

Towards a catalog of composite refactorings

Aline Brito  | Andre Hora | Marco Tulio Valente

Department of Computer Science, Federal University of Minas Gerais (UFMG), Belo Horizonte, Brazil

Correspondence

Aline Brito, Department of Computer Science, Federal University of Minas Gerais (UFMG), Belo Horizonte, Brazil.

Email: alinebrito@dcc.ufmg.br

Funding information

Conselho Nacional de Desenvolvimento Científico e Tecnológico; Coordenação de Aperfeiçoamento de Pessoal de Nível Superior; Fundação de Amparo à Pesquisa do Estado de Minas Gerais

Abstract

Catalogs of refactoring have key importance in software maintenance and evolution, since developers rely on such documents to understand and perform refactoring operations. Furthermore, these catalogs constitute a reference guide for communication between practitioners since they standardize a common refactoring vocabulary. Fowler's book describes the most popular catalog of refactorings, which documents single and well-known refactoring operations. However, sometimes, refactorings are composite transformations, that is, a sequence of refactorings is performed over a given program element. For example, a sequence of Extract Method operations (a single refactoring) can be performed over the same method, in one or in multiple commits, to simplify its implementation, therefore, leading to a Method Decomposition operation (a composite refactoring). In this paper, we propose and document a catalog with eight composite refactorings. We also implement a set of scripts to mine composite refactorings by preprocessing the results of refactoring detection tools. Using such scripts, we search for composites in a representative refactoring oracle with hundreds of confirmed single refactoring operations. Next, to complement this first study, we also search for composites in the full history of 10 well-known open-source projects. We characterize the detected composite refactorings, under dimensions such as size and location. We conclude by addressing the applications and implications of the proposed catalog.

KEYWORDS

composite refactoring, mining software repositories, refactoring, software evolution

1 | INTRODUCTION

Refactoring is a fundamental practice to keep software in a healthy shape. Developers have learned over the years that the lack of continuous refactoring can rapidly transform software projects into a “big-ball-of-mud”.^{1,2} As a result, maintenance—including both bug fixes and the implementation of new features—becomes very risky and slow. Therefore, to promote and facilitate the dissemination of this practice among developers, refactoring operations are usually documented in catalogs, like the one proposed by Fowler in his seminal book published in 1999.³ In this catalog, Fowler provides detailed documentation about dozens of refactorings, providing a name for each refactoring, describing the mechanics required to perform the source code transformation, and also giving illustrative examples of the proposed refactorings. However, most refactorings described in Fowler's catalog are restricted in time and scope. Particularly, they are described as source code transformations that can be performed by a single developer, in a short time frame (time constraint) and by impacting a limited number of program elements (scope constraint). This understanding of refactoring is also assumed by modern refactoring detection tools, such as RefactoringMiner^{4–6} and RefDiff.^{7–9} Indeed, these tools report refactorings at a very fine granularity level. For example, suppose that a given method $m()$ is implemented in classes A_1, A_2, \dots, A_n . Then, suppose that a Pull Up refactoring is performed to move the replicated method to a superclass B.

These tools report this refactoring as a sequence of the following unrelated operations:

```
Pull up: A1.m() to B Pull up: A2.m() to B Pull up: A3.m() to B
...
```

However, we claim that the best output would be reporting a single composite refactoring operation:

```
Pull up: A1.m(), A2.m(), A3.m(), ..., to B
```

This is just a trivial example of composite refactoring (in Section 3, we provide a more complex example). Indeed, composite refactorings were previously defined by Souza et al. as “two or more interrelated refactorings that affect one or more elements”.¹⁰ However, in their work, the authors focused on the role played by composite refactorings when removing code smells. In other words, they do not explore, document, and illustrate a comprehensive catalog of composite refactorings, which is exactly our goal in this paper.

We initially describe eight composite refactorings, in abstract terms and using illustrative examples. Then, we implemented a set of scripts to identify these composite refactorings. To conclude, we used these scripts to mine composite refactorings in two datasets, as follows:

Oracle study: In the first study, we mine composites in a large and representative refactoring oracle commonly used in the literature.^{4,11} Specifically, we look for occurrences of the refactorings in our catalog among 1.7 K confirmed single refactoring instances listed in this oracle. We identify that a significant rate of 60% of the refactorings of interest in this oracle are part of composite operations. We also characterize the detected composite refactorings, under dimensions such as size and scope.

Study in the wild: In the oracle study, we rely on a sample that includes selected refactorings from distinct projects. Therefore, as a complementary analysis, we also look for composite refactorings in the full history of 10 popular GitHub projects. As a result, we were able to identify and characterize 2886 instances of composite refactorings.

Therefore, our contributions are threefold: (1) we propose a comprehensive catalog of composite refactorings, (2) we implemented a set of scripts to detect the refactorings proposed in this catalog, and (3) we characterize a large sample of composite refactorings performed in real software projects, including a new viewpoint of a well-known refactoring oracle.

Paper structure: Section 2 shows a first example of composite refactoring, while Section 3 introduces the proposed catalog. We describe the results of the oracle study in Section 4, while Section 5 includes results in the wild, covering the full evolution history of 10 popular open-source projects. The results are then discussed in Section 6. Section 7 states threats to validity, and Section 8 presents related work. Finally, we conclude the paper in Section 9.

2 | AN EXAMPLE OF COMPOSITE REFACTORING

Figure 1 shows a real example of composite refactoring in Spring, a well-known Web framework for Java. As we can see, method `doDispatch(...)` was decomposed by performing six Extract Method refactorings. Moreover, in two cases, the Extract was followed by a Move Method. These operations were performed in a single commit.

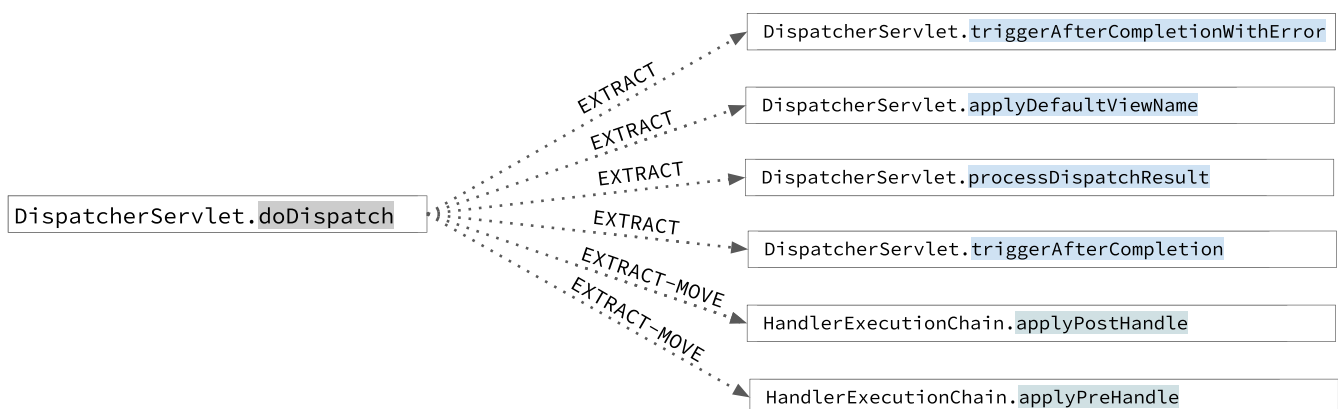


FIGURE 1 Example of composite refactoring from *Spring Framework*. Method `doDispatch` was decomposed by applying six Extract Method refactorings

When used in a context like this one, a refactoring detection tool, such as RefDiff⁷⁻⁹ or RefactoringMiner,^{4,5} detects these six single refactorings independently. However, it would be interesting to detect a high-level refactoring operation, that is, a composite refactoring grouping the six transformations. As we detailed in the next section, we call Method Decomposition refactoring this particular coarse-grained operation.

Techniques that may benefit from the detection of independent refactorings (like code visualization,¹²⁻¹⁴ code review,¹⁵⁻¹⁹ code authorship,^{20,21} bug-introducing detection,^{22,23} refactoring-aware tools,^{24,25} software mining approaches,²⁶⁻²⁹ to name a few) may also benefit from the detection of composite refactorings. As refactoring detection is the basis of such techniques, composite refactorings would bring to light novel operations not restricted to time and scope, therefore better representing the actual source code changes.

Before presenting our catalog, it is important to mention that composite refactorings are not limited to a single commit.¹⁰ For example, as stated by Fowler in the new version of his book on refactoring,³⁰ there are also *long-term* refactorings “that can take a team weeks to complete”.

3 | CATALOG OF COMPOSITE REFACTORINGS

In this section, we introduce the proposed catalog of composite refactorings. As customary in refactoring catalogs, we describe the proposed refactoring types and their mechanics. We also present an abstract example of each composite refactoring. There are two main groups of refactorings: (i) to decompose program structures (five composite refactorings) and (ii) to create program structures (three composite refactorings). To propose these composites, we basically leveraged our previous research on large refactorings operations. For example, we started this research by characterizing refactoring operations performed over time in Java projects.¹⁴ Later, we extended this work by including JavaScript projects and also performed a survey with the developers who were responsible for these operations.³¹

3.1 | Class decomposition

Motivation: According to Fowler,^{3,30} during software evolution, we might need to “move elements around”, aiming to improve modularity and cohesion and reduce coupling. Specifically, single Move Method operations should be performed “when classes have too much behavior or when classes are collaborating too much and are too highly coupled”. However, the overall solution may not be restricted to a single refactoring operation. Instead, we might need to move more than one method from a single source class. In this case, we say we performed a composite refactoring called Class Decomposition.

Mechanics: Figure 2 shows an example, in which class `Foo` lost multiple methods to classes `Bar` and `Baz`. The target class can be existing or new. Also, the Move operations can be followed by a Rename operation. In all cases, the final goal is to decompose the source class and make its implementation more cohesive. It is also worth noting that our definition does not require all move operations to be performed in a single commit. In other words, they can be spread over time, in multiple commits.

3.2 | Method decomposition

Motivation: We perform Extract Method operations when “you have to spend effort looking at a fragment of code and figuring out what it's doing”.^{3,30} In other words, Fowler advocates the improvement of understandability as the main reason to perform method extractions. However, the solution does not need to be limited to a sole operation. We could perform a sequence of two or more Extract Method operations over a single method. As a result, it generates a simpler one. Evidently, these refactorings also generate new methods. However, the goal is still the decomposition of the source method. In this case, we say we performed a composite refactoring called Method Decomposition.

Mechanics: Figure 3 shows an abstract example, in which methods `m1()` and `m2()` were extracted from method `m()`. After the extractions, the new methods can be moved to a distinct class, as happened with `m2()`. As usual, the operations can be performed in one or multiple commits.




FIGURE 2 Class decomposition

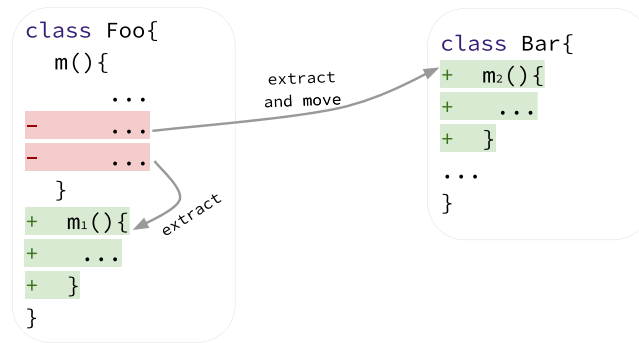


FIGURE 3 Method decomposition

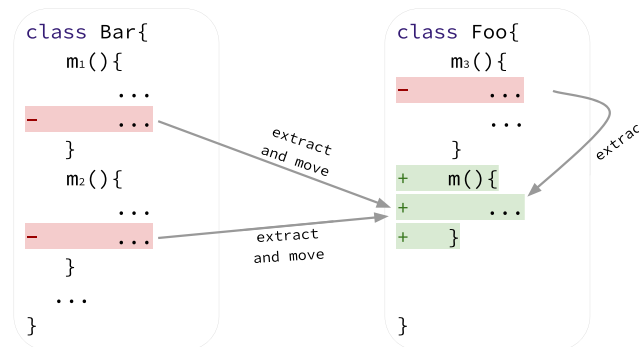


FIGURE 4 Method composition

3.3 | Method composition

Motivation: Extractions also can be performed to promote reuse and to remove duplication.^{3,30,32} Particularly, in such cases, we have similar fragments of code scattered over multiple locations. Therefore, a single Extract Method operation does not eliminate the duplicated issue. Instead, it may be necessary to apply multiple extractions to remove the duplicated code, generating a new method. In this case, we say we performed a composite refactoring called Method Composition.

Mechanics: Two or more Extract Method operations are performed over duplicated code, as illustrated in Figure 4. This code is then removed, and a new method is created, with the previously duplicated code. The operations also can be followed by Move Method operations, that is, the new method is placed in a distinct class.

3.4 | Composite inline method

Motivation: Inline Method—as originally proposed in Fowler's catalog—is reported as the opposite operation of Extract Method. The author suggests applying a set of Inline Method operations to remove trivial methods.³⁰ However, Inline Method is usually detected as a single operation by current refactoring detection tools.^{4,7} That is, such tools report independent Inline operations, even when they are part of the same group of operations. Therefore, we decided to include this refactoring in our catalog, since it matches our criteria for composite refactorings and is not properly explored and detected by current tools.

Mechanics: We expand a (simple) method body in its call sites, as shown in Figure 5. Then, we remove the source method. The calls may be located in methods from distinct classes.

3.5 | Composite pull up method

Motivation: Fowler also points to the need to move up or down methods in inheritance hierarchies.³⁰ In this context, we apply sequences of Pull Up Method to create a single and more general method in the superclass, therefore achieving code reuse. As in the case of Inline, we decided to include this refactoring in our catalog mainly because Pull Up operations are reported as individual and independent operations by current refactoring detection tools.

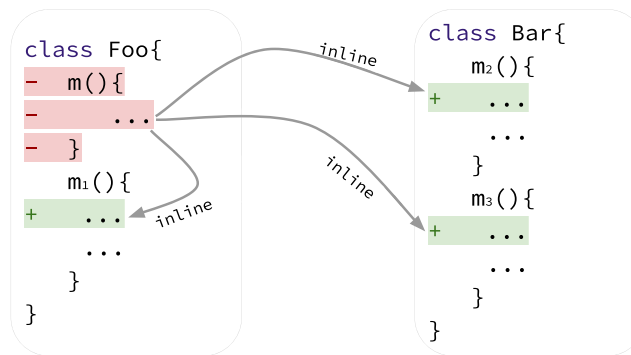


FIGURE 5 Composite inline method

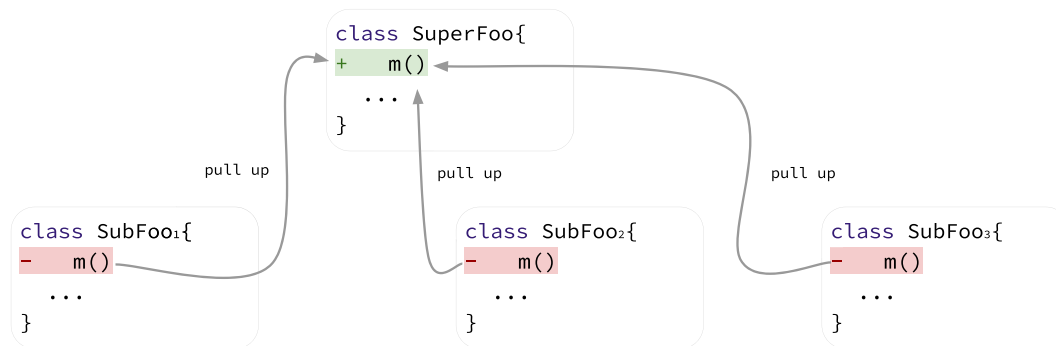


FIGURE 6 Composite pull up method

Mechanics: This operation refers to sequences of transformations performed to move methods from subclasses to their superclass. For example, consider a class `SuperFoo` with subclasses `SubFoo1`, `SubFoo2`, and `SubFoo3`, as presented in Figure 6. Suppose that a Pull Up operation is applied to move method `m()` from these subclasses to the superclass. Usually, this operation occurs in a single commit. First, a developer copies the method `m()` to the superclass, which can be an existing or new one. After that, the method is removed from the subclasses. In this context, the following three messages are issued by RefactoringMiner⁴:

Pull Up Method `public m() : void` from class `SubFoo1` to `public m() : void` from class `SuperFoo`

Pull Up Method `public m() : void` from class `SubFoo2` to `public m() : void` from class `SuperFoo`

Pull Up Method `public m() : void` from class `SubFoo3` to `public m() : void` from class `SuperFoo`

However, since essentially they are part of the same composite refactoring, we claim that these operations should have been reported using a single and comprehensive message, such as:

Pull Up method `public m() : void` From: `SubFoo1`, `SubFoo2`, and `SubFoo3` To: `public m() : void` in `SuperFoo`

3.6 | Composite push down method

Motivation: As an opposite scenario, we perform Push Down Method when a method is needed only in a few subclasses.^{3,30} Therefore, this refactoring promotes inheritance simplification. This operation—also present in Fowler's catalog—matches our criteria for composite refactoring. However, as in the case of Pull Up and Inline, it is reported as independent operations by current refactoring mining tools.

Mechanics: This operation moves a given method from the superclass to particular subclasses, as presented in Figure 7. After that, the method is removed from the superclass.

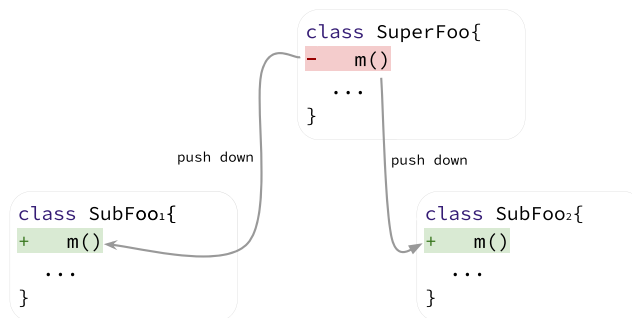


FIGURE 7 Composite push down method

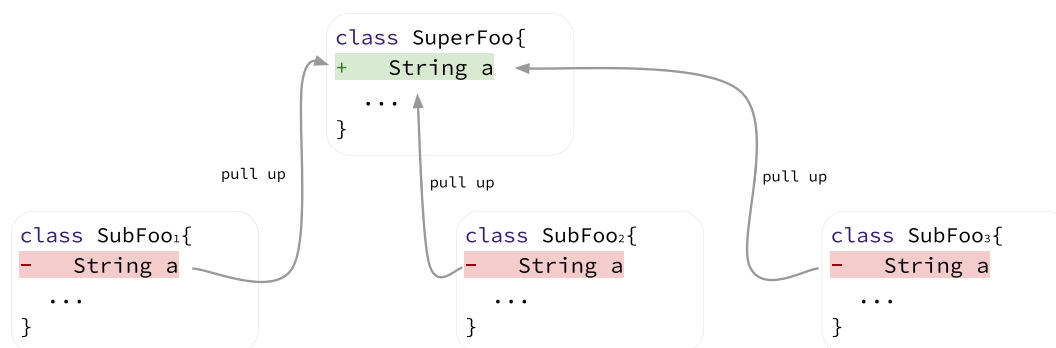


FIGURE 8 Composite pull up field

3.7 | Composite pull up field

Motivation: Often, we have duplicate data in inheritance hierarchies, for example, fields used for a similar purpose in distinct subclasses. In this case, we can perform a sequence of Pull Up Field to create a single one in the superclass, aiming to promote reuse.³⁰ Therefore, this operation also corresponds to our criteria for composite refactoring, and we say we performed a Composite Pull Up Field.

Mechanics: First, we declare the field in the superclass. Then, we remove the declaration in the subclasses, as shown in Figure 8. This operation can also be preceded by Rename Field, aiming to standardize the names before the movement to the superclass.

3.8 | Composite push down field

Motivation: Similar to Push Down Method, the goal involves moving data from a superclass to specific subclasses.³⁰ When this operation contemplates a sequence of fields movements, we say we performed a Composite Push Down Field.

Mechanics: First, we declare the field in the required subclasses. Then, we remove the declaration in the superclass, as shown in Figure 9.

3.9 | A final note on completeness

The original catalog of refactorings proposed by Fowler has dozens of refactorings. Therefore, the catalog of composites described in this section is much smaller (eight composites). On the one hand, this difference is expected because composites are coarse-grained and complex source code transformations, composed by atomic refactorings. On the other hand, it is also important to acknowledge that we do not claim on the completeness of the proposed catalog. Indeed, our central intention is to provide a comprehensive, well-documented, and easy to understand initial list of composite refactorings. In future studies, we can extend the list to include other types of composites. For example, most current instances are refactorings at the method level, since they are among the most frequent code elements affected by such operations.^{4,7} However, it is possible to extend the catalog by including, for example, operations to compose or decompose packages.

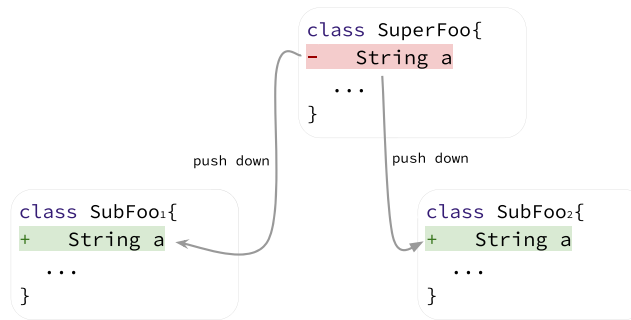


FIGURE 9 Composite push down field

4 | A FIRST ORACLE OF COMPOSITE REFACTORING

To investigate whether the proposed composite refactorings occur in real projects, we initially search for composite refactorings in one of the most representative refactoring oracles in the literature, curated by Tsantalis and other researchers.^{4,5,11} This oracle has been expanded over the years. The latest version includes more than 14 K refactoring operations from 185 Java projects. The oracle instances were validated by multiple authors and/or well-known tools. In other words, it is a trustworthy dataset for studying refactoring practices.

4.1 | Study design

4.1.1 | Research questions assessment

We propose two research questions:

(RQ1) What are the most common composite refactorings in the oracle? In the current version of the oracle, refactoring operations are reported as individual (i.e., noncomposite) ones. Thus, in this first RQ, our goal is to explore the oracle data from a new perspective, looking for occurrences of composite refactorings. In other words, we aim to provide a new oracle view, which is not based on individual refactoring operations. For this purpose, we first compute the frequency (i.e., the number of occurrences) of each composite instance.

(RQ2) What are the characteristics of composite refactorings in the oracle? The rationale of this second research question is to understand the main characteristics of the composite refactorings detected in RQ1. Therefore, for each composite instance, we compute information such as its scope (i.e., location of the entities before and after a refactoring operation) and size (i.e., number of individual refactoring operations).

4.1.2 | Dataset

In January 2022, we retrieved the most recent oracle version. Then, we selected only refactoring operations that could be part of composite operations.[†] For example, the original oracle includes refactorings such as Move Attribute and Rename Method, which are not related at all with the composite refactorings described in Section 3.


As presented in Table 1, our oracle sample includes 1725 individual refactoring instances. Most instances are Extract Method (976 occurrences) and Move Method operations (227 occurrences). These operations are detected in 450 commits from 166 projects, such as *infinispan/infinispan* (a tool for storing, managing, and processing data)[‡] and *gradle/gradle* (a build automation tool).[§] Figure 10 shows the distribution of the number of selected commits per project. As we can observe, the median is two commits, while the 90th percentile is about five commits. In the case of 78 projects (47%), there are only refactoring instances from a single commit. In other words, the oracle sample does not include the whole project's history.

4.1.3 | Detecting composite refactorings

We implement a set of scripts to detect the composite refactorings described in Section 3. Their input comprises a list of individual refactoring operations. Basically, these scripts operate by searching for clusters of refactoring operations R_1, R_2, \dots, R_n that can be replaced by a single

TABLE 1 Selected refactoring operations in the oracle

Operation	Projects	Commits	Occurrences	%
Extract Method	140	329	976	56.6
Move Method	53	73	227	13.2
Inline Method	48	64	127	7.4
Move and Rename Method	29	35	116	6.7
Extract and Move Method	29	35	114	6.6
Pull Up Method	24	28	74	4.3
Push Down Method	10	11	30	1.7
Pull Up Field	14	14	36	2.1
Push Down Field	11	11	25	1.4
All	166	450	1725	100

**FIGURE 10** Distribution of commits per project (oracle)**TABLE 2** Conditions to cluster two refactoring operations (r_1 and r_2) into a composite

Composite	Condition
Method Composition	$signature(r_1.target) = signature(r_2.target) \wedge type(r_1.target) = type(r_2.target) \wedge (r_1.refType, r_2.refType) \in \{extract, extract_move\}$
Method Decomposition	$signature(r_1.source) = signature(r_2.source) \wedge type(r_1.source) = type(r_2.source) \wedge (r_1.refType, r_2.refType) \in \{extract, extract_move\}$
Class Decomposition	$type(r_1.source) = type(r_2.source) \wedge (r_1.refType, r_2.refType) \in \{move, move_rename\}$
Composite Inline Method	$signature(r_1.source) = signature(r_2.source) \wedge type(r_1.source) = type(r_2.source) \wedge (r_1.refType, r_2.refType) \in \{inline\}$
Composite Pull Up Method	$signature(r_1.target) = signature(r_2.target) \wedge type(r_1.target) = type(r_2.target) \wedge (r_1.refType, r_2.refType) \in \{pull_up\}$
Composite Push Down Method	$signature(r_1.source) = signature(r_2.source) \wedge type(r_1.source) = type(r_2.source) \wedge (r_1.refType, r_2.refType) \in \{push_down\}$
Composite Pull Up Field	$name(r_1.target) = name(r_2.target) \wedge type(r_1.target) = type(r_2.target) \wedge (r_1.refType, r_2.refType) \in \{pull_up\}$
Composite Push Down Field	$name(r_1.source) = name(r_2.source) \wedge type(r_1.source) = type(r_2.source) \wedge (r_1.refType, r_2.refType) \in \{push_down\}$

composite refactoring CR. Therefore, we iterate over the list of refactorings detected in a system, grouping operations by considering the criteria described in Table 2.

For Method Decomposition, Class Decomposition, Composite Push Down Method, Composite Push Down Field, and Composite Inline Method (i.e., operations that break down code elements), we search for groups of refactorings that have as source the same code element. For Composite Pull Up Method, Composite Pull Up Field, and Method Composition), we look for refactorings that have as target the same code element.

Moreover, the source and target checking vary according to each composite refactoring. For composites at the level of methods, we verify the signature and the class. For example, for Method Composition, we group refactorings r_1 and r_2 into the same composite whenever the

TABLE 3 Inspected sample of composite refactorings (oracle)

Composite	Instances	Refactorings
Method Composition	4	48
Method Decomposition	4	22
Class Decomposition	4	52
Inline Method	4	9
Pull Up Method	4	11
Push Down Method	2	4
Composite Pull Up Field	4	10
Composite Push Down Field	2	4
All	28	160

signature of the target methods are the same (i.e., $signature(r1.target) = signature(r2.target)$) and the target methods are in the same class (i.e., $type(r1.target) = type(r2.target)$). We also check the respective refactoring types. In the case of Composite Pull Up Field and Composite Push Down Field, that is, composites at the level of fields, we verify the field's name and their respective class. Finally, for Class Decomposition, we group Move Method operations that originated from the same class (i.e., $type(r1.source) = type(r2.source)$).

It is also important to mention that our criteria for grouping refactoring operations do not include time constraints. Therefore, two or more refactorings can be part of the same composite, even though they were performed in distinct periods over the system's history. We made this decision motivated by two considerations. First, it is not trivial to set a threshold for the duration of the composites. Second, our main goal is to propose a catalog of composites, as well as to mine and analyze examples of these refactorings, even if they were performed in long time intervals.

After their execution, our scripts produce a list of composite refactorings, including a graph-based visualization and textual data. To validate the results, we manually inspected a sample of composite refactorings from the oracle. Specifically, we execute the following steps for each composite type:

1. We selected a random sample of four instances (of each composite).
2. For each selected instance:
 - We carefully analyzed the respective refactorings in the oracle, verifying whether the operations are correct. In other words, we check the refactoring type, source, and target, as well as basic information such as project name and commit.
 - In the last step, we verify if there are missing refactorings. In other words, we check if there are operations in the oracle that should be a part of the selected composite. For example, in `neo4j`, we detected a Method Composition that creates the method `createCountsTracker()` in class `CountsComputerTest`.⁴ For this case, we verify if there are extractions to the same target that were not properly detected by our scripts.

We manually inspected 28 composites, since for Composite Push Down Method and Composite Push Down Field, we detected only four instances in the oracle. Table 3 summarizes the results. The size of the selected composites ranges from 2 to 39 refactoring operations, covering 160 refactorings from the oracle. Overall, we did not identify errors by inspecting the sample of composite refactorings. In other words, we do not detect absent refactoring operations, that is, operations that were not clustered correctly by our scripts. The scripts and inspected sample are publicly available online (<https://github.com/alinebrito/composite-refactoring-catalog>).

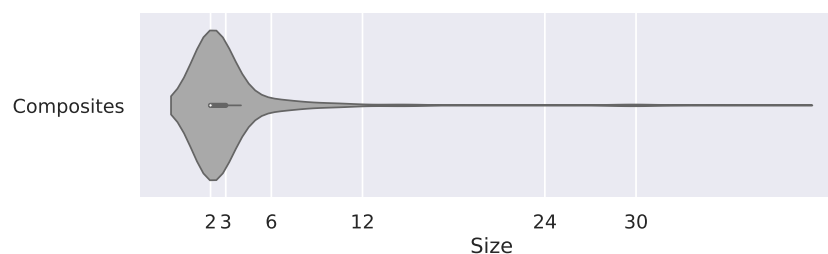
4.2 | Results

4.2.1 | (RQ1) What are the most common composite refactorings in the oracle?

Among 1725 single refactoring operations, an impressive number of 1043 (60.5%) are part of composite refactorings, as presented in Table 4. For example, 537 Extract Method or Extract and Move Method are part of Method Composition instances, which is the most frequent case. There are also significant rates of Method Decomposition (125 occurrences, 34.1%) and Class Decomposition (55 occurrences, 15%). However, composite refactorings in the inheritance hierarchy are infrequent. For example, there are only 15 composite refactorings formed by Push Down Method and Pull Up Method operations (4.1%). Also, there are a few occurrences of composites at the field level.

TABLE 4 Frequency of composite refactorings (oracle)

Name	Projects	Composites		
		Operations	Occurrences	%
Method Composition	37	537	142	38.8
Method Decomposition	37	295	125	34.1
Class Decomposition	37	277	55	15.0
Composite Inline Method	11	48	21	5.7
Composite Pull Up Method	7	33	13	3.6
Composite Push Down Method	2	4	2	0.6
Composite Pull Up Field	4	15	6	1.6
Composite Push Down Field	1	4	2	0.6
All	81	1043	366	100

**FIGURE 11** Distribution of the size of composite refactorings (oracle)

Overall, we detected 366 composite refactorings over 81 distinct projects. In 39 projects (48.1%), there is only a single composite instance. We identify most cases in a project called `Robovm`—127 instances grouping 389 single refactoring operations. Interestingly, this project also includes the largest composite refactoring instance, involving the composition of a method `has(...)`, which was created as the result of 30 Extract Method operations.[#] The new method contains only a single line of code:

```
public boolean has (CFString key) {
    return data.containsKey (key) ;
}
```

Therefore, this particular case of Method Composition was performed to remove code duplication (in this case, represented by a single line of code). It is worth mentioning that in the original oracle, this information was diluted over 30 individual and disconnected refactoring operations. By contrast, in our oracle view, they are represented by a single composite refactoring.

Summary of RQ1: Out of 1725 single refactoring operations, approximately 60% are part of composite refactorings. We detected the instances of composite refactoring in 81 projects. The most recurring cases are Method Composition (142 occurrences, 38.8%), Method Decomposition (125, 34.1%), and Class Decomposition (55, 15%).

4.2.2 | (RQ2) What are the characteristics of composite refactorings in the oracle?

We also investigate the main characteristics of the composite refactorings detected in RQ1 in terms of size and scope. Regarding their size, that is, the number of refactoring operations, most instances are small, as expected. As we can observe in Figure 11, about 84% of the detected composite refactorings have up to three refactoring operations (308 occurrences). The values range from 2 to 39 refactoring operations per composite.

Next, we detail the results for the most important composites:

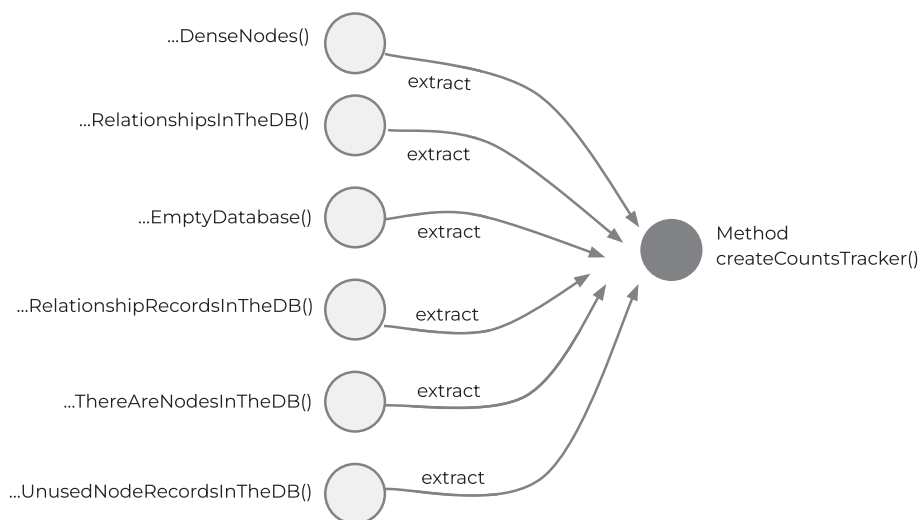


FIGURE 12 Example of method composition from Neo4j (oracle)

Class Decomposition. Among the 55 instances of Class Decomposition, 61.8% refers to classes losing up to two methods (29 occurrences) or three methods (5 occurrences). However, this category includes one of the largest composite refactorings in the oracle, where a developer from Graphhopper decomposed a class by moving 39 methods.^{||} Interestingly, in the commit message, the developer added a brief description regarding the motivation, which is related to a well-known design principle (use composition instead of inheritance³³):

“Refactoring of [class name]: use composition instead of inheritance”

Method Composition. The size of the 142 instances of Method Composition varies from 2 to 31 operations, with a median of two operations per composite. Furthermore, most Method Composition instances are *intra*class (121 occurrences, 85%), that is, the source methods are located in the same class of the target method. Figure 12 shows an example from Neo4j, where a developer extracted a method called `createCountsTracker()` from six methods.^{**} All refactorings happened in the scope of the same class `CountsComputerTest`. However, in the original oracle, these refactorings are reported as six distinct and unrelated operations. Finally, in 10 cases (7%), the composites are *inter*class, that is, developers compose methods by “merging” pieces of code coming from distinct classes. The remaining are *mixed* Method Decomposition, including the two categories.

Method Decomposition. About 95% of the instances of Method Composition (119 occurrences), which were detected in 37 projects, have up to three operations. The values range from 2 to 15 operations, distributed among *intra*class cases (91%), *inter*class cases (3%), and *mixed* ones (6%).

Other cases. In the oracle, there are only 127 occurrences of Inline Method. Consequently, we also found a few cases of composite inlines (21 instances, 5.7%), including at most four operations. The same applies to composite refactorings over inheritance hierarchies: Composite Pull Up Method (13 instances, 3.6%), Composite Push Down Method (2 instances, 0.6%), Composite Pull Up Field (6 instances, 1.6%), and Composite Push Down Field (2 instances, 0.6%).

Summary of RQ2: Most composite refactorings are small, including up to three operations. However, we also detect large instances, for example, 30 Extract Method operations to compose a single method. Regarding the scope of the operations, most Method Composition and Method Decomposition are *intra*class. In other words, developers usually extract multiple methods to the current classes.

5 | COMPOSITE REFACTORING IN THE WILD

In the first study, we look for composites in a well-known oracle. However, the refactoring instances selected for this oracle do not cover the complete history of each project. In other words, the oracle used in Section 4 only contains selected refactoring instances. Therefore, we might have missed operations in the reported composite refactorings simply because they were not selected for inclusion in the oracle. To tackle this issue, we decided to perform a complementary study, in which we search for composite refactorings in the complete history of ten popular GitHub-based projects.^{††}

5.1 | Study design

5.1.1 | Research questions assessment

As in the study described in Section 4, we propose two research questions:

(RQ3) *What are the most common composite refactorings in the wild?* Similarly to RQ1, we assess the frequency of each composite refactoring, but now in 10 popular GitHub projects.

(RQ4) *What are the characteristics of composite refactorings in the wild?* Similarly to RQ2, the rationale of this research question is to shed light on the main characteristics of composite refactorings while considering the complete development history of 10 projects.

5.1.2 | Dataset

To answer the proposed research questions, we relied on a set of real-world and popular projects. Specifically, we selected the top-10 Java projects on GitHub, ordered by their number of stars. We adopted this criterion because stars is a relevant metric to identify popular repositories.^{34,35} Moreover, in our sample, we only include projects that are software systems. For example, despite having a high number of stars, we did not include *kdn251/interviews* (a guide for interviews)^{††} and *iluwatar/java-design-patterns* (a set of code samples).^{§§} Table 5 describes the selected projects, including basic information, such as number of stars, commits, contributors, and short descriptions. The selected projects are from distinct domain areas, including web frameworks and animation libraries.

5.1.3 | Detecting composite refactorings

To detect composite refactorings, we need first to identify single refactoring operations. For this purpose, we used RefDiff, a well-known multi-language refactoring tool.^{7,8} As usual, in git-based mining tools, RefDiff detects refactorings by comparing a commit with its parent commit. To facilitate the usage of the tool, we first implemented a set of scripts that automate tasks such as downloading GitHub projects and retrieving the list of commits from the default branch. The scripts then rely on RefDiff to detect single refactoring operations. They also automatically exclude refactorings in noncore packages, such as “test” and “sample”. The final step concerns the detection of the composites defined in our catalog, using the scripts described in Section 4.

5.2 | Results

5.2.1 | (RQ3) What are the most common composite refactorings in the wild?

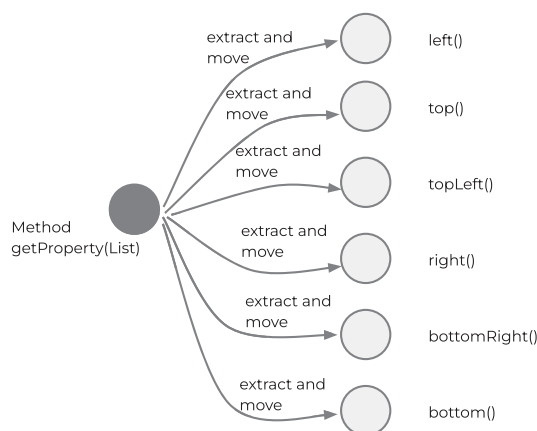
As presented in Table 6, we identify 2886 occurrences of composite refactorings. Most cases refer to Class Decomposition (957 occurrences, 33.2%); that is, 957 classes and interfaces have lost multiple methods. The values range from 8 classes in Lottie Android to 280 classes in Elasticsearch.

TABLE 5 Selected Java projects

Project	Stars	Comm.	Contr.	Short description
Spring Boot	56,717	33,692	831	Support framework
Elasticsearch	56,081	60,227	1651	Analytics engine
RxJava	45,055	5921	278	Event-based library
Spring Framework	43,943	22,728	551	Support framework
Google Guava	42,045	5609	265	Core Java libraries
Square Retrofit	38,539	1902	158	HTTP client
Apache Dubbo	35,968	4848	349	RPC framework
MPAndroidChart	33,811	2070	69	Chart library
Lottie Android	31,612	1321	106	Rendering library
Glide	31,578	2592	131	Image library

TABLE 6 Frequency of composite refactorings (in the wild)

Name	Wild		Oracle	
	Occur.	%	Occur.	%
Class Decomposition	957	33.2	55	15.0
Method Decomposition	683	23.7	125	34.1
Method Composition	582	20.2	142	38.8
Composite Pull Up Method	450	15.6	13	3.6
Composite Inline Method	129	4.5	21	5.7
Composite Push Down Method	85	2.8	2	0.6
Composite Pull Up Field	—	—	6	1.6
Composite Push Down Field	—	—	2	0.6
All	2886	100	366	100

**FIGURE 13** Example of method decomposition in Elasticsearch

Moreover, about 32% of the composites are from Elasticsearch, a popular search engine.¹¹ In this project, we detect 921 composites grouping 3310 single refactoring operations. Among them, most cases refer to Class Decomposition (280 occurrences, 30.4%).

There is also a significant number of Method Decomposition (683 occurrences, 23.7%), such as in the example of Figure 13. In this case, method `getProperty(List)` lost multiple pieces of code, after a developer performed six Extract and Move operations in a single commit.

Interestingly, all extracted methods were moved to the same class `GeoBoundingBox`. In the commit description,^{##} the maintainer points out the intention to centralize related logic:

“A lot of this logic can be centralized instead of having separated efforts to do the same things”

Summary of RQ3: In our extended dataset, the most common composite refactorings are Class Decomposition (957 occurrences, 33.2%), Method Decomposition (683 occurrences, 23.7%), and Method Composition (582 occurrences, 20.2%). There are also a few occurrences of composite refactorings related to inheritance, that is, Composite Pull Up Method and Composite Push Down Method.

Comparison with the oracle results (RQ1): In Table 6, we also report the results obtained with the oracle sample, aiming to facilitate comparison. As we can notice, the frequency of composites follows a similar tendency, that is, the top three cases are exactly the same: Class Decomposition, Method Decomposition, and Method Composition. However, in the oracle, the order is the reverse (e.g., Method Composition is the most frequent composite).

5.2.2 | (RQ4) What are the characteristics of composite refactorings in the wild?

Regarding their size—as measured by the number of single refactorings in each composite—most instances in the extended dataset are also small. Figure 14 presents the size distribution per project, after removing outliers, since they tend to distort the plot's aspect. In all projects, the median is two or three operations. However, there are also large composites; for example, the largest case includes dozens of operations, in which several methods were moved from a single class.^{III} In the following paragraphs, we detail the characteristics and give examples of each composite refactoring.

Class Decomposition. Figure 15 summarizes the size results of Class Decomposition. As we can notice, most composites of this type are small. About 62% of the cases involve up to three operations, such as in the example in Figure 16. In this example, a class of Lottie Android lost three methods in two commits. However, there are also large instances. For example, in Google Guava one developer moved each method from class `EmptyImmutableMap` to a distinct one, that is, he performed a composite refactoring composed of ten operations.***

Method Decomposition. As in the study described in Section 4, we also separate the composites into *intra*class (i.e., extractions to the same class of the fragmented method), *inter*class (i.e., when the extracted methods are moved to distinct classes), and *mixed* (i.e., both cases), as shown in Table 7. As we can observe, most extractions are in the *intra*class category (317 occurrences, 46%). However, another significant part of the

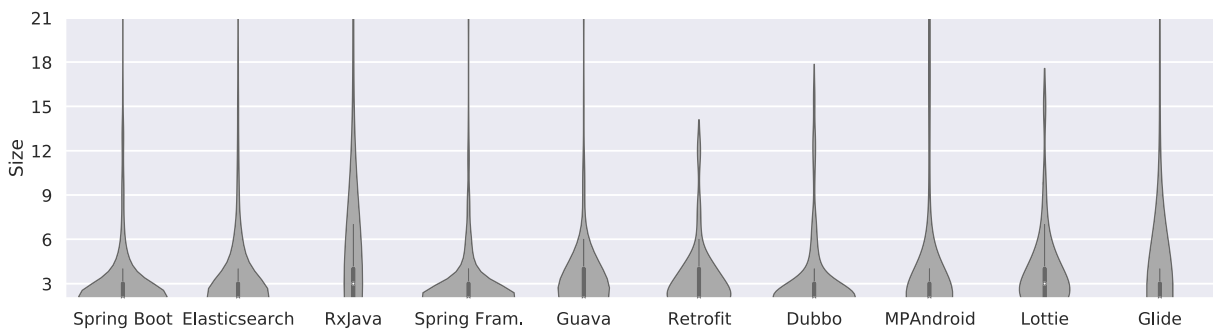


FIGURE 14 Distribution of the size of composite refactorings per project

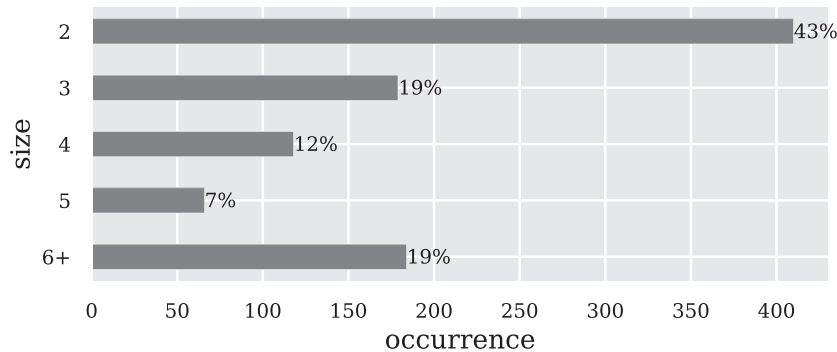


FIGURE 15 Number of operations by composite refactoring (class decomposition)

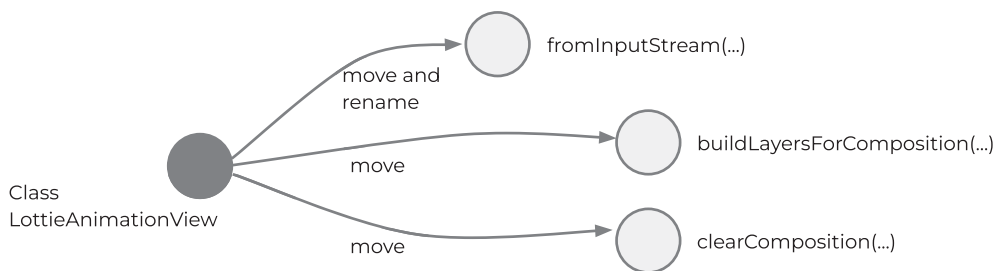



FIGURE 16 Example of class decomposition in Lottie Android

TABLE 7 Characteristics of method decomposition (in the wild)

Project	Occur.	Intraclass		Interclass		Mixed	
		Occur.	%	Occur.	%	Occur.	%
Spring Boot	148	107	72	21	14	20	14
Elasticsearch	234	75	32	118	50	41	18
RxJava	5	2	40	3	60	0	0
Spring Framework	152	77	51	38	25	37	24
Guava	10	3	30	5	50	2	20
Retrofit	6	2	33	2	33	2	33
Dubbo	51	22	43	20	39	9	18
MPAndroidChart	35	5	14	25	71	5	14
Lottie Android	14	5	36	5	36	4	29
Guide	28	19	68	1	4	8	29
All	683	317	46	238	35	128	19

**FIGURE 17** Number of operations by composite refactoring (method decomposition)**TABLE 8** Characteristics of method composition (in the wild)

Project	Occur.	Intraclass		Interclass		Mixed	
		Occur.	%	Occur.	%	Occur.	%
Spring Boot	79	34	43	42	53	3	4
Elasticsearch	219	49	22	161	74	9	4
RxJava	8	2	25	6	75	0	0
Spring Framework	151	67	44	66	44	18	12
Guava	17	5	29	10	59	2	12
Retrofit	5	1	20	4	80	0	0
Dubbo	39	20	51	15	38	4	10
MPAndroidChart	30	5	17	22	73	3	10
Lottie Android	12	1	8	10	83	1	8
Guide	22	13	59	9	41	0	0
All	582	197	34	345	59	40	7

results is *interclass* (238 occurrences, 35%); that is, all extracted methods are kept in the current class. We also investigate the number of extractions, that is, the size of the Method Decomposition instances. As presented in Figure 17, most cases refer to methods decomposed using two Extract operations (527 occurrences, 77%) or three operations (108 occurrences, 16%).

Method Composition. Among 582 instances of this composite refactoring, frequently, the extracted code is moved to distinct classes (345 occurrences, 59%), that is, they are *interclass*, as shown in Table 8. Figure 18 shows the results considering the size: as we can observe, most




FIGURE 18 Number of operations by composite refactoring (method composition)




FIGURE 19 Number of operations by composite refactoring (composite pull up method)




FIGURE 20 Size of composite refactorings (composite push down method)

cases involve up to three operations (467 occurrences, 80%). Dubbo includes an outlier, in which a developer extracted a utility method called `isEmptyMap(Map)` from seven other methods.^{†††} The extracted method has the following code:

```
public static boolean isEmptyMap (Map map) {
    return map == null || map.size() == 0;
}
```

Composite Pull Up Method. In this category, the number of operations follows the same tendency detected in RQ2, or example, most cases comprise two (311 occurrences, 69%) or three operations (65 occurrences, 15%), as reported in Figure 19. However, 8% of the occurrences have five or more operations. Spring Boot includes an example, in which a developer moved method `matches(...)` from five subclasses to the super-class `SpringBootCondition`.^{†††} In the commit description, the developer mentioned his intention, which relates to the improvement of the inheritance hierarchy:

“Create common [name class] base class... This removes the need for [class name] and simplifies many of the existing condition implementations.”

Composite Push Down Method. Figure 20 presents the results regarding the size of this type of composite refactoring. Overall, most cases comprise operations to move a method to at most three subclasses (82 occurrences, 97%).

Composite Inline Method. Regarding the number of affected elements, most Composite Inline operations involve two or three operations (109 occurrences, 85%), as presented in Figure 21. Elasticsearch includes a large instance, in which a developer removed method `cast(Input, Output)` by performing 23 Inline Method operations.^{§§§} In the commit description, the developer explained his motivation in the following way:




FIGURE 21 Size of composite refactorings (composite inline method)




FIGURE 22 Distribution of the age of composite refactorings performed over multiple commits (wild, 448 instances)

“Remove [functionality name] from [class name] as mutable state... this is no longer necessary as each cast is only used directly in the semantic pass after its creation...”

Age of Composites. In the oracle study (Section 4), we detect a single composite performed over multiple commits. However, in the wild study, a significant part of the composites is performed over time (448 instances, 15.5%). In these cases, we also assess age by computing the number of days between the most recent and the oldest commit in a composite. Figure 22 shows the distribution of the results. As we can notice, there are composites performed in a single day (3.3%, 15 composites), and also, there are composites performed over months. Among the 448 composite instances, 75% are performed up to 468 days (about 15 months), with a median age of 186 days (approximately 6 months). The 90th percentile is 835 days. However, it is difficult to generalize these results. For example, open-source projects are subjected to multiple periods of inactivity.³⁶

Summary of RQ4: In our extended dataset, most composite refactorings are also small, that is, they are formed by two or three operations (2258 composites, 78%). Regarding the scope of the operations, most Method Decomposition are *intra*class (46%), while most Method Composition are *inter*class (59%). In other words, when decomposing methods, developers usually extract them to their current classes. In contrast, when removing code duplication, developers frequently extract methods from multiple classes.

Comparison with the oracle results (RQ2): In the study described in Section 4, we also investigate composite characteristics. Regarding the size, the results are similar. For example, most composites have up to three operations (84% in the oracle vs. 78% in the extended dataset). The notable difference between both datasets refers to composites over multiple commits. The oracle only contains selected refactoring instances, that is, it does not cover the whole projects' history. Due to this fact, we identified a single composite over time, that is, a composite performed over more than one commit. In contrast, in RQ4, we detect 448 composites spread over two or more commits (15.5%). Finally, regarding the refactorings' scope, the results also follow a similar tendency. For example, most Method Decomposition operations are *intra*class in both samples. However, in the oracle, there is a higher frequency of *intra*class operations (91% in the oracle vs. 46% in the extended dataset).

6 | DISCUSSION AND IMPLICATIONS

In this paper, we proposed a catalog of eight composite refactorings, that is, refactoring operations composed of simple code transformations. Three of these refactorings—Class Decomposition, Method Decomposition, and Method Composition—are new, in the sense that they are not

documented in Fowler's catalog. The other refactorings—Pull Up Method, Push Down Method, Pull Up Field, Push Down Field, and Inline Method—are also described in Fowler's catalog. However, we decided to include them in our catalog for two key reasons: (a) they imply the realization of multiple source code transformations that affect multiple program elements and (b) they are not properly detected by refactoring detection tools, such as RefactoringMiner⁴ and RefDiff.⁷ However, to avoid potential conflicts, for these instances, we added the prefix “Composite” in their names. Regarding their popularity, the three new composites—Class Decomposition, Method Decomposition, and Method Composition—represent about 77% of the results in the wild study (Section 5). In the oracle study, 88% of the instances refer to these new cases (Section 4). These values are highlighted in Table 6.

Essentially, the main contribution of our study is the catalog, the set of scripts to identify the described composite refactorings, and a new perspective of the well-known refactoring oracle proposed by Tsantalis and other researchers.^{4,11}

We claim that this contribution can have two practical implications. First, as usual, our catalog highlights the importance and existence of composite refactorings. In other words, a catalog is a fundamental artifact to promote and disseminate the usage of composite refactorings among software practitioners. In fact, our studies showed that developers rely on composite refactorings during maintenance tasks. Therefore, the catalog and oracle can contribute to increasing the usage and application of such refactorings.

As a second practical implication, we showed that composite refactorings are not properly identified by refactoring detection tools, such as RefactoringMiner⁴⁻⁶ and RefDiff.⁷⁻⁹ Typically, these tools detect the parts of composite refactorings as independent operations. For this reason, we decided to implement a set of scripts to detect the eight composite refactorings in our catalog. Consequently, we also claim the concept of composite refactoring can be used to improve the results of empirical software engineering studies on refactoring practices. Finally, our scripts and catalog can also help to improve the user experience provided by refactoring-aware code review tools,²⁵ by supporting the detection of refactorings at a higher abstraction level.

7 | THREATS TO VALIDITY

Generalization of results. We characterized composite refactorings in terms of size and location. Our findings are based on a relevant oracle of refactoring operations and 10 real-world Java systems hosted on GitHub. However, they—as common in empirical software engineering—cannot be generalized to other scenarios, such as closed software systems or other programming languages.

Catalog of composite refactorings. Our catalog includes eight composite refactorings, which describe a sequence of operations to compose or decompose source code elements. We acknowledge that the current version of our catalog is not complete and final. However, any catalog of refactorings can increase over time due to new insights, research, and development demands. For example, the first catalog proposed by Fowler has 68 refactoring operations.³ After 18 years, in the second edition of his book, he introduced 15 new refactorings.³⁰ We also followed the idea of Fowler's book,^{3,30} using a single and popular programming language to guide the documentation and to provide illustrative examples. As mentioned by the author, it is “*better to use a single language so they can get used to a consistent form of expression*”.³⁰ In fact, we plan to extend our study in the future, by including, for example, composite refactorings at the package level. We also intend to explore other programming languages and refactoring types.

Detection of single refactorings. Before detecting composite refactorings, we first need to identify single operations. In our first study—described in Section 4—we rely on a well-known refactoring oracle, in which single refactoring instances were validated by multiple authors or tools.^{4,11} Therefore, our results are based on a trustworthy sample. For the second study—described in Section 5—we rely on RefDiff⁷ to mine refactorings in ten popular projects. According to recent results, the precision of RefDiff is high, reaching 96.4% for Java.⁷ In this second study, we also cleaned up the dataset; for example, we remove packages that are not part of the core system (e.g., *test*, *docs*, *sample*), and we removed constructors since they are essentially initialization structures. Finally, as natural during software evolution, commits can include temporary or unintentional operations, such as reverted commits due to test fails and experimental code. To mitigate this threat, we focus only on the main branch evolution.

Detection of composite refactorings. Regarding composite detection, we implement a set of scripts, as described in Section 4.1.3. The input comprises a list of single refactoring operations, including details such as path, refactoring type, and entities names. A possible threat is the possibility of errors in the implementation of our tool and parsers. For the oracle analysis, we extract this information from textual data. We also rely on well-known Python libraries to mitigate this threat, for example, retrieving the data by regex expression. Also, we inspected a sample of 28 composite refactorings to check the results (see details in Section 4.1.3), when we did not identify any error in the process of clustering refactoring operations as composites. Our verification included 160 single refactoring operations from the oracle created and curated by Tsantalis et al.^{4,5,11} Finally, we are making publicly available the datasets and scripts used to detect composite refactorings.

8 | RELATED WORK

We organized related work in three subsections: (a) field studies regarding sets of related refactoring operations, (b) studies about catalogs, and (c) other studies on refactoring.

8.1 | Batch and composite refactorings

Refactoring was already studied in scenarios such as code review,^{19,25,37,38} code understanding,^{32,39–41} and education.^{42,43} However, these studies do not propose catalogs of refactorings operations to improve software practices. Most of them also focus on single refactoring operations.

There are two central types of studies regarding groups of related refactoring operations, studies on *batch* refactorings, and studies on *composite* refactorings, which is the concept we explored in this paper. Batch refactorings refer to a set of single refactoring operations, which are then grouped considering criteria such as time,^{44,45} version system,⁴⁶ and developers.^{46,47} As mentioned by Cedrim et. al.,⁴⁶ “*the way the batches are synthesized is open-ended, that is, different developers can have different views of how to create a batch*”. Similarly, composites are defined as sequences of atomic refactoring operations.^{4,10,48} This concept is explored in contexts like domain specific languages for describing refactoring⁴⁸ and code smells.^{10,49}

Sousa et. al.¹⁰ originally defined composite refactorings as “*two or more interrelated refactorings that affect one or more elements*”. The detection of composite refactoring relies on three distinct heuristics. The first heuristic—called *element-based*—combines single refactoring operations by the scope. The scope can be, for instance, a single class. For example, the authors show a composite refactoring from this category, which includes the movement of attributes, movement of methods, and extract superclass operations. The *commit-based* heuristic links refactoring operations performed in a single commit. Finally, the third heuristic—named *range-based*—connects refactorings by location (e.g., if a refactoring crosscuts two classes named C_1 and C_2 , both are part of the location). As a consequence, an instance of a composite can include mixed operations at distinct levels, that is, classes, attributes, and methods. In summary, the study considers some criteria to cluster composites, also reusing previous heuristics.^{46,47} However, they do not introduce and document a catalog of composite refactorings (as we do in this paper), and a significant part of the study investigates the relevance of composites for removing code smells. Although we are reusing the definition, in this paper, we explore another perspective, that is, our key goal is to propose and document a catalog of composite refactorings. Moreover, we also show the importance of composite refactorings by mining and characterizing their occurrence in two datasets: a sample with hundreds of confirmed single refactoring operations and the history of ten well-known open-source projects.

There are also studies focusing on subcategories of composite refactorings. For example, “incomplete composites”,⁵⁰ that is, when the composite refactoring “*is not able to entirely remove a smelly structure*”.

8.2 | Catalog of refactorings

Recently, Bibiano et. al.⁴⁹ investigated “complete composites”, that is, sets of refactoring operations that remove the whole occurrence of four code smell types. The study includes 618 complete composites formed by well-known refactoring operations. Differently from our study, the identification of composite refactorings relies on a range-based heuristic defined in previous studies,¹⁰ which groups refactorings affecting the same location. The authors also present a catalog of complete composites to remove code smells. This catalog includes five complete composites, which are sequences of Move Method or Extract Method operations. For example, their catalog focuses on the removal of Long Method (i.e., large and complex methods) and Feature Envy (i.e., a method that uses several methods from a distinct class). For each instance, the authors discuss side-effects, that is, when a composite removes a target code smell but introduces other ones. Among the five complete composites from their catalog, three instances refer to extract operations to remove long methods. However, in the first case, the extraction contributes to introducing a Feature Envy. The second one does not reduce the method's size, that is, it is necessary to perform new extract operations to remove the smell. Finally, the third instance introduces a long parameter list before the extraction. In our study, we propose a catalog of eight types of composite refactoring, which are formed by distinct refactoring types. We cluster refactoring operations by considering the source or target code elements. In other words, our scripts identify sequences of single refactoring operations to compose or decompose a source code element, regardless of the presence of a code smell. In fact, there are several reasons to refactor a given source code element, which do not necessarily involve smell removal.^{32,40}

Tsantalis et. al.⁴ also present a brief discussion regards composites. The authors introduce a new version of RefactoringMiner, which detects Extract Class—also defined in Fowler's catalog.^{3,30} According to the authors, composite refactorings “*are composed of basic ones*”. Therefore, Extract Class matches this concept, since it comprises a set of Move Method and Move Field operations aiming to generate a new class. In our catalog, Class Decomposition can include a set of move operations to a new class or existing one. However, the focus refers to the decomposition of the source class.

Fowler proposes a popular and widely used catalog of refactoring operations.^{3,30} The recent version includes new composite refactorings, such as Inline Class and Collapse Hierarchy.^{###} In the case of Inline Class, we eliminate a class by moving all elements to distinct ones. Therefore, it is a subcategory of Class Decomposition. However, the current refactoring detection tools do not support this composite.^{4–9} In Collapse Hierarchy, we eliminate subclasses by moving all elements to the superclass. Therefore, Composite Pull Up Method can be a part of this operation. The current version of RefactoringMiner detects this refactoring.⁴ However, it is not properly explored in the literature. For example, the oracle used in this paper includes only a single instance of Collapse Hierarchy.

8.3 | Other studies on refactoring

In previous papers, we explored the reasons for refactorings performed over time.^{14,31} We analyzed characteristics such as time, refactoring types, and authorship. In particular, we relied on a graph-based abstraction—called refactoring graph—to mine refactoring operations performed over the history of ten GitHub projects. The insights from this study helped us to propose the catalog of composite refactorings described in this current paper.

Finally, we reinforce findings from a recent study that points out a significant rate of multiple extractions to decompose methods in a single commit.⁵¹ The authors show that Extract Method operations are frequently performed by developers, who create methods for distinct purposes, such as testing, validation, and setup. However, the study does not document a catalog of composite refactorings. Its goal is to characterize method extractions, for example, their content, size, and degree. Similarly, in our study, Method Decomposition and Method Composition—composites formed by Extract Method and Extract and Move Method operations—are among the top three most frequent composites. In the study in the wild, for example, we detected 1265 occurrences. Among them, 275 composites (21.7%) involving extractions are performed over multiple commits.

9 | CONCLUSION

We introduce a catalog of composite refactorings. According to our definition, a composite refactoring can be spread in multiple commits. Our catalog includes eight instances that describe sequences of operations that compose or decompose program elements: Method Composition, Method Decomposition, Class Decomposition, Composite Pull Up Method, Composite Push Down Method, Composite Pull Up Field, Composite Push Down Field, and Composite Inline Method.

In order to show that the proposed refactorings occur in real scenarios, we searched for occurrences of each instance in two datasets. First, we focus on a well-known refactoring oracle. In this first study, we identify that about 60% of the selected sample is part of a higher-level composite refactoring. Then, we mine the history of 10 popular GitHub projects, in which we detected 2886 instances of composite refactorings.

Future work might include an extension of the proposed catalog with other composite refactorings. We also plan to extend our mining study with projects implemented in other programming languages, such as C, JavaScript, and Go. The scripts are publicly available online (<https://github.com/alinebrito/composite-refactoring-catalog>).

ACKNOWLEDGMENTS

This research is supported by grants from FAPEMIG, CNPq, and CAPES.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in GitHub at <https://github.com/alinebrito/composite-refactoring-catalog>.

ORCID

Aline Brito  <https://orcid.org/0000-0002-7088-4935>

ENDNOTES

* github.com/spring-projects/spring-framework/commit/3642b0f3

† By construction, the discarded refactorings cannot be part of the composites included in our catalog. However, we acknowledge they can be part of future composites (in this case, therefore, we will need to update the current oracle).

‡ <https://github.com/infinispan/infinispan>

§ <https://github.com/gradle/gradle>

¶ https://github.com/alinebrito/composite-refactoring-catalog/blob/main/results/oracle/neo4j/neo4j/results/composition_extract_method/view/subgraph_atomic_4.md

<https://github.com/robovm/robovm/commit/bf5ee44b>

|| <https://github.com/graphhopper/graphhopper/commit/7f80425b>

** <https://github.com/neo4j/neo4j/commit/5fa74fbb>

†† Since operations at the field level are infrequent, and it is also unsupported by the current RefDiff tool version, we decide not to include them in this complementary study.

‡‡ <https://github.com/kdn251/interviews>

§§ <https://github.com/iluwatar/java-design-patterns>

¶¶ <https://github.com/elastic/elasticsearch>

<https://github.com/elastic/elasticsearch/commit/769650e0>
||| <https://github.com/ReactiveX/RxJava/commit/10325b90>
*** <https://github.com/google/guava/commit/d8f98873>
††† <https://github.com/apache/dubbo/commit/458a4504>
‡‡‡ <https://github.com/spring-projects/spring-boot/commit/840fdeb5>
§§§ <https://github.com/elastic/elasticsearch/commit/022d3d7d>
¶¶¶ <https://martinfowler.com/articles/refactoring-2nd-ed.html>
<https://refactoring.com/catalog>

REFERENCES

1. Sarkar S, Ramachandran S, Kumar GS, Iyengar MK, Rangarajan K, Sivagnanam S. Modularization of a large-scale business application: A case study. *IEEE Softw.* 2009;26(2):28-35.
2. Foote B, Yoder JW. *Big ball of mud*; 1997.
3. Fowler M. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley; 1999.
4. Tsantalis N, Ketkar A, Dig D. RefactoringMiner 2.0. *IEEE Transactions on Software Engineering (TSE)*; 2020.
5. Tsantalis N, Mansouri M, Eshkevari LM, Mazinanian D, Dig D. Accurate and efficient refactoring detection in commit history. In: 40th international conference on software engineering (icse); 2018:483-494.
6. Atwi H, Lin B, Tsantalis N, et al. PYREF: refactoring detection in python projects. In: 2021 IEEE 21st International Working Conference on Source Code Analysis and Manipulation (SCAM); 2021:1-6.
7. Silva D, da Silva JP, Santos G, Terra R, Valente MT. Refdiff 2.0: a multi-language refactoring detection tool. *IEEE Trans Softw Eng.* 2020;1(1):1-17.
8. Silva D, Valente MT. RefDiff: detecting refactorings in version histories. In: 14th international conference on mining software repositories (msr); 2017: 1-11.
9. Brito R, Valente MT. RefDiff4Go: detecting refactorings in Go. In: 14th brazilian symposium on software components, architectures, and reuse (sbcars); 2020:101-110.
10. Sousa L, Cedrim D, Garcia A, et al. Characterizing and identifying composite refactorings: concepts, heuristics and patterns. In: 17th international conference on mining software repositories (msr); 2020:186-197.
11. Tsantalis N, Mansouri M, Eshkevari L, Mazinanian D, Ketkar A. Refactoring oracle. <http://refactoring.encs.concordia.ca/oracle>, Online; accessed January 2022.
12. AlOmar E, Mkaouer M, Ouni A. *Knowledge Management in the Development of Data-Intensive Systems*: Auerbach Publications; 2021.
13. Merino L, Ghafari M, Anslow C, Nierstrasz O. A systematic literature review of software visualization evaluation. *J Syst Softw (JSS)*. 2018;144: 165-180.
14. Brito A, Hora A, Valente MT. Refactoring graphs: assessing refactoring over time. In: 27th international conference on software analysis, evolution and reengineering (saner); 2020:367-377.
15. Paixao M, Ucha A, Bibiano AC, et al. Behind the intents: an in-depth empirical study on software refactoring in modern code review. In: 17th international conference on mining software repositories (msr); 2020:125-136.
16. Bacchelli A, Bird C. Expectations, outcomes, and challenges of modern code review. In: 35th international conference on software engineering (icse); 2013:712-721.
17. Sadowski C, Söderberg E, Church L, Sipko M, Bacchelli A. Modern code review: a case study at Google. In: 40th international conference on software engineering: software engineering in practice (icse-seip); 2018:181-190.
18. Bacchelli A, Bird C. Expectations, outcomes, and challenges of modern code review. In: 35th international conference on software engineering (icse); 2013:712-721.
19. Ge X, Sarkar S, Witschey J, Murphy-Hill E. Refactoring-aware code review. In: Symposium on visual languages and human-centric computing (vl/hcc); 2017:71-79.
20. Avelino G, Passos L, Hora A, Valente MT. A novel approach for estimating truck factors. In: 24th international conference on program comprehension (icpc); 2016:1-10.
21. Hattori L, Lanza M. Mining the history of synchronous changes to refine code ownership. In: 6th international working conference on mining software repositories (msr); 2009:141-150.
22. Rahman F, Posnett D, Hindle A, Barr E, Devanbu P. BugCache for inspections: hit or miss? In: 19th international symposium on the foundations of software engineering (fse); 2011:322-331.
23. Kim S, Zimmermann T, Pan K, Whitehead EJJ. Automatic identification of bug-introducing changes. In: 21st international conference on automated software engineering (ase); 2006:81-90.
24. Shen B, Zhang W, Zhao H, Liang G, Jin Z, Wang Q. IntelliMerge: a refactoring-aware software merging technique. *Progr Lang.* 2019;3(170):170:1-170:28.
25. Brito R, Valente MT. RAID - Refactoring aware and intelligent diffs. In: 29th international conference on program comprehension (icpc); 2021: 265-275.
26. Higo Y, Hayashi S, Kusumoto S. On tracking Java methods with git mechanisms. *J Syst Softw (JSS)*. 2020;165.
27. Grund F, Chowdhury S, Bradley N, Hall B, Holmes R. CodeShovel: constructing method-level source code histories. In: 43rd international conference on software engineering: companion proceedings (icse); 2021:1510-1522.
28. Hora A, Silva D, Robbes R, Valente MT. Assessing the threat of untracked changes in software evolution. In: 40th international conference on software engineering (icse); 2018:1102-1113.
29. Spadini D, Aniche M, Bacchelli A. PyDriller: python framework for mining software repositories. In: 26th software engineering conference and symposium on the foundations of software engineering (fse); 2018:908-911.

30. Fowler M. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley; 2018.
31. Brito A, Hora A, Valente MT. Characterizing refactoring graphs in Java and JavaScript projects. *Empir Softw Eng*. 2021;26(6):1-43.
32. Silva D, Tsantalis N, Valente MT. Why we refactor? Confessions of GitHub contributors. In: 24th international symposium on the foundations of software engineering (fse); 2016:858-870.
33. Gamma E, Helm R, Johnson R, Vlissides J. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley; 1994.
34. Silva H, Valente MT. What's in a GitHub star? Understanding repository starring practices in a social coding platform. *J Syst Softw*. 2018;146:112-129.
35. Borges H, Hora A, Valente MT. Understanding the factors that impact the popularity of GitHub repositories. In: 32nd international conference on software maintenance and evolution (icsme); 2016:334-344.
36. Coelho J, Valente MT, Milen L, Silva LL. Is this GitHub project maintained? measuring the level of maintenance activity of open-source projects. *Inf Softw Technol*. 2020;1:1-35.
37. Hayashi S, Thangthumachit S, Saeki M. Rediffs: refactoring-aware difference viewer for Java. In: 20th working conference on reverse engineering (wcre); 2013:487-488.
38. Ge X, Sarkar S, Murphy-Hill E. Towards refactoring-aware code review. In: 7th international workshop on cooperative and human aspects of software engineering (chase) ACM; 2014:99-102.
39. Wang Y. What motivate software engineers to refactor source code? evidences from professional developers. In: International conference on software maintenance (icsm); 2009:413-416.
40. Pantiuchina J, Zampetti F, Scalabrino S, et al. Why developers refactor source code: a mining-based study. *ACM Trans Softw Eng Methodol*. 2020;37(4):1-32.
41. Golubev Y, Kurbatova Z, AI Omar EA, Bryksin T, Mkaouer MW. One thousand and one stories: a large-scale survey of software refactoring. In: 29th acm joint meeting on european software engineering conference and symposium on the foundations of software engineering (fse); 2021:1303-1313.
42. Lopez C, Alonso JM, Marticorena R, Maudes JM. Design of e-activities for the learning of code refactoring tasks. In: 2014 16th international symposium on computers in education (siete); 2014:35-40.
43. Hebig R, Ho-Quang T, Jolak R, et al. How do students experience and judge software comprehension techniques? In: 28th international conference on program comprehension (icpc); 2020:425-435.
44. Murphy-Hill E, Parnin C, Black AP. How we refactor, and how we know it. In: 31st international conference on software engineering (icse); 2009:287-297.
45. Murphy-Hill E, Parnin C, Black AP. How we refactor, and how we know it. *Trans Softw Eng*. 2012;38(1):5-18.
46. Cedrim D. Understanding and improving batch refactoring in software systems. *Ph.D. Thesis*: PUC-Rio; 2018.
47. Bibiano AC, Garcia EFDOA, Kalinowski M, et al. A quantitative study on characteristics and effect of batch refactoring on code smells. In: 13th international symposium on empirical software engineering and measurement (esem); 2019:1-11.
48. Li H, Thompson S. A domain-specific language for scripting refactorings in erlang. In: 15th fundamental approaches to software engineering (fase); 2012:501-515.
49. Bibiano AC, Assuno W, Coutinho D, et al. Look ahead! revealing complete composite refactorings and their smelliness effects. In: 37th international conference on software maintenance and evolution (icsme); 2021:1-12.
50. Bibiano A, Soares V, Coutinho D, et al. How does incomplete composite refactoring affect internal quality attributes. In: 28th international conference on program comprehension (icpc); 2020:149-159.
51. Hora A, Robbes R. Characteristics of method extractions in java: a large scale empirical study. *Empir Softw Eng*. 2020;25:1798-1833.

AUTHOR BIOGRAPHIES

Aline Brito is a PhD candidate in the Computer Science Department at the Federal University of Minas Gerais (UFMG), where she also received her Master's Degree in Computer Science. She received her Bachelor's Degree in Computer Engineering from the Pontifical Catholic University of Minas Gerais (PUC Minas). Brito also was a software developer for 5 years. Her research interests include software quality analysis, software maintenance and evolution, and software repository mining. Contact her at alinebrito@dcc.ufmg.br; alinebrito.com.

Andre Hora is a professor in the Computer Science Department at the Federal University of Minas Gerais (UFMG). His research interests include software evolution, software repository mining, and empirical software engineering. Hora received a PhD in Computer Science from the University of Lille. He was a Postdoctoral researcher at the ASERG/UFMG group during 2 years and a software developer at Inria/Lille during 1 year. Contact him at andrehora@dcc.ufmg.br; www.dcc.ufmg.br/~andrehora.

Marco Tulio Valente is an associate professor in the Computer Science Department at the Federal University of Minas Gerais (UFMG), where he also heads the Applied Software Engineering Research Group (ASERG). His research interests include software architecture and modularity, software maintenance and evolution, and software quality analysis. Valente received his PhD in Computer Science from the Federal University of Minas Gerais. He is a Researcher I-D of the Brazilian National Research Council (CNPq) and holds a Researcher from Minas Gerais State scholarship, from FAPEMIG. Contact him at mtov@dcc.ufmg.br; www.dcc.ufmg.br/~mtov.

How to cite this article: Brito A, Hora A, Tulio Valente M. Towards a catalog of composite refactorings. *J Softw Evol Proc*. 2023;e2530. doi:[10.1002/smr.2530](https://doi.org/10.1002/smr.2530)