischer, Martin Pinzger, and Harald C. Gall. Populating a Release History Database from Version Control and Bug Tracking Systems. In *19th International Conference on Software Maintenance (ICSM 2003), The Architecture of Existing Systems, 22-26 September 2003, Amsterdam, The Netherlands*. 23.
- [54] Marios Fokaefs, Nikolaos Tsantalis, Eleni Stroulia, and Alexander Chatzigeorgiou. JDeodorant: identification and application of extract class refactorings. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011*. ACM, 1037–1039.
- [55] Martin Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, Boston, MA, USA.
- [56] M. Fowler, K. Beck, J. Brant, W. Opdyke, and D. Roberts. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley.
- [57] Mathew Hall, Neil Walkinshaw, and Phil McMinn. Supervised software modularisation. In *28th IEEE International Conference on Software Maintenance, ICSM*. 472–481.
- [58] Frank E Harrell Jr, with contributions from Charles Dupont, and many others. *Hmisc: Harrell Miscellaneous*. <https://CRAN.R-project.org/package=Hmisc> R package version 4.0-3.
- [59] Brian Henderson-Sellers, Larry L. Constantine, and Ian M. Graham. Coupling and cohesion (towards a valid metrics suite for object-oriented analysis and design). *Object Oriented Systems* 3 (1996), 143–158.
- [60] David Kawrykow and Martin P. Robillard. Improving API Usage Through Automatic Detection of Redundant Code. In *Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering (ASE '09)*. 111–122.
- [61] Miryung Kim, Thomas Zimmermann, and Nachiappan Nagappan. A Field Study of Refactoring Challenges and Benefits. In *Proceedings of the 20th International Symposium on Foundations of Software Engineering*.
- [62] Miryung Kim, T. Zimmermann, and N. Nagappan. An Empirical Study of Refactoring Challenges and Benefits at Microsoft. *Software Engineering, IEEE Transactions on* 40, 7 (July 2014), 633–649.

- [63] Sunghun Kim and Michael D. Ernst. Which warnings should I fix first?. In *Proceedings of the joint meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/FSE)*. 45–54.
- [64] Sunghun Kim, Thomas Zimmermann, E. James Whitehead Jr., and Andreas Zeller. Predicting Faults from Cached History. In *29th International Conference on Software Engineering (ICSE 2007)*. 489–498.
- [65] Giri P. Krishnan and Nikolaos Tsantalis. Unification and refactoring of clones. In *Proceedings of the IEEE Conference on Software Maintenance, Reengineering and Reverse Engineering*. 104–113.
- [66] A. Kuhn, S. Ducasse, and T. Girba. Semantic Clustering: Identifying Topics in Source Code. *Information and Software Technology* 49, 3 (2007), 230–243.
- [67] D. Lawrie and D. Binkley. Expanding identifiers to normalize source code vocabulary. In *2011 27th IEEE International Conference on Software Maintenance (ICSM)*. 113–122.
- [68] Dawn Lawrie, Henry Feild, and David Binkley. Syntactic Identifier Conciseness and Consistency. In *SCAM'06*. 139–148.
- [69] Dawn Lawrie, Christopher Morrell, Henry Feild, and David Binkley. What's in a Name? A Study of Identifiers. In *ICPC'06*.
- [70] Dawn Lawrie, Christopher Morrell, Henry Feild, and David Binkley. Effective identifier names for comprehension and memory. *ISSE* 3, 4 (2007), 303–318.
- [71] António Menezes Leitão. Detection of Redundant Code Using R2D2. *Software Quality Journal* 12, 4 (01 Dec 2004), 361–382.
- [72] Bin Lin, Simone Scalabrino, Andrea Mocci, Rocco Oliveto, Gabriele Bavota, and Michele Lanza. Investigating the use of code analysis and nlp to promote a consistent usage of identifiers. In *2017 IEEE 17th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 81–90.
- [73] Mehran Mahmoudi, Sarah Nadi, and Nikolaos Tsantalis. Are Refactorings to Blame? An Empirical Study of Refactorings in Merge Conflicts. In *26th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2019*. 151–162.
- [74] Onaiza Maqbool and Haroon A. Babri. Hierarchical Clustering for Software Architecture Recovery. *IEEE Transactions on Software Engineering* 33, 11 (2007), 759–780.
- [75] A. Marcus and D. Poshyvanyk. The Conceptual Cohesion of Classes. In *Proc. of 21st IEEE ICSM*. IEEE CS Press, Budapest, Hungary, 133–142.
- [76] Andrian Marcus, Denys Poshyvanyk, and Rudolf Ferenc. Using the Conceptual Cohesion of Classes for Fault Prediction in Object-Oriented Systems. *IEEE Trans. Software Eng.* 34, 2 (2008), 287–300.
- [77] John Micco. Flaky tests at Google and how we mitigate them. *Online*] <https://testing.googleblog.com/2016/05/flaky-tests-at-google-and-how-we.html> (2016).
- [78] Naouel Moha, Yann-Gaël Guéhéneuc, Laurence Duchien, and Anne-Françoise Le Meur. DECOR: A Method for the Specification and Detection of Code and Design Smells. *IEEE Transactions on Software Engineering* 36 (2010), 20–36.
- [79] Raimund Moser, Pekka Abrahamsson, Witold Pedrycz, Alberto Sillitti, and Giancarlo Succi. Balancing Agility and Formalism in Software Engineering. Chapter A Case Study on the Impact of Refactoring on Quality and Productivity in an Agile Team, 252–266.
- [80] Raimund Moser, Witold Pedrycz, and Giancarlo Succi. A comparative analysis of the efficiency of change metrics and static code attributes for defect prediction. In *30th International Conference on Software Engineering (ICSE 2008)*. 181–190.
- [81] Emerson Murphy-Hill, Chris Parnin, and Andrea P. Black. How We Refactor, and How We Know It. *Transactions on Software Engineering* 38, 1 (2011), 5–18.
- [82] Matheus Paixão, Mark Harman, Yuanyuan Zhang, and Yijun Yu. An Empirical Study of Cohesion and Coupling: Balancing Optimization and Disruption. *IEEE Trans. Evolutionary Computation* 22, 3 (2018), 394–414.
- [83] Fabio Palomba, Gabriele Bavota, Massimiliano Di Penta, Fausto Fasano, Rocco Oliveto, and Andrea De Lucia. On the diffuseness and the impact on maintainability of code smells: a large scale empirical investigation. *Empirical Software Engineering* 23, 3 (2018), 1188–1221.
- [84] Fabio Palomba, Andrea De Lucia, Gabriele Bavota, and Rocco Oliveto. Anti-Pattern Detection: Methods, Challenges, and Open Issues. *Advances in Computers* 95 (2015), 201–238.
- [85] Jevgenija Pantiuchina, Gabriele Bavota, Michele Tufano, and Denys Poshyvanyk. Towards just-in-time refactoring recommenders. In *Proceedings of the 26th Conference on Program Comprehension, ICPC 2018*. 312–315.

- [86] Anthony Peruma, Mohamed Wiem Mkaouer, Michael J. Decker, and Christian D. Newman. An Empirical Investigation of How and Why Developers Rename Identifiers. In *Proceedings of the 2Nd International Workshop on Refactoring (IWor 2018)*. 26–33.
- [87] Kata Praditwong, Mark Harman, and Xin Yao. Software Module Clustering as a Multi-Objective Search Problem. *IEEE Transactions on Software Engineering* 37, 2 (2011), 264–282.
- [88] Baishakhi Ray, Daryl Posnett, Vladimir Filkov, and Premkumar T. Devanbu. A large scale study of programming languages and code quality in github. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 155–165.
- [89] B. Rosner. *Fundamentals of Biostatistics* (7th edition ed.). Brooks/Cole, Boston, MA.
- [90] Chanchal K. Roy. Large scale clone detection, analysis, and benchmarking: An evolutionary perspective. In *12th IEEE International Workshop on Software Clones, IWSC 2018*.
- [91] Simone Scalabrino, Mario Linares-Vásquez, Rocco Oliveto, and Denys Poshyvanyk. A Comprehensive Model for Code Readability. *Journal of Software: Evolution and Process* 30, 6, e1958.
- [92] Danilo Silva, Nikolaos Tsantalis, and Marco Tulio Valente. Why we refactor? confessions of GitHub contributors. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016*. 858–870.
- [93] Jaime Spacco, David Hovemeyer, and William Pugh. Tracking defect warnings across versions. In *Proceedings of the 2006 international workshop on Mining software repositories*. ACM, 133–136.
- [94] Davide Spadini, Fabio Palomba, Tobias Baum, Stefan Hanenberg, Magiel Bruntink, and Alberto Bacchelli. Test-driven code review: an empirical study. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019*. 1061–1072.
- [95] Donna Spencer. *Card sorting: Designing usable categories*. Rosenfeld Media.
- [96] Konstantinos Stroggylos and Diomidis Spinellis. Refactoring—Does It Improve Software Quality?. In *Proceedings of the 5th International Workshop on Software Quality (WoSQ '07)*. IEEE Computer Society, Washington, DC, USA, 10–. <http://dx.doi.org/10.1109/WOSQ.2007.11>
- [97] Gábor Szoke, Gábor Antal, Csaba Nagy, Rudolf Ferenc, and Tibor Gyimóthy. Bulk Fixing Coding Issues and Its Effects on Software Quality: Is It Worth Refactoring?. In *Source Code Analysis and Manipulation (SCAM), 2014 IEEE 14th International Working Conference on*. IEEE, 95–104.
- [98] Armstrong Takang, Penny Grubb, and Robert Macredie. The effects of comments and identifier names on program comprehensibility: an experiential study. *Journal of Programming Languages* 4, 3 (1996), 143–167.
- [99] Nikolaos Tsantalis and Alexander Chatzigeorgiou. Identification of Move Method Refactoring Opportunities. *IEEE Transactions on Software Engineering* 35, 3 (2009), 347–367.
- [100] Nikolaos Tsantalis, Matin Mansouri, Laleh M. Eshkevari, Davood Mazinianian, and Danny Dig. Accurate and Efficient Refactoring Detection in Commit History. In *Proceedings of the 40th International Conference on Software Engineering (ICSE '18)*. 483–494.
- [101] Nikolaos Tsantalis, Davood Mazinianian, and Shahriar Rostami. Clone refactoring with lambda expressions. In *Proceedings of the 39th International Conference on Software Engineering, ICSE 2017, Buenos Aires, Argentina, May 20-28, 2017*. 60–70.
- [102] Michele Tufano, Fabio Palomba, Gabriele Bavota, Rocco Oliveto, Massimiliano Di Penta, Andrea De Lucia, and Denys Poshyvanyk. When and Why Your Code Starts to Smell Bad (and Whether the Smells Go Away). *IEEE Trans. Software Eng.* 43, 11 (2017), 1063–1088.
- [103] Michele Tufano, Jevgenija Pantiuchina, Cody Watson, Gabriele Bavota, and Denys Poshyvanyk. On learning meaningful code changes via neural machine translation. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019*. 25–36.
- [104] Arie Van Deursen, Leon Moonen, Alex Van Den Bergh, and Gerard Kok. Refactoring test code. In *Proceedings of the 2nd international conference on extreme programming and flexible processes in software engineering (XP2001)*. 92–95.
- [105] Carmine Vassallo, Giovanni Grano, Fabio Palomba, Harald Gall, and Alberto Bacchelli. A large-scale empirical exploration on refactoring activities in open source software projects. *Science of Computer Programming* 180, 1 (2019), 1–15.
- [106] Yi Wang. What motivate software engineers to refactor source code? evidences from professional developers. In *Software Maintenance, 2009. ICSM 2009. IEEE International Conference on*. 413–416.
- [107] Ying Zhang, Gang Huang, Xuanzhe Liu, Wei Zhang, Hong Mei, and Shunxiang Yang. Refactoring Android Java Code for On-demand Computation Offloading. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA '12)*.

