UNIVERSITÉ DE MONTRÉAL

AUTOMATIC DETECTION AND CLASSIFICATION OF IDENTIFIERS

LALEH M. ESHKEVARI
DÉPARTEMENT DE GÉNIE INFORMATIQUE ET GÉNIE LOGICIEL
ÉCOLE POLYTECHNIQUE DE MONTRÉAL

UNIVERSITÉ DE MONTRÉAL

ÉCOLE POLYTECHNIQUE DE MONTRÉAL

Cette thèse intitulée :

AUTOMATIC DETECTION AND CLASSIFICATION OF IDENTIFIERS

présentée par : M. ESHKEVARI Laleh
en vue de l'obtention du diplôme de : Philosophiæ Doctor
a été dûment acceptée par le jury d'examen constitué de :

M. KHOMH Foutse, Ph.D., président
M. ANTONIOL Giuliano, Ph.D., membre et directeur de recherche
M. GUÉHÉNEUC Yann-Gaël, Doctorat, membre et codirecteur de recherche
M. MERLO Ettore, Ph.D., membre
Mme HILL Emily, Ph.D., membre externe

To Kianoush and Kevin

ACKNOWLEDGEMENTS

RÉSUMÉ

Le lexique du code source joue un rôle primordial dans la maintenabilité des logiciels. Un lexique pauvre peut induire à une mauvaise compréhension du programme et à l'augmentation des erreurs du logiciel. Il est donc important que les développeurs maintiennent le lexique de leur code source en renommant les identifiants afin qu'ils reflètent les concepts qu'ils expriment. Dans cette thèse, nous étudions le lexique et proposons une approche pour détecter et classifier les renommages des identifiants dans le code source.

La détection des renommages est basée sur la combinaison de deux techniques : la différenciation des codes sources et l'analyse de flux de données. Tandis que le classificateur de renommage utilise une base de données ontologique et un analyseur syntaxique du langage naturel pour classer les renommages selon la taxonomie que nous avons défini. Afin d'évaluer l'exactitude et l'exhaustivité du détecteur de renommage, nous avons réalisé une étude empirique sur l'historique de cinq programmes Java open-source. Les résultats de cette étude rapportent une précision de 88% et un rappel 92%. Nous avons également mené une étude exploratoire qui analyse et discute comment les identifiants sont renommés, selon la taxonomie proposée, dans les cinq programmes Java de l'étude précédente. Les résultats de cette étude exploratoire montrent qu'il existe des renommages dans chaque dimension de notre taxonomie.

Afin d'appliquer l'approche proposée aux programmes PHP, nous avons adapte notre détecteur de renommages pour prendre en compte les caractéristiques inhérentes à ces programmes. Une étude préliminaire effectuée sur trois programmes PHP montre que notre approche est applicable aux programmes PHP. Cependant, ces programmes ont des tendances de renommages différentes de celles observées dans les programmes Java. Cette thèse propose deux résultats. Tout d'abord, la détection et la classification des renommages et un outil, qui peut être utilisé pour documenter les renommages. Les développeurs seront en mesure de, par exemple, rechercher des méthodes qui font partie de l'interface de programmation car celles-ci impactent les applications clientes. Ils pourront également identifier les incohérences entre le nom et la fonctionnalité d'une entité en cas de renommage dit risqué comme lors d'un renommage vers un antonyme. Deuxièmement, les résultats de nos études nous fournissent des leçons qui constituent une base de connaissance et de conseils pouvant aider les développeurs à éviter des renommages inappropriés ou inutiles et ainsi maintenir la cohérence du lexique de leur code source.

ABSTRACT

Source code lexicon plays a paramount role in software maintainability: poor lexicon can lead to poor comprehensibility and increase software fault-proneness. For this reason, developers should maintain their source code lexicon by renaming identifiers when they do not reflect the concepts that should express. In this thesis, we study lexicon and propose an approach to detect and classify identifier renamings in source code. The renaming detection is based on a combination of source code differencing and data flow analysis, while the renaming classifier uses a ontological database and a natural language parser to classify renamings according to a taxonomy we defined. We report a study—conducted on the evolution history of five open-source Java programs—aimed at evaluating the accuracy and completeness of the renaming detector. The study reports a precision of 88% and a recall 92%. In addition, we report an exploratory study investigating and discussing how identifiers are renamed in the five Java programs, according to our taxonomy. Moreover, we report the challenges and applicability of the proposed approach to PHP programs and report our preliminary results of renamings detection and classification for three programs. This thesis provides two outcomes. First, the renaming detection and classification approach and tool, which can be used for documenting renamings. Developers will be able to, for example, look up methods that are part of the public API (as they impact client applications), or look for inconsistencies between the name and the implementation of an entity that underwent a high risk renaming (*e.g.*, towards the opposite meaning). Second, pieces of actionable knowledge, based on our qualitative study of renamings, that provide advice on how to avoid some unnecessary renamings.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LISTE DES ANNEXES

CHAPTER 1   INTRODUCTION

Program comprehension is one of the important activities during software maintenance. Previous studies showed that majority of the time during maintenance is spent on understanding source code (von Mayrhauser *et al.*, 1997; Standish, 1984; Tiarks, 2011) . Researchers agree on the paramount role of source code lexicon in software comprehensibility and maintainability because, very often, documentation is either scarce or outdated. Hence, when performing change tasks, developers must rely on source code lexicon. Recent studies have related the quality of source code identifiers with the overall software quality (Takang *et al.*, 1996; Lawrie *et al.*, 2007, 2006b). Also, researchers have developed approaches to help developers using appropriate identifiers, consistent with requirements and other high-level artifacts (De Lucia *et al.*, 2011) as well as approaches to support program comprehension tasks by splitting and expanding identifiers composed of multiple terms, including abbreviations and acronyms (Corazza *et al.*, 2012; Enslen *et al.*, 2009; Lawrie et Binkley, 2011; Madani *et al.*, 2010; Guerrouj *et al.*, 2011).

As source code evolves, identifiers evolve too (Abebe *et al.*, 2009). Previous work (Lawrie *et al.*, 2006a; Malpohl *et al.*, 2000) investigated the evolution of the structure of identifiers and the presence and stability of domain terms (Haiduc et Marcus, 2008). Much as any other programming element, identifiers are added, deleted, or modified, *i.e.*, renamed. Renaming, the modification of an entity name, happens for a variety of reasons. Renaming may occur to improve program understanding, to better reflect the developers' programming style, to better convey domain or application concepts or when name of an entity is not (anymore) consistent with its functionality. To the best of our knowledge, there has been no study for understanding renamings, in particular why, how, and when developers rename identifiers.

The thesis of this dissertation is:

> Renaming is an activity that occurs with different frequency during the life cycle of a program. Renaming is inevitable and its purpose is to increase consistency. Detection and linguistic analysis of identifier renamings provides valuable insight on how, why, when developers rename identifiers. Tool supports, programming language, and naming convention are factors that impact renamings frequency.

We propose a methodology to automatically detect identifier renamings across different versions of a program and to automatically classify renamings according to a novel taxonomy. The taxonomy provide a view of identifier renamings across different dimensions: (i) what

kind of identifier was renamed (*e.g.*, class name), (ii) whether one or more terms composing the identifier were added/removed/changed (*e.g.*, a term is added when renaming `files` to `srcFiles`), and (iii) how terms were changed with respect to their semantics (*e.g.*, towards opposite meaning when renaming `disable` to `enable`) and grammar (*e.g.*, from adjective to noun when renaming `localDeclaration` to `location`). We use a lightweight file differencing tool to identify changed source code lines and map the lines and the declared entities across versions to identify candidate renamings. We apply data-flow analysis on programming entities and extract def-use for the entities participating in candidate renamings to further reduce false positives. For classification, uses WordNet [1] (Miller, 1995) and the Stanford Part-of-Speech Analyzer (Toutanova et Manning, 2000) to classify the detected renamings according to our taxonomy. Our methodology is implemented via a tool suite called REPENT. The tool is composed of two main components for the detection and classification of the identifier renamings. In the following section we report the result of an online survey we designed to seek developers of open-source/industrial projects opinion about renamings.

1.1 Developer Survey on Identifier Renamings

This section motivates the need for automatic renaming detection and classification. For this purpose, we designed an online survey to understand the importance of renaming, *i.e.*, to what extent developers of open-source/industrial projects perform identifier renaming, under what circumstances, and whether they believe that identifier renaming requires automatic documentation.

We invited 739 developers, via e-mail using a convenience sampling (Groves *et al.*, 2009), involving (i) original developers of the five Java programs that we study and (ii) other developers from the industry and open-source communities. 71 developers responded to the survey resulting in a response rate close to 10% as expected (Groves *et al.*, 2009). Although we profile survey participants based on their background, their identity is kept anonymous for confidentiality purposes.

We observe that renaming is tangled with many development activities—most of the participants perform renaming while performing other refactorings (90% of surveyed developers); changing or adding functionality (89% and 65%, respectively); understanding existing code (51%) or fixing bugs (42%). Sometime renaming is even performed apart from other development activities (17%). Renaming is an activity that 39% of participants perform from a few times per week to almost every day and 46% perform it a few times per month. Participants

---

1. `http://wordnet.princeton.edu`

mainly use automatic tools to perform renaming (72%), although 20% say that they rename manually and 8% do both. However, although tool support is available, 92% of the participants consider renaming not straightforward and only 24% think that in most cases renaming has no cost. Indeed, the largest fraction of the surveyed developers (67%) believes that the cost of renaming depends on the particular case and that it requires time and effort. For example, participants underline that renaming identifiers that belong to non-local scope may break backward compatibility, increase integration cost, or impair program understanding for those already familiar with the old name. Figures with detailed results are reported in Appendix.

In the following we summarize the results of the survey and we illustrate them with comments from the participants. We complement the survey output with examples that we collected from online discussions of the analyzed programs (issue reports, mailing lists, and commit notes).

How often do developers rename?   Renaming is an activity that participants perform from almost every day (21%), a few times per week (18%), a few times per month (46%), to once per month (14%). A developer commented: *"There's a balance to be struck: - identifiers are communication, and as the code is refactored it is critical that identifiers continue to correctly describe their purpose - changing identifiers tends to break APIs, and sometimes they're used for unintended purposes, over-frequent change is not good."*

Is renaming straightforward?   When we asked participants whether renaming has a cost, only 8% answered that renaming is straightforward. 24% of participants think that in most cases renaming has no cost, often due to the availability of automatic tool support. Indeed the majority of participants (72%) use automatic tool support to perform renaming, although 20% rename manually and 8% use a mix of both, *i.e.*, rename manually and automatically. 32% of participants believe that the cost of renaming depends on the particular case:

"Renaming identifiers that belong to non-local context (e.g., public or protected methods) has a potentially massive cost associated with breaking the interfaces between components. Otherwise it is typically a rather cheap and non-disruptive exercise that may have end benefit of more readable and consistent code. Another element of cost and risk is when the identifiers are being bound to at runtime only (e.g., when classes are loaded by name or methods are bound by name). It is not always easy to trace all such use cases in a large system." Indeed, renaming an entity that is part of a public API of a program has a higher cost as it breaks

backward compatibility and increases the integration cost of the program in client programs. 10% of participants believe that in most cases renaming has a cost, and finally 25% answer that renaming defiantly requires time and effort. Another example where renaming has a cost is when the team uses code reviews, as developers must schedule a code review and justify their decision. A developer indicated that code reviews impact the frequency of renaming *"because you appear negatively to the boss when asking for a review on a 'too minor improvement"'.* The cost of renaming also includes the cost of finding a proper name and assuring that the new name reflects the purpose of the entity in all scenarios that it is used. Quotes like *"I have the feeling that your method name is not good [..]"* for method `getBufferForWrite` in an Eclipse issue report (issue #332248) indicates that, indeed, developers spend time understanding the rationale behind names that are chosen by other teammates.

Already postponed a renaming? It also appears that, although necessary, some renamings are delayed. After discussing the difference between the term "delete" and "remove", an ArgoUML developer concluded that: *"[..] maybe I shall rename these after next release"* (issue #2938).

We asked participants to share reasons for which they recall having decided not to rename an entity. 52% recall the reason to be the potential impact on other systems. A developer explains: *"As a middleware developer, providing a stable API is paramount for clients. There are numerous cases where we would not rename a class or method despite an obviously better name being proposed, in order to minimize the cost of integrating new versions."* 35% recall that the renaming was too risky, *i.e.*, it might have introduced a bug—a developer recalls: *"I encountered a problem when my colleague wrote Java code which uses reflection. I avoided renaming some classes/methods which will be inspected by the reflection, since doing so can introduce unpredictable bugs."*

25% of participants answered that the high impact of the renaming on the system was the show-stopper and finally, 25% recall deciding not to rename because of the high effort required: *"I'm not touching poorly-worded APIs which are shared across multiple projects - the cost of the change does not justify it [..]."*

Participants also shared that the impact on other developers is sometimes decisive: *"If too many people in the company know a thing by name X it's sometimes better to keep it even when name Y is more descriptive."*

Other factors impacting the decision to undertake a renaming are insufficient domain knowledge (85% of participants), code ownership (79%), and close deadline (76%).

How can REPENT help in such a context? Detecting and classifying renamings with REPENT—for example generating parts of commit notes when renaming occurs—can be used by developers while backtracking bugs or understanding changes of program entities. REPENT allows developers to differentiate and thus document and retrieve all or only certain types of renamings—*e.g.*, renamings towards opposite meaning as they deserve more attention and can be flagged to make sure that they reflect the developer's intentions. The documentation of renamings is also useful as a starting point for documenting API changes in release notes. Last but not least REPENT can reduce the unnecessary cost of some renamings by informing developers about names that were already changed in the past.

We asked participants whether they consider useful automatically documenting renaming and 52% of them were positive. A developer elaborates: *"It depends on how this was implemented, but if it were field-level history, e.g., like svn records history for a file, then I'm all for it [..]";* *"Tracking changes to public api is imperative in large fast moving teams."*

1.2 Our Contributions

We apply REPENT on five open-source Java programs. Despite the renaming cost, risks, and implications for program understanding, only a small percentage of the renamings is actually documented—1% of the renamings. This novel observation explains why a high number of surveyed developers (52%) considered automatic documentation of renamings useful; one of the surveyed developers explains that *"if there were an easy way to look renaming up, it would potentially be informative when backtracking for problems, or even just trying to understand someone else's code. You can often learn a lot about what something does by looking at how other people disagree regarding what it does."* The developer echoes George Santayana: *"Those who cannot remember the past are condemned to repeat it"* (Santayana, 1905). We share their point of view and argue that documenting renamings is important to track changes in vocabulary and to create traceability links between entities over time, and to avoid using names that will be later changed.

The proposed taxonomy and thus renaming classification is language independent. However, renamings detection requires parsing and data flow analysis and thus is bound to parser. To investigate the applicability of the proposed approach on other programming languages, we perform a preliminary study investigating the occurrence of renamings in a dynamic type language. We chose PHP as it is a dynamically typed language and one of the most popular Web scripting language [2], accounting for more than 80% of existing web sites. We analyzed

---

2. https://en.wikipedia.org/wiki/PHP#Usage

one month of daily commits of three open-source PHP programs. The results are motivating in term of precision and recall for detection and classification of identifiers in PHP programs. However, we need to analyze more data to make a concrete conclusion.

The contribution of this thesis can be summarized as:

– An empirical study to understand the developers point of view on renamings. The results of the survey help us to know why developers rename identifiers. It provides evidence that developers believe renaming is costly and they avoid renamings when the associated cost is too high.
– A novel taxonomy to identify semantic changes of identifier renamings.
– Detection of identifier renamings to understand the frequency of renamings, and thus answer when identifiers renamed and what are the most renamed entities.
– Classification of identifier renamings along orthogonal dimensions of proposed taxonomy. Through the quantitative results of the study, we identify instances of semantic changes in renamings and thus we can answer how identifiers are renamed, while the qualitative results enable us to answer why developers rename identifiers.
– Adaption of our automatic detection for PHP programs. The results of this study enable us to evaluate the feasibility of proposed methodology for detection and classification of renamings in other programming language.

Our contribution published in the following journal and conference papers:

MSR 2011: *Taxonomy and preliminary study of Identifier renaming*
Laleh Eshkevari, Venera Arnaoudova, Massimiliano Di Penta, Rocco Oliveto, Yann-Gaël Guéhéneuc, and Giuliano Antoniol. An Exploratory Study of Identifier Renamings. In Proceedings of the Working Conference on Mining Software Repositories (MSR), 2011.

TSE 2014: *Extension of taxonomy, improving renaming detection, and empirical study*
Venera Arnaoudova, *Laleh Eshkevari*, Massimiliano Di Penta, Rocco Oliveto, Giuliano Antoniol, and Yann-Gaël Guéhéneuc. REPENT : Analyzing the Nature of Identifier Renamings". In: IEEE Transactions on Software Engineering (TSE), 40 (5), 2014, pp.502–532.

ICPC 2014 *Resolving include statements in PHP programs*
Laleh Eshkevari, James R. Cordy, Massimiliano Di Penta, and Giuliano Antoniol. Identifying and locating interference issues in PHP applications: the case of WordPress. In Proceedings of 22nd International Conference on Program Comprehension (ICPC), 2014.

1.3   Organization of the Thesis

The organization of this dissertation is presented below:

Chapter 2 discusses the related literature on role of identifiers in software quality, analysis of changes and refactoring. **Chapter 3** describes the taxonomy we defined to classify identifier renamings. **Chapter 4** describes the proposed approach to detect identifier renamings. **Chapters 5** reports and discusses the results of renamings detection, accuracy and completeness of our approach. **Chapter 6** describes our approach for classifying renamings according to our taxonomy. **Chapter 7** reports and discusses the results of our empirical study to classify renamings. **Chapter 8** reports our challenges and preliminary results on applicability of our automated approach on PHP programs. **Chapter 9** concludes the dissertation by summarizing our findings and discussing future work.

CHAPTER 2   LITERATURE REVIEW

Following our thesis, we present the related work on: (1) the role of identifiers in software quality, and (2) approaches for detecting refactoring.

2.1   Role of Identifiers in Software Quality

There is quite a consensus among researchers (Caprile et Tonella, 2000; Deissenbock et Pizka, 2005; Lawrie *et al.*, 2006a; Enslen *et al.*, 2009) on the role played by identifiers on program comprehension, maintainability, and quality in general. In particular, researchers studied the usefulness of identifiers to recover traceability links (Antoniol *et al.*, 2002; Maletic *et al.*, 2005), measure conceptual cohesion and coupling (Marcus *et al.*, 2008; Poshyvanyk et Marcus, 2006), and, in general, high quality identifiers are considered an asset for source code understandability and maintainability (Takang *et al.*, 1996; Lawrie *et al.*, 2007, 2006b).

As suggested by (Deissenbock et Pizka, 2005), identifiers should be consistent and concise. Unfortunately, verifying consistency and conciseness is a difficult task and thus approaches have been developed to detect consistency and conciseness violations by identifying usages of synonyms and holonyms (Lawrie *et al.*, 2006a). We share the concern expressed in previous studies on identifier quality as a support for various software engineering tasks. However, we are focusing our study on identifier renaming based on a newly proposed taxonomy. We concur with Lawrie *et al.* (2006a) that synonyms can indeed affect consistency. However, we also believe that renaming towards synonyms or towards other semantically-related terms—such as hypernyms, hyponyms, and antonyms—should be investigated as they likely point to program understanding issues.

2.2   Analysis of Changes and Refactorings

Several authors have proposed automated approaches to detect different kinds of refactorings. Renamings is one type of refactoring and thus we share with the following works detection of renaming refacotrings.

At design level, Xing et Stroulia (2006) propose UMLDIFF to detect refactorings. UMLD-IFF works with class diagrams; it inputs two class diagrams and it produces as output an XML design differencing file. By querying such a XML file, it would be possible to detect simple (*e.g.*, rename class/method/field, pull-up/push-down method/field) and composite refactoring actions (*e.g.*, replace inheritance with delegation). For evaluating the detection

technique, the Xing *et al.* conducted a study on 11 versions of HTMLUnit and 31 versions of JFreeChart. The results of the study showed that all refactorings documented by the program developers as well as some that were not documented were identified. UMLDIFF analyses the design artefacts. As its purpose and objectives differs from our proposal, a comparison with our approach would not be appropriate.

Demeyer *et al.* (2000) detect object-oriented refactorings based on a set of heuristics defined in terms of changes of object-oriented metrics measuring two successive software versions of Smalltalk programs. They validated the approach on several successive versions of three cases studies implemented in Smalltalk.

Dig *et al.* (2006) propose REFACTORING CRAWLER for detecting sequences of refactorings between consecutive versions of Java programs. REFACTORING CRAWLER identifies seven types of refactoring. Among others, they detect—as does REPENT—-package, class, and method renaming. The detection algorithm consists of a fast syntactic analyzer followed by a more computationally intensive semantic analyzer. The syntactic analyzer finds similar text fragments between two versions of source code based on Shingle encoding (Broder, 1997) as candidates of refactorings. The semantic analyzer further filters the candidate pairs to reduce false positives.

Weissgerber et Diehl (2006) propose a signature-based approach to identify refactorings. The approach starts with collecting and pre-processing data from the version control system. Next, it identifies added and removed entities (classes, interfaces, methods, and fields) in each transaction. Those entities are then compared based on their signatures and potential refactorings are identified. The approach then ranks and filters potential refactorings based on the similarity of the entity body in the old and new version. To measure the similarity between the two versions, the approach first tests for string equality. Then if it fails, the approach uses the result of a token-based code clone detection algorithm, *i.e.*, CCFINDER (Kamiya *et al.*, 2002). Among the detected refactorings, the approach detects "Rename Method" and "Rename Class" refactorings, as REPENT does.

Prete *et al.* (2010) propose REF-FINDER as a way to detect atomic refactorings and then based on logic templates reconstruct more complex refactorings (such as extract method). REF-FINDER detects method renamings based on method body similarity.

It is important to point out that the approaches described above have been conceived to detect refactorings in general. They detect only a subset of the renamings detected by REPENT, and do not perform a classification of the detected renamings.

Malpohl *et al.* (2000) propose RENAMING DETECTOR for detecting identifier renamings. The

tool uses three main components: Parser, Symbol Analyzer, and Differencer. RENAMING DETECTOR analyzes each file for extracting identifier declarations and references. Next, it matches the declarations in two versions of a file. To increase accuracy, variable types and references are compared for matching the identifiers. Malpohl *et al.* evaluated the technique on two consecutive versions of the tool itself. They report a 100% precision rate for of the 77 analyzed file pairs. We share with Malpohl *et al.* the general idea as well as the use of data-flow analysis in the renaming detection process. However, our approach for the detection of renamings is substantially different. Specifically, it is a multi-stage approach, in which an initial filtering localizing changes based on differencing analysis is then followed by a data-flow analysis on candidate renamings, aimed at filtering out false positives. This allows us a better scalability, and hence the ability to analyze the evolution of large projects such as JBoss or Eclipse-JDT.

Neamtiu *et al.* (2005) propose an approach for understanding code evolution using AST matching. They analyze open source programs written in C and provide a release digest that summarized changes between two subsequent releases of a file. Renamings of types and variables are part of the changes detected by the authors. Those renamings correspond to our class, field, local variable, and parameter renamings. The approach does not handle method renamings, nor local variable and parameter renamings when the latter are in a renamed method because the approach is based on the hypothesis that in C functions are relatively stable over time. Thus, to detect local variable and parameter renamings, Neamtiu *et al.* compare the ASTs of two methods with the same name. Such an assumption would be unrealistic for Java programs. As our results show, method renamings represent 26% of the renamings for the 5 programs.

Fluri *et al.* (2007) propose a tree differencing algorithm, CHANGEDISTILLER, for extracting the changes from two consecutive versions of Java files. Renaming is a type of change that CHANGEDISTILLER can detect together with many others. The algorithm compares the ASTs of the files and computes the edit operations to transform the AST of the old version of a file to the AST of the new version of the same file. Fluri *et al.* used bi-gram string similarity for calculating the similarity between two leaf nodes (*i.e.*, identifier names). Thus, the detection of renamings is based on the declaration, whereas in our case we consider def-uses when available. As in our case, the detection of renamings depends on the pre-established thresholds. To evaluate their approach, Fluri *et al.* built a benchmark that consists of 1,064 manually classified changes of eight methods extracted from three open source programs. Only four of the classified changes are renamings, specifically one method and three parameters.

Kawrykow et Robillard (2011) measured the impact of non-essential differences on approaches aimed at detecting change couplings based on association rule discovery (Zimmermann *et al.*, 2004). Kawrykow and Robillard treated the renaming of an entity as an "essential" change, while the updated (*i.e.*, impacted) statements of this renaming are considered as "non-essential". Although the end goal is different from ours, detecting renamings is part of both approaches. To detect renamings Kawrykow and Robillard propose DIFFCAT, which is built on the approach of Fluri *et al.* (2007). That is, they used CHANGEDISTILLER to detect method and field renamings. However, they further enhanced it to also detect class, local variable, and additional field renamings. Since the goal of DIFFCAT is different from ours, it favors precision at the expense of recall. We compare REPENT performances with DIFFCAT on renamings detected in a random sample of revisions of dnsjava and JBoss. The results shows that REPENT has better performance in terms of precision and recall. Details of the comparison is provided in Section 5.4.

2.3   Summary

There have been quite a lot of woks on detection of refactorings at different level of abstraction. Renaming is also considered a refactoring, and thus we share with all these works rename refactoring. Our proposed approach is specific to detection of renaming and thus we identify renamings of all software entities. Comparison with the state of the art tool DIFFCAT shows that our proposed technique outperform in terms of precision and recall on a set of randomly selected versions of dnsjava and JBoss.

Previous works studied the importance of lexicon in software comprehension and quality. As suggested by Deissenbock *et al.* identifiers must be sufficiently distinctive yet must relate to one another and to the context in which they appear (Deissenbock et Pizka, 2005). Use of synonym and homonyms and their impact on consistency have been investigated in literature. It is interesting to further explore use of other semantic relations such as hypernyms, hyponyms, and antonyms. We propose a technique to classify renamings along all semantic relations. Such analysis will provide a better understanding on how identifiers are changed, and thus bring the attention of developers towards inconsistencies such as renaming towards antonym. Next chapter will provide details on the proposed taxonomy.

CHAPTER 3   RENAMING TAXONOMY

Renamings are classified along the dimensions of a taxonomy that extends and refines the taxonomy proposed in our previous work (Eshkevari *et al.*, 2011). We built the taxonomy based on a grounded-theory approach (Strauss, 1987; Glaser, 1992)

considering dimensions that we believe apply to source code identifiers and the terms that compose them, the latter being hard or soft words in the following (Lawrie *et al.*, 2006a) Specifically, we built the taxonomy by looking at identifier renamings, which we manually validated in our previous work (Eshkevari *et al.*, 2011), and grouping them into categories. The manual analysis required multiple iterations in order to converge and to consider all the dimensions of the proposed taxonomy.

The taxonomy comprises four dimensions, namely *entity kinds*, *forms of renaming*, *semantic changes*, and *grammar changes*. The first dimension distinguishes renamings based on the programming paradigm, whereas the last three dimensions distinguish renamings based on different natural language aspects. A summary of the taxonomy is reported in Table 3.1. The dimensions are orthogonal and thus each renaming will be classified in each dimension of the taxonomy. However, there are implicit relations between levels of the different dimensions. For example, classifying an identifier renaming in *form of renaming* as *formatting only* implies that in *semantic change* and *grammar change* it will be classified as *none*.

Concretely, in the field renaming `invParamsPtr` → `invalidParamReferencesPtr` the term `inv` is expanded to become `invalid`; `Params` changed to `Param`; `Reference` was added; and `Ptr` stayed unchanged (see Fig. 3.1). According to our taxonomy, this renaming will be classified as follows:

– **Entity kind:** *Field*,
– **Form of renaming:** *Complex* as two terms are changed, and one term is added,
– **Semantic change:** *Preserve meaning* as the term renaming `inv` → `invalid` is an expansion, and *add meaning* as the term `Reference` is added,
– **Grammar change:** *Part of speech change* as the term renaming `Params` → `Param` implies a change from plural to singular.

In the rest of this section we describe the different dimensions of the taxonomy.

Figure 3.1 Example of classifying a renaming based on the proposed taxonomy.

3.1 Entity Kinds

The first dimension of the taxonomy concerns the kind of the renamed entity, *i.e.*, whether the renamed entity is a package, type (*i.e.*, class, interface, or enumeration), field (*i.e.*, class variables, instance variables, constants, or enumeration constants), constructor, method, getter (*i.e.*, field accessor), setter (*i.e.*, field modifier), parameter, or local variable. We distinguish getters and setters from other methods using naming conventions, *i.e.*, the method name must start with the keywords "get" ("is" for boolean return type) or "set" and of a field name.

3.2 Forms of Renamings

The second dimension provides a classification of the renaming to determine whether one or more terms were changed, whether the renaming was solely related to text formatting, or whether it consisted in term reordering.

Simple: Those are renamings where one term has been changed, *e.g.*, `predeclareStatements` → `predeclare` where the term `Statements` has been removed and `override` → `overriding` where the only term composing the old name has been renamed.

Complex: Those are renamings where more than one composing terms have been renamed, *e.g.*, `IsAssignmentWithNoEffectMASK` → `Assignment HasNoEffect`.

Formatting only: Those are renamings where no term renaming occurs, but rather changing letter cases or adding/removing term separators. Examples include `getJRMPPort` → `getJrmpPort` and `JavaExtension` → `JAVA_EXTENSION`.

Term reordering: Those renamings consist in exchanging the position of terms composing the old identifier. Examples: `setDelaySocketClose` → `setSocketCloseDelay` and `pojoNoInterfacesIntro` → `noInterfacesPOJOIntro`.

3.3 Semantic Changes

The semantic changes dimension concerns changes (or not) in the meaning of the identifier due to the addition/removal of terms to/from the identifier or due to changing one or more terms with terms having different (or same) meaning. As a result, identifiers change while (i) preserving their meaning, (ii) changing their meaning, (iii) adding meaning, or (iv) removing meaning. There is no semantic change when only the format or order of terms change.

3.3.1 Preserve Meaning

Renamings falling in this category preserve the meaning of the identifier.

Synonym: The old and new terms have the same meaning (according to a given ontology). For example, in the renaming `isPotentialMatch` → `isPossibleMatch`, the two terms are synonyms.

Synonym phrase: One or more terms in the old identifier are renamed to one or more terms in the new identifier while preserving the meaning. Example: `javadocNotVisibleReference` → `javadocHiddenReference` where the renamed terms (`visible` and `hidden`) are antonyms

and one of them is negated (`visible` and `not visible`).

Spelling error: Examples of correction and introduction of spelling errors are `actionMesasage` → `actionMessage` and `sourceField` → `fiieldInfo` respectively. While one can easily understand the rationale of spelling error correction, spelling error introduction can happen as a side effect of a renaming.

Expansion: A renaming towards expansion occurs when a term—often not belonging to the English dictionary—is expanded into a (longer) term, often belonging to the English dictionary: `setAuthMechanism` → `setAuthenticationMechanism` and `collab` → `collaboration`.

Abbreviation: A renaming towards abbreviation is the opposite of a renaming towards *expansion* and occurs when a term—often belonging to the English dictionary—is contracted into a shorter term, often not belonging to the English dictionary: `packageName` → `pkgName` and `operationDesc` → `opDesc`.

3.3.2 Change in Meaning

Renamings falling in this category include cases in which the renaming changes the meaning of the old identifier.

Opposite: The new term has the opposite meaning of the old term (antonym), *e.g.*, `disableLookups` → `enableLookups`.

Opposite phrase: One or more terms in the old identifier are renamed to one or more terms in the new identifier with an opposite meaning: `isNotPrimitiveType` → `isPrimitiveType` where the term `Primitive` is negated. Also, it can happen that a term is replaced by a synonym and is negated: `isNonModifiableContainer` → `canUpdateContainer`.

Whole-part: The new and the old terms hold a whole-part relation (holonym/meronym); respective examples are `Point` → `Line` (fictitious example) and `body` → `node`.

Whole-part phrase: When more than one whole-part relation exists in the renamed identifiers. An example in this category is `Path` → `FileAndDirectory` (fictitious example[1]).

Unrelated: The old and new terms have unrelated meanings. It is the case for the terms `expression` and `script` in the identifier renaming `expressionModel` → `scriptModel`.

---

1. The example is fictitious as we did not classify any of the renamings detected in the five programs as *whole-part phrase*. However, for the sake of completeness and because results may be different for different programs we decided to keep this level of the taxonomy.

3.3.3 Narrow Meaning

Specialization: The meaning of the old identifier is narrowed when a term is renamed to its hyponym, *e.g.*, `thrownExceptionSize` → `boundExceptionLength`, where the new term `Length` is a hyponym of the old term `Size`.

Specialization phrase: Adding a term that specifies another term narrows the meaning of the old identifier: `item` → `todoItem` and `type` → `authType`. To do so, we consider nouns and adjectives modifiers (as specifiers) added before a term because, based on our qualitative analysis done using grounded theory, these are the most common modifiers that developers use to specify other terms.

3.3.4 Broaden Meaning

Generalization: Opposite to *specialization*, here the old and new terms hold a generalization relation, *i.e.*, the new term is a hypernym of the old term, *e.g.*, `getAccessRestriction` → `getAccessRuleSet`, where the term `Rule` is a hypernym of term `Restriction`.

Generalization phrase: Opposite to *specialization phrase*, here a specifying term is removed. Examples: `eventName` → `name` and `getInitialRepetitions` → `getRepetitions`. As in specialization phrase, we consider nouns and adjectives as modifiers.

3.4 Add Meaning

Add meaning renamings happen when an identifier is renamed by adding one or more terms and such a change does not fall into any of the cases discussed above, *i.e.*, in which the meaning is kept or changed. In other words, the added terms add meaning to the identifier, rather than changing (*e.g.*, generalizing, specializing, or negating) the current meaning: `_delete` → `deletePossible` and `flags` → `typeAndFlags`.

3.5 Remove Meaning

Remove meaning renamings happen when an identifier is renamed by removing one or more terms and, again, the change does not fall into any of the above cases. That is, the term removal also removes meaning from the identifier. Examples: `includeRule` → `rule` and `removedPackagePath` → `packagePath`.

3.5.1 None

No change in any of the terms in the identifier implies no semantic change, *i.e.*, the *semantic change* will be *none.* This is when the change is only in letter cases or adding/removing term separators.

3.6 Grammar Changes

The grammar changes dimension concerns changes in the part of speech of terms. We further classify part of speech changes into verb conjugation changes, singular to plural changes (and vice versa), or other (*e.g.*, change from noun to adverb).

3.6.1 Part of Speech Change

Those renamings occur when the part of speech of any term composing the old identifier changes. We consider the part of speech set contained in the Penn Treebank Tagset (Marcus *et al.*, 1993). Thus, a grammar change occurs when an adjective is changed to a verb, as in `getUpdatedSize` → `updateFigGroupSize`. We further focus our attention on nouns and verbs as, to the best of our knowledge, they represent the most critical part of identifiers. Specifically, nouns are shown to be the most important in terms of meaning (Capobianco *et al.*, 2013), whereas verbs are used in naming methods and thus changes of verbs can imply changes in functionality (Abbott, 1983; Bruegge et Dutoit, 2003).

3.6.2 None

There is no grammar change when the modified terms' part of speech remains the same. For example, renaming method `isPotentialMatch` to `isPossibleMatch` does not imply any grammar change as both terms, `Potential` and `Possible`, are tagged as adjective. Moreover, renamings classified as *formatting only* in the *forms of renaming* will also be classified as *none* in the *grammar change*, *e.g.*, `getJRMPPort` → `getJrmpPort`.

3.7 Summary

In this chapter we provide the taxonomy of renamings classification along four orthogonal dimensions: *entity kinds*, *forms of renaming*, *semantic changes*, and *grammar changes*. While first dimension aims to provide insight on how often programming entities renamed the other three dimensions aim to provide a view of linguistic transformations of identifiers. The following chapter explains renamings detection in details.

Table 3.1 Summary of the identifier renaming taxonomy

Entity kinds	Package	
	Type	
	Field	
	Constructor	
	Method/Getter/Setter	
	Parameter	
	Local Variable	
Forms of renaming	Simple	
	Complex	
	Formatting only	
	Term reordering	
Semantic changes	Preserve meaning	Synonym
		Synonym phrase
		Spelling error correction/introduction
		Expansion
		Abbreviation
	Change in meaning	Opposite
		Opposite phrase
		Whole-part
		Whole-part phrase
		Unrelated
	Narrow meaning	Specialization
		Specialization phrase
	Broaden meaning	Generalization
		Generalization phrase
	Add meaning	
	Remove meaning	
	None	
Grammar changes	Part of speech change	Singular/Plural
		Verb conjugation
		Other
	None	

CHAPTER 4   RENAMING DETECTION

In this chapter we explain the detection component of the REPENT in details.

Fig. 4.1 shows the processing steps of REPENT at a high level of detail to outline the renaming detection and classification processes. REPENT first detects a set of candidate renamings by means of file context *diff* and filters out false renamings using def-uses pairs, textual analysis, and heuristics. Finally, REPENT classifies renamings along the taxonomy dimensions.

4.1   Methodology

The first step of the process, step 1 in Fig. 4.1, is detailed in Fig. 4.2. In this step, REPENT compares two source files by applying a line differencing algorithm, the Unix context diff algorithm, which produces as output a set of line mappings. REPENT uses the mapped source code lines to compare and map entities declared into mapped lines and identify candidate renamings.

The comparison of source files is performed between two consecutive versions of a file or, in case of renamed files, between the file with the old name and the one with the new name. To identify file renamings, REPENT first builds the list of candidate file renamings consisting of (i) all possible couples formed by one file removed in the change set and one file added in the same change set and (ii) explicit renamings in the versioning system (SVN only). For CVS a change set is computed by grouping files based on the commit time stamp (less than 200 seconds between commits belonging to the same change set), commit note, and committer name (Zimmermann *et al.*, 2004). Then, REPENT evaluates each couple and selects the best option if the difference between the two files is reasonably low. Specifically, we use the Unix *diff* algorithm to compare the number of changed lines between two files and we consider them as a file renaming if the difference does not exceed a relative threshold of 60%. The value for the threshold is estimated based on the central tendency of explicitly renamed files as logged by the versioning file system.

The output of the comparison between two files is a mapping between lines of the old and new files; four cases must be considered:

1. **One-to-one line mapping:** one line of the old file is mapped onto one line of the new file. Fig. 4.3 shows an example where line 12 is mapped onto line 14. These two mapped lines are completely unrelated. Indeed, this is a case where the line mapping
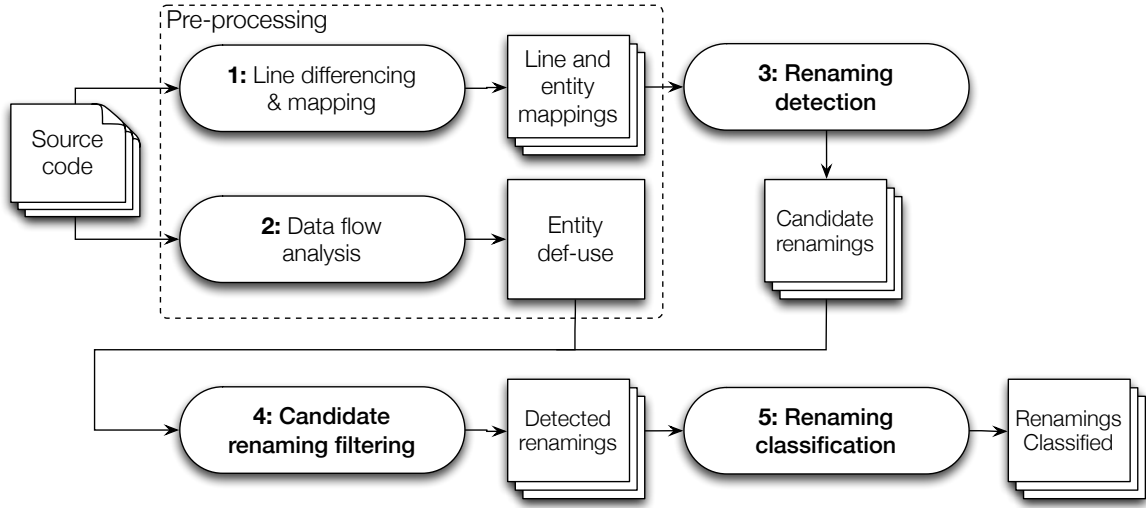
Figure 4.1 REPENT: Renaming detection and classification process.
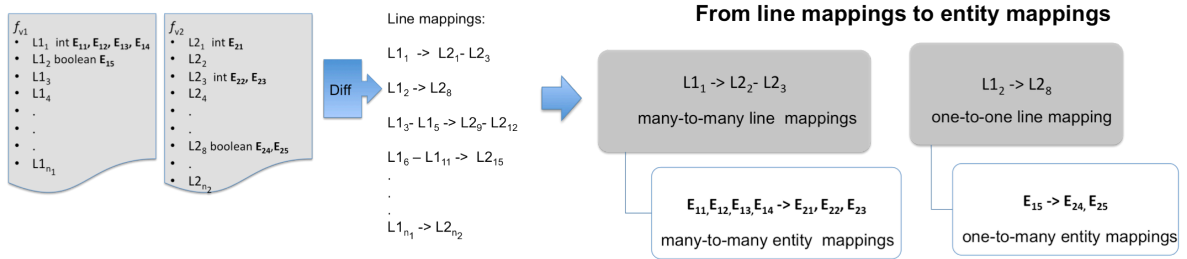
Figure 4.2 Details of the renaming detection process.

fails to map the lines correctly. To overcome this limitation of line mapping, we perform a cross validation step (see Section 4.2).

2. **One-to-many line mapping:** one line of the old file is mapped onto multiple lines of the new file.

3. **Many-to-one line mapping:** multiple lines of the old file are mapped onto one line of the new file.

4. **Many-to-many line mapping:** multiple lines of the old file are mapped onto multiple lines of the new file. Fig. 4.4 shows an example where lines 203 to 206 of the old file are mapped onto lines 203 to 206 in the new file.

Once REPENT creates lines mappings, it maps entities that are declared into mapped lines, *i.e.*, it creates entities mappings. To this end, REPENT first parses source code, creates an Abstract Syntax Tree (AST) using Eclipse Java Development Tools (JDT), and identifies the line numbers of all declared entities. Next, given the line mappings and the entities declarations for each source code line, REPENT creates an entity mapping. Again, four cases are possible:

1. **One-to-one entity mapping:** there is exactly one entity of a same kind that is declared in the old line(s) (*e.g.*, local variable) as well as in the new line(s). In this case, the old entity is mapped onto the new entity.

2. **One-to-many entity mapping:** there is only one entity declared in the old line(s), while there are many entities declared in the new line(s). In this case, the old entity is mapped onto several new entities.

3. **Many-to-one entity mapping:** this is the converse of the previous case, *i.e.*, several old entities are mapped onto one new entity.

4. **Many-to-many entity mapping:** there are several entities declared in the old line(s) and in the new line(s). In this case, the old entities are mapped onto the several new entities.

The entity mapping computed at this (early) stage has to be considered as possible mappings and thus as candidates renamings.

Fig. 4.3 shows an example of a one-to-one line mapping corresponding to a one-to-one entity mapping in Tomcat. The developer added a method `terminate` at line 14 of the new file. The line mapping algorithm maps line 12 in the old file—containing the declaration of method `loadNative`—onto line 14 of the new file. REPENT discovers only one entity declared in

the old file, as well as in the new one and builds the mapping, *i.e.*, candidate renaming, `loadNative` to `terminate`[1].

Fig. 4.4 shows a less trivial example where many entities are declared into the mapped lines (*i.e.*, many-to-many line mapping and many-to-one entity mapping). Thus, in this case, four entities (`STATE_INITIAL`, `STATE_INITIALIZED`, `STATE_STARTED`, and `STATE_STOPED`) are mapped into one (`STATE_PRE_INIT`).

For entities that are part of candidate renamings REPENT performs data flow analysis. REPENT first builds a symbol table considering the entities scope, signature, and line number. Then, it identifies modifications and uses within and across files by resolving the imports of the files.

4.2   Candidate Renaming Filtering

The set of mapped entities computed in the previous step may contain false positives. For example, if a code fragment is moved from the top to the bottom of its file then the context *diff* may not trace it—produces incorrect results and several entity mappings may be created. The mapped entities cannot however be discarded outright if the old entity also exists in the new files. For example, a developer can move the body of a method into a new one (with a new method name); replacing the body of the old method with a call to the new method. Thus, REPENT must first assign a score to each entity mapping, then check if the entity in the old file also exists in the new one, and if so, compute a score for this new pair, and based on this latter score (compared to the other scores), keep or prune the other entity mappings. This latter step is referred in the following as "cross validation consistency check".

Fig. 4.5 reports details on the REPENT filtering strategy. Two cases may occur: (i) both entities in the candidate renaming have def-uses or (ii) at least one of the entities has no def-use.

The second case is particularly common for getters and setters as they are often automatically generated and not necessarily used in the program.

For entities with def-uses, REPENT calculates the score of two entities involved in a candidate renaming as the textual similarity between the lines where the entity def-uses have been detected. For the entities without def-uses, the score is the textual similarity between the two entity declarations.

Finally, once scores are available, cross validation consistency check is performed. The fol-

---

1.   This mapping is not desired and that the candidate renaming will be later filtered out due to the low similarity between the two entity declarations.

Old version

1	public class AprImpl {
...	...
12	private void loadNative() {
...	...
...	}
...	// Temp - testing only, will be moved
...	//to separate file
...	public void main(String args[]) {
...	...
...	}
...	}

<<line mapping>>

New version

1	public class AprImpl {
...	...
14	public native int terminate();
...	...
44	public void loadNative(String libPath) {
...	try {
...	if (libPath==null)
...	libPath="jni_connect";
...	// XXX use load() for full path
...	if (libPath.indexOf("/") >=0 ||
...	libPath.indexOf("\\") >=0) {
...	System.load(libPath);
...	} else {
...	System.loadLibrary(libPath);
...	}
...	} catch(RuntimeException ex) {
...	ex.printStackTrace();
...	}
...	}
...	}

loadNative → terminate

<<entity mapping>>

Figure 4.3 Example of one-to-one entity mapping.

Old version

	...	...
	...	public static final int ACC_INIT_END=1;
	...	public static final int ACCOUNTS=2;
	...	// State
O1	203	public static final int STATE_INITIAL=0;
	204	public static final int STATE_INITIALIZED=1;
	205	public static final int STATE_STARTED=2;
	206	public static final int STATE_STOPED=3;
	...	...
	...	// ------- local variables -------
O2	...	private int state=STATE_INITIAL;
		...
		/**
		* Adds a new Context to the set managed
		* by this ContextManager.
		* @param ctx context to be added.
		*/
		public void addContext(Context ctx)
		throws TomcatException {
		...
O3		if(state == STATE_INITIAL)
		return;
		...
		}
	...	
		}

<<line mapping>>

New version

	...	...
	...	public static final int ACC_INIT_END=1;
	...	public static final int ACCOUNTS=2;
	...	// State
	203	/** Server is not initialized
	204	*/
N1	205	public static final int STATE_PRE_INIT=0;
	206	/** Server was initialized, engineInit() was called.
	...	* addContext() can be called.
	...	*/
	...	public static final int STATE_INIT=1;
	...	...
	...	// ------- local variables -------
N2	...	private int state=STATE_PRE_INIT;
	...	...
	...	public final void initContext(Context ctx)
	...	throws TomcatException {
N3	...	if (state != STATE_PRE_INIT) {
	...	...
	...	}
	...	}
	...	...
	...	/**
	...	* Adds a new Context to the set managed
	...	* by this ContextManager.
	...	* @param ctx context to be added.
	...	*/
	...	public void addContext(Context ctx)
	...	throws TomcatException {
	...	...
N4	...	if(state == STATE_INITIAL)
	...	return;
	...	...
	...	}
	...	
		}

STATE_INITIAL STATE_INITIALIAZED STATE_STARTED STATE_STOPPED → STATE_PRE_INIT

<<entity mapping>>

Figure 4.4 Example of many-to-one entity mapping.
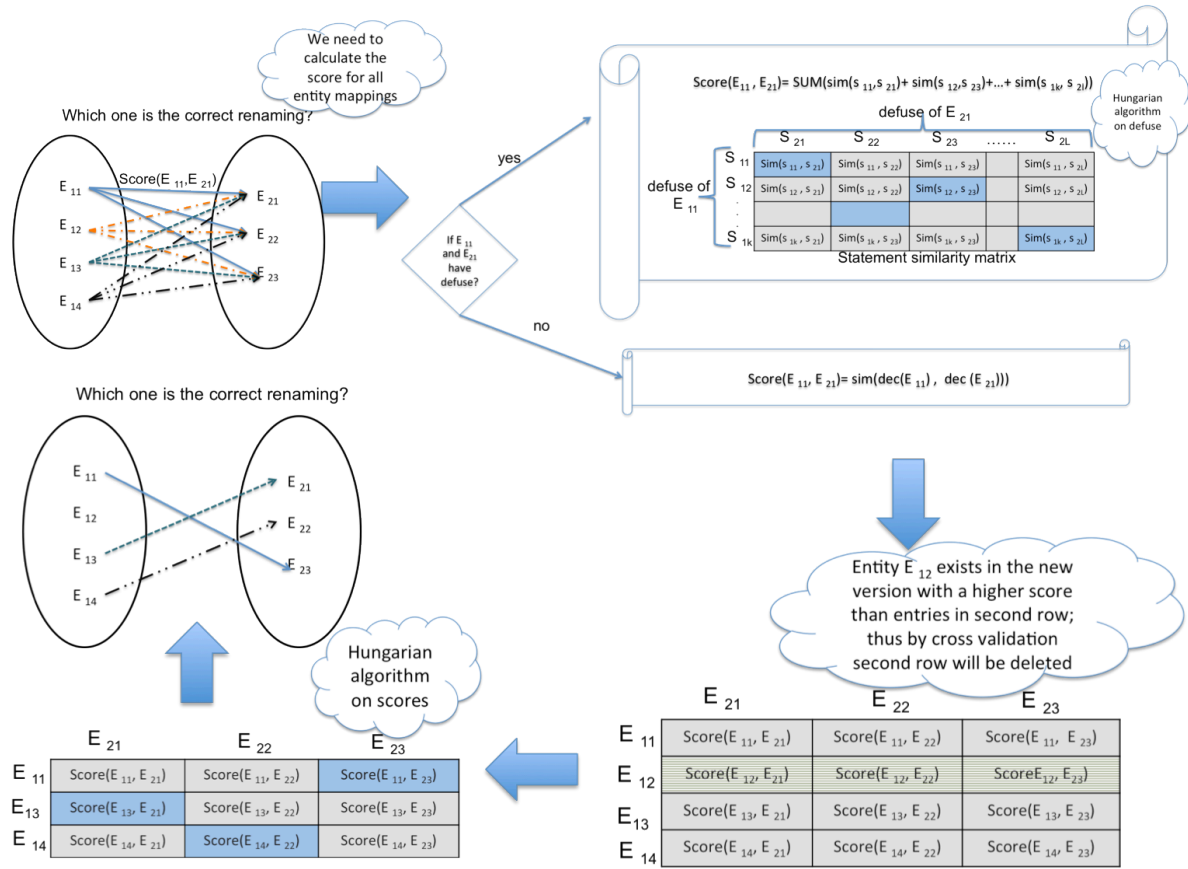
Figure 4.5 Candidate renaming filtering process.

lowing subsections provide details on how REPENT computes such scores and filters out likely false renamings.

Computing the Score between Entities with Def-uses

Let us assume that there are n statements in the old file where a given entity, say E_l, is either defined or used (or both). Also, let us assume that E_l has a candidate mapping with the entity E_k, and that in the new file there are m statements where the entity E_k is defined or used. To assign a score to the matching (E_l, E_k) REPENT creates an $n \times m$ *statement score matrix*. A matrix entry (i, j) contains the similarity score between statement s_i and s_j respectively of the old and new release. Before computing the score, REPENT removes the name of the entities from both statements (*i.e.*, s_i and s_j), to remove any bias introduced by similar entity names. The assigned score is based on the Normalized Levenshtein edit Distance (NLD) (Levenshtein, 1966), defined as:

$$NLD(s_i, s_j) = \frac{LD(s_i, s_j)}{length(s_i) + length(s_j)} \tag{4.1}$$

where $LD(s_i, s_j)$ is the Levenshtein edit Distance between s_i and s_j. The score assigned to the statement pair (s_i, s_j), cell (i, j), of the *statement score matrix* is computed as $1 - NLD(s_i, s_j)$.

REPENT uses heuristics to prune low scores and thus reduce false positives. If the similarity is lower than a given threshold, named *Statement Similarity Threshold (SST)* (see Appendix), the selection is filtered out by setting the (i, j) matrix entry to zero. REPENT uses different threshold values depending on the kind of entity. The reason behind this choice is the different nature of the definitions and uses of different entities.

As shown in Fig. 4.5, REPENT applies the Hungarian algorithm on the statement score matrix to identify the best possible matching between statements. The Hungarian algorithm
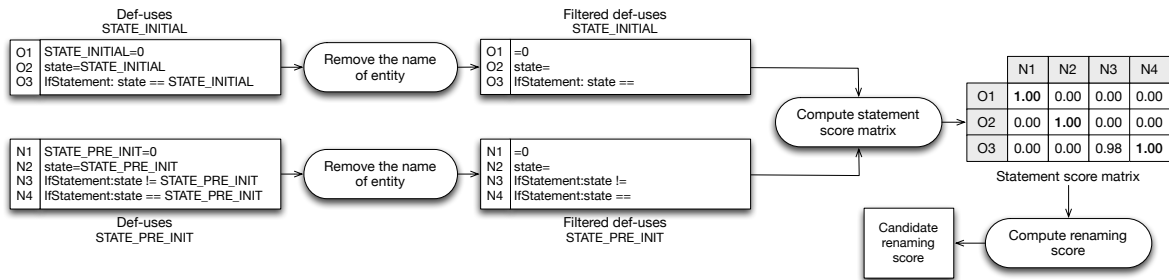
Figure 4.6 Computing the score of a candidate renaming in presence of entity def-uses.

(Kuhn, 1955) is an optimization algorithm used to solve the assignment problem, *e.g.*, assigning tasks to people, which given a score/cost matrix, will find the best assignment, *i.e.*, maximizing/minimizing the score/cost between matrix lines and columns.

The result of the Hungarian algorithm is used to assign a score to the pair of entities (E_l, E_k) as follows. If the number of mapped statements with scores higher than zero is higher than a predefined threshold—named *Number of Matched Statements Threshold (NST)*—then the score between the two entities is the sum of the similarities between the mapped def-uses; otherwise, the score is set to zero.

Consider the example in Fig. 4.4, which describes a many-to-one entity mapping where entities have def-uses. In this example, REPENT computes the scores for the following four possibilities:

STATE_INITIAL $\rightarrow$ STATE_PRE_INIT
STATE_INITIALIZED $\rightarrow$ STATE_PRE_INIT
STATE_STARTED $\rightarrow$ STATE_PRE_INIT
STATE_STOPED $\rightarrow$ STATE_PRE_INIT

Fig. 4.6 illustrates the computation of the score for the first possibility, *i.e.*, STATE_INITIAL $\rightarrow$ STATE_PRE_INIT. In particular, once identified the def-uses, the entity names are removed and the similarity is computed between all the possible pairs of def-uses for the old and new entities. Such similarities are then stored in the *statement score matrix*.

The application of the Hungarian algorithm on the statement score matrix for the candidate renaming STATE_INITIAL $\rightarrow$ STATE_PRE_INIT creates the following statement matches: O1 $\rightarrow$ N1, O2 $\rightarrow$ N2, and O3 $\rightarrow$ N4, where Oi (Ni) is the i^{th} statement of the older (newer) entity. The similarity scores of the matched statements for fields should be more than 0.8 according to Table .1. The heuristic to prune false renamings for fields states that one must have a number of matched statements higher or equal to 40% (see Appendix) of the longest list of def-use statements for a candidate matching to be considered (this is to say, 40% of the largest score matrix dimension).

In the example described above, the matrix is a three by four matrix, which means that at least two def-uses must be matched ($4 \times 40\% = 1.6 \approx 2$). Actually, three statements have been matched, thus the score for the two entities (STATE_INITIAL and STATE_PRE_INIT) is the sum of the scores mapped by the Hungarian algorithm: $3.00 (= 1.00 + 1.00 + 1.00)$.

Computing the Score between Entities without Def-uses

If at least one of the entity has no def-uses, the computation of the similarity is based on the textual similarity between the two entity declaration statements, ds_1—the declaration statement of the entity in the old file—and ds_2—the declaration statement of the entity in the new file. Also in this case, if the similarity $(1 - NLD(ds_1, ds_2))$ is lower than a predefined threshold—named *Declaration Similarity Threshold: (DST)*—we discard the candidate renaming by setting its similarity to the minimum, *i.e.*, to zero.

To better understand how the similarity is computed in this particular case, let us go back to the example of Fig. 4.3. In this example, only one candidate renaming has been identified, *i.e.*, `loadNative` $\rightarrow$ `terminate`. Because both methods have no def-uses, we simply compute the similarity between the declarations, *i.e.*, between `private void loadNative()` and `public native int terminate()`. The similarity between these two entities $(1 - NLD(ds_1, ds_2)$, where ds_1 is the declaration statement of `loadNative` and ds_2 is the declaration statement of `terminate`) is 0.63 which is lower than the fixed DST threshold (0.7); thus the score between `loadNative` and `terminate` is set to zero.

Computing the Score between Candidate Package Renamings

Using a package means importing that package or types declared in that package. Thus, a score based on def-uses is not suitable for package renamings as it does not allow one to distinguish between a case where types were moved from one package to another and a case where the package was actually renamed. In both cases, the use of the package will change in an identical manner; however, only the second situation is an actual package renaming. Thus, REPENT first computes the score of candidate package renamings using the version control system (to identify renamed, added, and deleted folders); then it filters out candidate package renamings that do not match a change in the folder structure of the version control system by setting their score to zero.

Filtering Out Likely False Renamings

As noted above, REPENT must also cross-check if entities exist in both files before promoting a candidate renaming to a real renaming. In essence, two further steps are needed: cross validation and entity score evaluation.

Cross validation consistency check   This step is necessary as there may be cases in which the line mapping is not precise. These cases often occur when source code fragments

are moved within the same file (as in Fig. 4.3). Since line mapping relies on the Unix *diff*, and since *diff* relies of the context, *i.e.*, surrounding statements, such moving of code may result in wrong line mapping.

To cope with this imprecision, REPENT cross-checks each entity of a candidate renaming—*i.e.*, in the old file REPENT looks for an entity with the same name as the new entity and in the new version REPENT looks for an entity with the same name as the old entity—and computes the score between the two entities, using the algorithms described in Sections 4.2 and 4.2. If this score is greater than the score of the candidate renaming, then the renaming is discarded, *i.e.*, its score is set to zero. In the example in Fig. 4.3 the cross-checking identifies that `loadNative` → `terminate` is not an actual renaming, since the method `loadNative` is still present in the new file. Instead, in the example shown in Fig. 4.4, the cross-check fails since neither of the entities in the older version of the file appear in the new file and the entity of the new file does not appear in the old file.

Entity score evaluation   For each candidate renaming the scores of all possibilities are stored in $p \times q$ *entity score matrix*, where p and q are the number of old and new entities, respectively. Once such a matrix is produced for a candidate renaming, we apply the Hungarian algorithm to collect the assignments between the entities. Clearly, if a row (or a column) of the score matrix contains all zero values, the related entity is not renamed.

Fig. 4.7 shows the computation of the entity score matrix for the example in Fig. 4.4. This is a 4×1 entity score matrix where only one entry has a score greater than zero. Thus, when applying the Hungarian algorithm we obtain that the best score is 3.00, *i.e.*, the score between `STATE_INITIAL` and `STATE_PRE_INIT`. Therefore, REPENT concludes that `STATE_INITIAL` has been renamed to `STATE_PRE_INIT`.

4.3   Summary

REPENT uses line differencing algorithm, the Unix context diff algorithm, to identify line mapping between two consecutive versions of files. File renamings are detected by comparing

Figure 4.7 REPENT: Filtering false positives renamings.

the added and deleted files in the repositories of the programs under analysis. Line mapping is used to build the first renamings candidates list by extracting entity declarations of the mapped lines. Using data flow analysis REPENT extracts for each software entity its def/use statements. False positive are removed from renamings candidates by comparing the def/use statements of mapped entities. In the next chapter we introduce characteristic of the five open source Java programs we used to evaluate the accuracy of the proposed approach along with the results of renamings..

CHAPTER 5   RESULTS OF DETECTION

This chapter reports the results of an empirical study conducted to evaluate the accuracy and completeness of REPENT renaming detection. The *goal* of this study is to analyze the detection accuracy of REPENT with the *purpose* of investigating to what extent undocumented renamings can be identified. The *perspective* of the study is that of researchers, who are interested in investigating how REPENT is suitable to identify renamings. The evaluation has been carried out in the *context* of the source code history of five Java open-source programs (ArgoUML, dnsjava, Eclipse-JDT, JBoss, and Tomcat).

ArgoUML[1] is a UML modeling tool. dnsjava[2] is a Java Domain Name System (DNS) implementation. Eclipse-JDT is a set of plug-ins that adds the capabilities of a full-featured Java IDE to the Eclipse[3] platform. JBoss AS[4], in the following simply referred to as JBoss, is a Java application server. Tomcat[5] is an implementation of a servlet container and Java Server Page (JSP) engine. Table 5.1 reports the main characteristics of the analyzed programs: the analyzed time periods, size ranges in KLOCs, numbers of files, numbers of analyzed revisions, and numbers of committers. ArgoUML and Eclipse are versioned under CVS, while all other Java programs are versioned under SVN.

The following sections detail the procedure and results of the accuracy evaluation of RE-PENT, in particular we discuss precision and recall.

---

1. `http://argouml.tigris.org`
2. `http://www.xbill.org/dnsjava`
3. `http://www.eclipse.org`
4. `http://www.jboss.org/jbossas`
5. `http://tomcat.apache.org`

Table 5.1 Characteristics of the analyzed programs.

Program	Analyzed period	KLOCs (range)	Files (total)	File revisions	Committers
ArgoUML	1998-2012	1-20	300	68,400	42
dnsjava	1998–2011	9-35	365	1,415	2
Eclipse-JDT	2001–2006	2,089–6,949	5,758	54,571	50
JBoss	1999–2011	2,000-1,200	40,003	25,028	422
Tomcat	1999–2006	5–315	12,205	46,498	79

5.1 Research Questions and Study Procedure

To evaluate the accuracy of REPENT, we address the following two research questions:

RQ-DP: *How accurate is the set of renamings detected by* REPENT*?* This research question aims at estimating the accuracy of the detection approach, measured in terms of *precision.* Since Section 5.2 presents an empirical study on how developers rename identifiers, precision indicates the accuracy of the renamings used in such a study.

RQ-DR: *How complete is the set of renamings detected by* REPENT*?* This research question aims at estimating the completeness, measured in terms of *recall*, of REPENT with respect to the set of renamings performed by the developers of the analyzed programs. Recall gives an estimate of the representativeness of the analyzed renamings reported in Section 5.2.

The following subsections detail how we evaluate the accuracy of REPENT.

5.1.1 Evaluating the Precision of the Detection Approach: Manual Validation

This section explains in details the process of evaluating the precision for detection. Precision is computed by manually validating a sample of renamings from the analyzed programs. We reuse the oracle built from our previous work, *i.e.*, manually validated renamings for Tomcat, to calibrate thresholds (see Appendix).

We then evaluate the approach on all programs by validating a statistically representative sample for each. Sampling separately for each program allows us to evaluate the detection also on programs with low number of renamings (*e.g.*, dnsjava), which otherwise would have less chances to be selected (*e.g.*, with respect to JBoss) if the total population was considered.

To estimate the size of the representative samples we choose a confidence level of 95% and a confidence interval of ±5% Sheskin (2007). Thus, we can be 95% sure that the precision of the approach on the detected renamings for each program will be the precision estimated for the sample ±5%.

Once the sample sizes have been determined, we use a stratified random sampling to select the renamings to be validated.

This means that for each program we first group renamings based on the kind of entity being renamed, *i.e.*, the first dimension of our taxonomy (*e.g.*, type, method). Then, we estimate the proportion of each group with respect to the total population of detected renamings for that particular program and we use the same proportion for the sample. For example, if type renamings are 5% of the total population of detected renamings in ArgoUML, 5% of the sample must be type renamings. Finally, we randomly select the sample for each group.

The sample size and the number of detected renamings are reported in Appendix .

The advantage of using stratified random sampling is in ensuring that all groups are represented Black (1999). If random sampling is used instead, the chances that a package renaming is selected for validation for Eclipse-JDT for example are almost zero (4 over 12,557).

The manual validation was conducted as follows. For each chosen detected renaming, two evaluators independently inspected the source code of both versions of the file between which the entity was renamed. Then, the evaluators marked the renaming as true positive (TP) or false positive (FP). In all cases in which the two authors provided a different classification for the renaming, the inconsistency was discussed and solved. When needed, a third author also reviewed such candidate renamings in which case the classification was obtained by a majority vote. All available details (comments, uses, neighboring entities) contributed to the decision-making.

Whenever the lack of knowledge prevented us from taking a decision, the renaming was removed and replaced by a new one; the process was iterated up to the required sample size.

Finally, the precision is computed as the fraction of detected renamings in the validated sample that the validators classified as TP. In other words, given the subset of detected renamings sampled for validation, TPS the set of those classified as true positives, and FPS the set of those classified as false positives, the precision Pr is given by:

$$Pr = \frac{|TPS|}{|TPS| + |FPS|}$$

5.1.2 Evaluating the Recall of the Detection Approach: Comparison with Documented Renamings

To evaluate the recall, ideally one should have the knowledge of all actual renamings that occurred in a program. Unfortunately, such information is not available for open-source programs.

With approximately two hundred thousand file revisions, it is very tedious, time consuming, and error prone to manually extract all entity renaming. There are (relatively few) cases in which developers documented renamings in the versioning system commit notes of Java programs. Hence, we estimate the recall as the proportion of such documented renamings also detected by REPENT.

To identify documented renamings, we filter the commit logs and consider only the commits whose note contain the keyword "renam". Then, we complement the automatic filtering with

a manual analysis of the identified commit notes, with the aim of pruning out false positives. Typical cases of false positives are the commits in Eclipse-JDT related to changes to the refactoring feature, which includes a renaming feature. Other false positives are related to renaming of files not containing source code, *e.g.*, images or documentation files. Then, we analyze the source code of the files involved in the documented renaming to locate renamed entities, and hence verify whether such renamings are detected by our approach. Hence, given DCR, the set of documented renamings identified as described above and DR the set of detected renamings, the recall Rc is the proportion of documented renamings that are also detected by REPENT:

$$Rc = \frac{|DR \cap DCR|}{|DCR|}$$

5.2 Results

This section analyzes the results achieved aiming at answering our research questions **RQ-DP** and **RQ-DR**.

5.2.1 RQ-DP: How accurate is the set of renamings detected by REPENT?

Table 5.2 reports, for the analyzed programs, the precision of REPENT computed for each kind of entity as well as for the overall sample of renamings.

We observed an average precision of about 88%, as expected slightly lower than the one computed when calibrating the thresholds (about 92%, see Table .1). The number of detected package renamings is very low, thus at most 1 package renaming was sampled per program, which results in a 0% or 100% precision and explains the lowest precision reported in Table 5.2, *i.e.*, 67% for package renamings. REPENT has a somehow low precision—compared to the rest of the entity kinds—for parameters renamings (76%). The set used to calibrate the thresholds may not have a sufficiently large set of parameter renaming and thus thresholds may need to be recalibrated. The accuracy of REPENT may also be impacted by methods not being called in the program, or getters/setters automatically generated and again not used. In such cases REPENT relies on the fixed DST string matching threshold. Higher values may improve precision but again at the cost of recall.

5.2.2 RQ-DR: How complete is the set of renamings detected by REPENT?

In the following we describe in details the computation of recall of detection.

Table 5.2 Estimated precision *Pr* for renaming detection of different entities (95% ±5 confidence).

	ArgoUML	dnsjava	Eclipse-JDT	JBoss	Tomcat	Overall
Package	0%	-	100%	100%	-	67%
Type	100%	100%	100%	100%	100%	100%
Constructor	100%	100%	100%	100%	100%	100%
Field	100%	93%	95%	94%	73%	93%
Method	100%	98%	94%	94%	93%	94%
Getter/Setter	100%	83%	97%	96%	79%	90%
Parameter	90%	54%	92%	87%	67%	76%
Local variable	95%	86%	93%	84%	90%	92%
Overall	**97%**	**78%**	**94%**	**91%**	**80%**	**88%**

Table 5.3 Comparison with documented renamings.

Program	Files involved	Documented renamings
ArgoUML	77	4
dnsjava	113	229
Eclipse-JDT	140	52
JBoss	146	50
Tomcat	66	2

Table 5.3 reports the number of files involved in the commits whose log message suggest possible renamings. Documented renamings refer to the number of renamings that we found in the files committed with the log messages that were either documenting a renaming in a vague manner (*e.g.*, *"renamed some stuff"*) or explicitly (*e.g.*, *"rename Name.fromStringNoValidate(String) to Name.fromStringNoException(String)"*).

Table 5.4 reports—for each kind of entity—the detected proportion of documented renamings. We conclude that although sometimes renamings are documented, this is not a general rule. This result further motivates the use of REPENT as a renaming re-documentation tool. Table 5.4 also shows that documented renamings often pertain to types and, thus, to constructors.

For ArgoUML, the number of documented renamings is very low. We detect three out of four

Table 5.4 Detected documented renamings and recall *Rc* of different entities.

	ArgoUML		dnsjava		Eclipse-JDT		JBoss		Tomcat		Overall	
Package	-		-		100%	(1/1)	-		-		100%	(1/1)
Type	-		94%	(58/62)	18%	(4/22)	95%	(20/21)	-		78%	(82/105)
Constructor	-		100%	(134/134)	100%	(4/4)	95%	(18/19)	-		99%	(156/157)
Field	100%	(1/1)	100%	(4/4)	100%	(3/3)	100%	(1/1)	-		100%	(9/9)
Method	-		100%	(1/1)	100%	(7/7)	100%	(5/5)	-		100%	(13/13)
Getter/Setter	67%	(2/3)	-		-		100%	(2/2)	100%	(2/2)	86%	(6/7)
Parameter	-		100%	(28/28)	100%	(7/7)	-		-		100%	(35/35)
Local variable	-		-		88%	(7/8)	100%	(2/2)	-		90%	(9/10)
Overall	75%	(3/4)	98%	(225/229)	63%	(33/52)	96%	(48/50)	100%	(2/2)	92%	(311/337)

renamings. The renaming our approach fails to detect is a complex combination of renaming and refactoring activities, where the renamed getter method was abstract in the old version and became a concrete method in the new version. Also, the field associated with the getter was moved from the superclass to a subclass.

In Tomcat there are only two documented renamings and REPENT detects both of them.

For JBoss, REPENT only fails to identify two documented renamings, *i.e.*, one class and one constructor. Both entities are defined in the same file and the file was renamed as well. REPENT misses these renamings as the difference between the original and renamed files is greater than the 60% relative threshold for detecting renamed files.

For Eclipse-JDT, REPENT fails to identify one local variable and 18 class renamings. mainly because Eclipse-JDT used (for the analyzed period) CVS. Therefore, as explained in Section 4.1, we grouped commits using the heuristic of Zimmermann *et al.* Zimmermann et Weisgerber (2004)

However, sometimes commits belonging to the same change occur in different days and developers do not always use consistent commit notes. As a consequence, REPENT fails to identify some file renamings.

Finally, dnsjava was (surprisingly) the program containing the highest number of documented renamings, despite being the smallest one. In this case, our approach detected 98% of the documented renamings, *i.e.*, it fails to detect only four class renamings. These classes have no def-uses. Although REPENT detects these class renamings as candidate renamings, it filters them as false positive, since their similarities are less than the declaration similarity threshold for type renamings (see Table .1).

Table 5.5 summarizes REPENT precision and recall. The first column reports the data set

used for the evaluation; the second column corresponds to the size of the data set (*i.e.*, numbers of renamings); the last column reports the measures. Precision was evaluated over a representative sample from the detected renamings while recall was evaluated with respect to renamings documented by developers. The documented renaming set is extracted from the repository of the programs under analysis. Although its size is not as large as the sample used to evaluate the precision, it is an unbiased oracle as the entries are reported by the developers.

> **REPENT detection accuracy:** Overall, we found that REPENT reaches high precision (88%) and recall (92%) thus being suitable for most of the foreseeable tasks.

5.3 Threats to Validity

This section discusses the threats to validity that can affect the studies performed. We discuss the most important threats that can affect this kind of study, *i.e.*, *construct validity*, *conclusion validity*, and *external validity.*

Construct validity threats concern the relationship between theory and observation. As for the renaming detection study, construct validity threats are due to the detection of renamed files, the estimation of precision, and recall. For programs using SVN versioning system, when files are not explicitly renamed we compare deleted versus added files for two consecutive revisions. We use the Unix *diff* algorithm to compare the number of changed lines between all possible combinations of added and deleted files. We select the best possible combination (*i.e.*, smallest number of lines changed) if it does not exceed a relative threshold of 60%. The value for the threshold is estimated based on the central tendency of explicitly renamed files as logged by the versioning system files. For CVS, we first group files based on the commit date, log message, and committer ID Zimmermann *et al.* (2004), and then apply the same heuristic used for SVN, considering as deleted files the files that appear for the first time in the "Attic" directory.

Table 5.5 REPENT precision and recall.

Data set	Size	Accuracy (measure)
Sample over detected renamings	1723	88% (*Pr*)
All documented renamings	337	92% (*Rc*)

As for precision, the manual validation could be affected by subjectiveness or human error. Specifically, regarding the classification, we may be affected by the lack of domain knowledge—i.e., the original developers of the five programs may have classified differently some of the renamings. To mitigate those threats, the validation was performed by two persons independently and, in case of different classification, the renaming was discussed, and a third person was also asked to perform the classification. As for the recall, we are aware that the sample of documented renamings may not be fully representative of the entire set of renamings performed in a project. First, this happens because developers do not always document renamings in commit notes, especially if they are performed together with other changes. Second, most of the documented renamings are related to types and thus to constructors, *i.e.*, to entities whose names can impact on other developers' activities.

Internal validity threats are related to factors, internal to our study, that can affect our results. Such a threat is mainly due to the calibration of thresholds. Indeed, different calibration could have produced different results, and also indirectly affected the subsequent renaming classification study. Appendix explains how thresholds have been empirically determined. Clearly, the calibration has been performed based on a thorough validation on a sample set of detected renamings of Tomcat, when using low thresholds. The use of data from one of the projects to calibrate thresholds was also used in other studies Bavota *et al.* (2011); Koschke *et al.* (2006); Thummalapenta *et al.* (2010). However, this does not guarantee that the choice is optimal for the other projects.

Conclusion validity threats concern the relationship between the experimentation and the outcome. Our study is an exploratory study in which we do not make use of statistical tests to reject specific hypotheses. The only issue related to conclusion validity is the representativeness of the sample used to validate the renaming detection precision. To evaluate the detection accuracy, we performed for each program a stratified random sampling across the kinds of entities considering a confidence level of 95% and a confidence interval of at least $\pm 5\%$.

External validity threats concern the generalizability of our results. Both the evaluation of the renaming detection approach and the exploratory study were conducted on data from a subset of the evolution history of five Java open-source projects. Although we have chosen a pretty variegated set of projects—belonging to different domains and development organization, and having different size—it could happen that replicating the study on other projects could lead to different results, *e.g.*, different performances of the renaming detection tool. A different matter is the application of the proposed approach on programming languages different from Java. Our preliminary results shows that the proposed approach is applicable

to other languages, but some adaptation may be needed.

5.4   Comparison with Existing Approaches

In the following we compare REPENT performances with existing approaches that detect renamings as part of their refactoring detection.

Comparison with DiffCat: We compared REPENT to DIFFCAT on renamings detected in a random sample of revisions of dnsjava and JBoss. For each randomly selected revision, we applied both tools and manually validated the detected renamings. We stopped sampling once both approaches reached a total of 100 renamings. Table 5.6 shows the results of the comparison[6]. The last column of Table 5.6 reports the false negatives (FN), *i.e.*, the number of true renamings that each approach does not detect. Overall, REPENT outperforms DIFFCAT in terms of precision and number of detected true renamings.

REPENT uses file context *diff* to reduce the search space of entities involved in potential renamings, the actual validation of renamings is mainly based on the def-use analysis or, when not available, on the similarity between declarations. File context *diff* can be replaced with any other differencing tool including CHANGEDISTILLER. However, the advantage of the file context *diff* is its speed and a simple comparison between file context *diff* and CHANGEDISTILLER showed that both tools are sensitive to cases where chunks of code are moved within a file. In addition, the comparison with DIFFCAT, which is built on CHANGEDISTILLER, shows that REPENT is more accurate in terms of detected renamings.

We share with all of the above approaches their general ideas and goal. We also use parsing and differencing technologies, though in a different combination, to achieve an approximated, lightweight, robust, and scalable approach. The novelty of our work is a renaming taxonomy directly conceived to better represent renamings on orthogonal dimensions, and a classifier that—by relying on WordNet and on the Stanford NLP—classified renamings according to the proposed taxonomy. Thus REPENT detects finer-grain details about renamings, such

---

6.   The detected and validated renamings can be found in the replication package.

Table 5.6 Accuracy of REPENT and DIFFCAT on a random sample of revisions.

	Precision		TP / Detected		FN	
	REPENT	DIFFCAT	REPENT	DIFFCAT	REPENT	DIFFCAT
dnsjava	99%	96%	104 / 105	99 / 103	16	21
JBoss	81%	72%	108 / 133	77 / 107	33	63

as the grammatical renaming type or the semantic type. Furthermore, our approach does not require a compilable program to work.

Comparison with Refactoring Crawler: REFACTORING CRAWLER was applied on various releases of Eclipse UI, Struts, and JHotDraw with accuracy as high as 85%. The syntactic analyzer is fast, however it may produce false positives; therefore, it is necessary to filter out the false positives via the semantic analyzer. As reported in the paper, the semantic analyzer does not scale up for real world applications with tens of thousands of entities. Detecting renamings is the common part between REFACTORING CRAWLER and our proposal. In addition to the renamings detected by REFACTORING CRAWLER, our technique also detects field, parameter, and local variable renamings. Moreover, we do not have any limitations for detecting renaming of abstract methods and methods declared in interface. Our technique can compare any two versions of Java source code and does not require the source code to be compilable unlike REFACTORING CRAWLER.

5.5 Summary

In this chapter we provide details of the studied programs, results of REPENT detection component and the accuracy of the results in terms of precision and recall. We discussed in details the threats to the validity of our study. The results show that renamings is a frequent activity during software evolution and thus confirms our thesis. Having an overall high precision and recall for detection of renamings assures that the input to the classification is a trust able and thus motivates the classification of detected renamings. Next two chapters explains how renamings are classified using our taxonomy and the results of classification for the five Java programs.

CHAPTER 6 RENAMING CLASSIFICATION

In this chapter we provide details of our renaming classifications.

The classification process of REPENT is summarized in Fig. 6.1. It entails a sequence of phases: (i) identifier splitting (Section 6.1), (ii) mapping of identifier terms (Section 6.2), and (iii) combining part of speech and semantic analyses (Section 6.3). Each phase is detailed in the following.

The renaming classifier heavily relies on tools—ontological databases such as WordNet (Miller, 1995) and natural language parsers such as the Stanford Part-of-Speech Analyzer (Toutanova et Manning, 2000)—explicitly conceived to process natural language documents rather than source code. As pointed out by Hindle *et al.* (Hindle *et al.*, 2011), such tools can be far from optimal when applied to source code. However, at the moment they represent, to the best of our knowledge, the most suitable technology for our purposes. In future, REPENT could be further improved by replacing or combining WordNet with a domain-specific ontology.

For example, in their recent work Yang and Tan (Yang et Tan, 2013) propose an approach to mine semantically related words in a project or multiple projects from the same domain. Similar work has been done by Howard *et al.* (Howard *et al.*, 2013) where the authors mine semantically similar words across projects from multiple domains. However, in the current status we could not apply the aforementioned approaches, because they would have required us to manually validate all the mined semantic relations, which would have required a deep domain knowledge for the projects considered in our study (which we do not have).

6.1 Identifier Splitting

This step aims at splitting both the old and new names into their composing terms. RE-PENT uses a Camel Case splitting algorithm. The output of this phase is the lists of terms composing the old name *i.e.*, $t_{1,1}, t_{1,2}, \ldots, t_{1,n_1}$ and the new name *i.e.*, $t_{2,1}, t_{2,2}, \ldots, t_{2,n_2}$, where n_1 and n_2 are the number of terms composing the old and new names respectively. For example, the identifier `getChildCount` is split into `get, child,` and `count`. More sophisticated identifier splitting approaches such as *Samurai* (Enslen *et al.*, 2009), *TIDIER* (Madani *et al.*, 2010; Guerrouj *et al.*, 2011), *Normalize* (Lawrie et Binkley, 2011), or *LINSEN* (Corazza *et al.*, 2012) can be plugged in. However, the current implementation of REPENT favors speed over accuracy; a Camel Case splitter is much faster than, for example, TIDIER (Madani *et al.*, 2010; Guerrouj *et al.*, 2011). Moreover, previous studies found that for Java, the
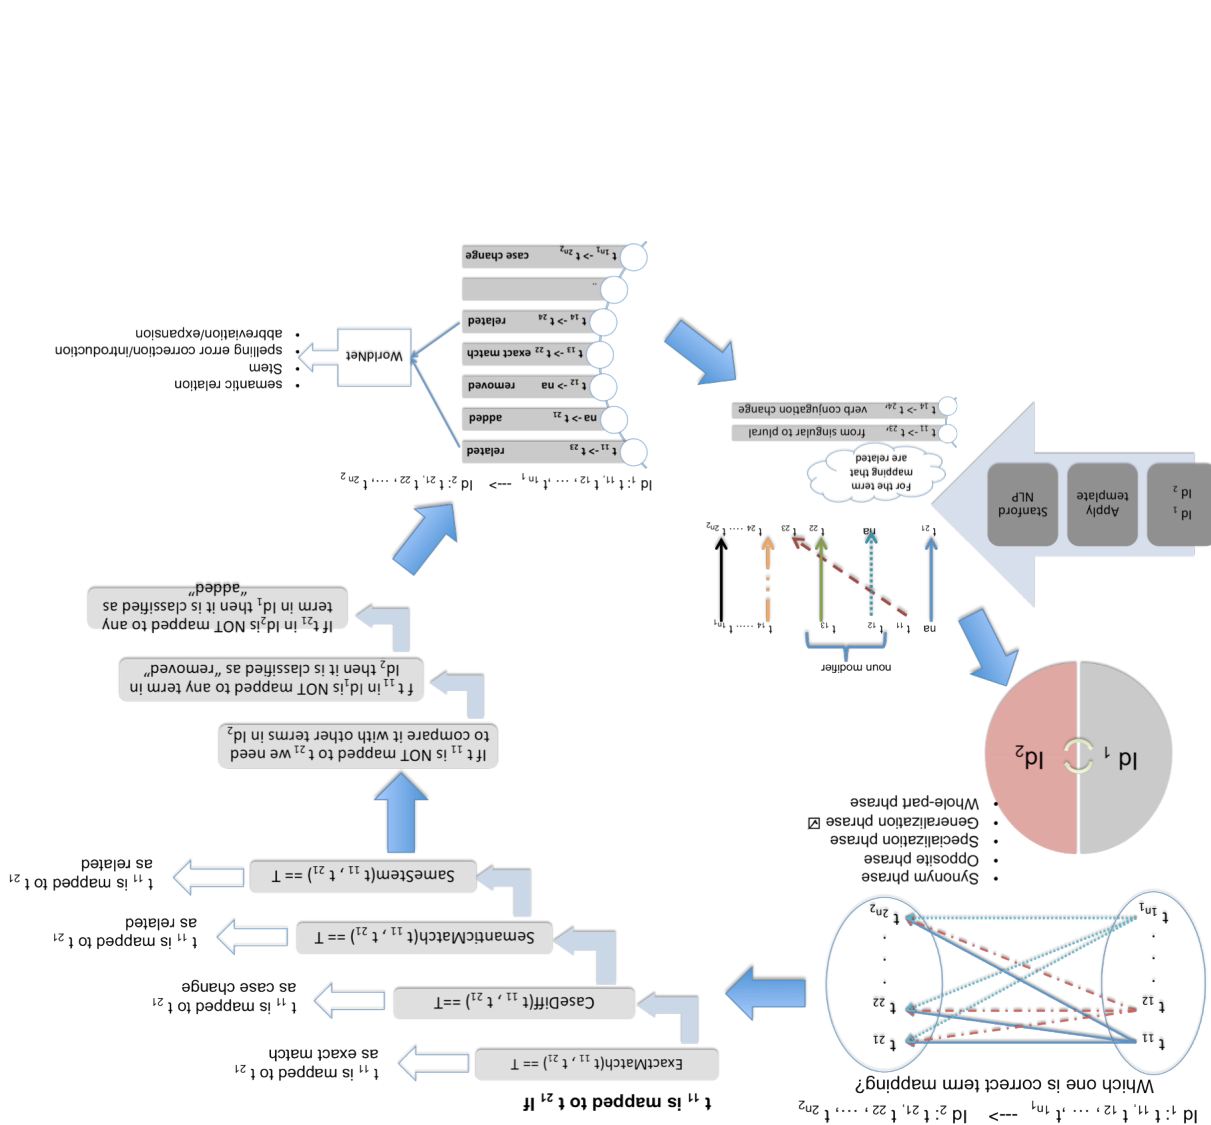
Figure 6.1 REPENT: Renaming classification process.

identifier splitting/expansion accuracy does not vary substantially between Camel Case and more sophisticated approaches (Madani *et al.*, 2010).

6.2 Mapping Terms

The second phase aims at mapping the n_1 terms $t_{1,1}, t_{1,2}, \ldots, t_{1,n_1}$ composing the old name onto the n_2 terms $t_{2,1}, t_{2,2}, \ldots, t_{2,n_2}$ of the new name. The mapping phase is, in turn, divided into two main steps depicted in Fig. 6.2. In the reported example REPENT must map the name `getStorage` onto `getMemoryBlock`.

First, REPENT discovers changed and unchanged terms plus added and deleted terms. To this aim, each term composing the old and new names is thought of as a source code line. In our example, the terms `get`, and `Storage` compose the lines of the first (old) file, while `get`, `Memory`, and `Block` represent the lines of the second (new) file. Then, a *diff* algorithm identifies churns of unchanged, added, removed, and changed terms (lines) between the two versions of a name (file), using an algorithm that solves the longest common subsequence (LCS) problem (Cormen *et al.*, 1990). In the example in Fig. 6.2, after applying such an algorithm, the renaming of `getStorage` to `getMemoryBlock` is considered as the removal of the term `Storage` and the adding of the terms `Memory`, and `Block`. The term `get` is identified as unchanged.

In the second step, REPENT performs a fine-grained analysis of changed terms (*i.e.*, the term `Storage` from the old name and the terms `Memory`, and `Block` from the new name in the example shown in Fig. 6.2). Such an analysis is based on Algorithm 1 that builds a term-by-term mapping and classifies it. A term $t_{1,i}$ of the old name is mapped onto a term $t_{2,j}$ of new name according to multiple criteria, encoded in the function $matching(t_1, t_2, matchType)$. Given two terms and a matching criterion, this function returns *true* if the terms match according to the matching criterion, *false* otherwise. Specifically, the matching is performed using, sequentially, the following criteria:

1. **Exact match**: if the two terms exactly match, *e.g.*, `get` and `get`.

2. **Case difference**: if the two terms only differ by the alphabetic letter case, *e.g.*, `Book` and `book`. If this does not happen, all terms are converted into lower case letter, and the subsequent criterion are matched.

3. **Semantic match**: if the two terms have any semantic relation according to the upper ontology WordNet. Words in WordNet are organized based on their relations. Synonyms are grouped into unordered sets, called synsets, which in turn are related using semantic and lexical relations. In the example reported in Fig. 6.2 the terms `Storage`

and `Memory` belong to the same synset. The semantic relations considered by REPENT are synonym, hyponym, hypernym, antonym, meronym, and holonym. REPENT first identifies semantic relation between two mapped terms and if there is no such semantic relation it looks for a semantic relation between words in the synsets of the two mapped terms. This process repeats for up to three synsets of each word in the synsets of two mapped terms.

For example, when looking if an antonym relation exists between two terms, REPENT first checks if there is an antonym relation. However, if this is not the case, REPENT further analyzes their respective synsets for antonym relation between two words, each belonging to the synset of the original terms. If an antonym relation is found, it will be considered as a relation of level 1. In the opposite case, *i.e.*, no relation is found, REPENT further looks for antonym relation between the words of the synsets of the synsets, *i.e.*, by doing a transitive closure up to level 3.

4. **Is stem**: if the two terms have the same stem according to the Porter (Porter, 1980) stemming. This rule is applied only if the *semantic match* rule fails. Indeed, if both terms are defined in WordNet, *e.g.*, `synchronization` and `synchronizing`, then they will be related according to the semantic match. If any of the two terms is not defined in WordNet, as it is the case of `invoc` from the identifier renaming `invocationType` → `invocType`, then the rule **is stem** is applied.

Algorithm 1 builds a mapping of terms of the old name onto any (not yet mapped) term of the new name, repeatedly traversing the terms of the new name, moving from the position of the term of the old name and using the above matching criteria, in the order in which they are mentioned.

After the term mapping has been performed, REPENT identifies mapped terms on which the renaming classification will focus, *i.e.*, all mapped terms that are not trivially mapped according to the **exact match** criterion. For example, in the identifier renaming `getStorage` → `getMemoryBlock`, both identifiers contain `get`

that is an exact match and thus is removed from further consideration.

After terms of the old name have been mapped onto terms of the new name, REPENT classifies the renamings at term level, as:

1. **Removed:** terms of the old name not mapped onto any term of the new name are classified as removed, as it is the case for the term `statement` from the identifier renaming `statementLength` → `length`.

2. **Added:** terms of the new name not mapped onto any term of old name are classified

as added. In the example reported in Fig. 6.2 the terms `Block` is identified as an added term.

3. **Matched:** terms of the old name mapped onto terms of the new name according to Algorithm 1 with an exact match, *e.g.*, the term `get` in the example reported in Fig. 6.2.

4. **Change case:** terms of the old name mapped onto terms of the new name according to Algorithm 1 with a case difference match, as this is the case for the term `jar` from the renaming `pJARFile` $\rightarrow$ `jarFile`.

5. **Related:** terms of the old name mapped onto terms of the new name according to Algorithm 1 with a **semantic match** or a **is stem** match. In the example reported in Fig. 6.2 the terms `Storage` and `Memory` are classified as related since there exists a semantic relation (synonym) between them.

Algorithm 1 Algorithm for mapping and classifying the n_1 terms of the old name onto the n_2 terms composing the new name.

---

for $matchType$ in ($exact$, $case\_difference$,$semantic$, is_stem) **do**

    for $x \leftarrow 1$ to n_1 **do**
      if $\neg mapped_1[x]$ **then**
        $y1 \leftarrow x$, $y2 \leftarrow x$;
        while $y1 > 0$ or $y2 \leq n_2$ **do**
          for y **in** $(y1, y2)$ **do**
            if $matching(t_{1,x}, t_{2,y}, matchType)$ AND $\neg mapped_2[y]$ AND $y > 0$ AND $y \leq n_2$ **then**

$$mapped_1[x] \leftarrow y;$$
$$mapped_2[y] \leftarrow x;$$

            end if
          end for
        $y1 - -$, $y2 + +$;
        end while
      end if
    end for
end for

REPENT uses the mapped terms to classify the renaming in dimension *forms of renaming* as follows:

Simple: when only one term is added, removed, or changed.

Complex: when more than one terms are added, removed, or changed.

Formatting only: the following two conditions hold: (i) all term mappings are matched and/or

change case and (ii) the two identifiers are the same when underscore and camel case are ignored.

Term reorder: the following two conditions hold: (i) at least two terms of the old identifier are matched to two terms in the new identifier while possibly changing case and (ii) the two identifiers are not the same when underscore and camel case are ignored.

REPENT refines **related** matches via WordNet to find semantic relations between terms, *i.e.*, synonymy, hyponymy, hypernymy, antonymy, meronymy, or holonymy and can thus classify the renaming in dimension *semantic change* as *synonymy*, *specialization*, *generalization*, *opposite*, or *whole-part* when the corresponding relation exists.

If no semantic relation is found REPENT checks whether there is a *spelling error* correction/introduction. REPENT assumes there is a spelling error if the following three conditions hold: (i) one of the two terms does not exist in WordNet but the other does, (ii) there is only small (string) difference between the two

(Levenshtein distance is 2 or smaller—see Appendix for a discussion on the threshold value), and (iii) one

term is not included in the other (to avoid misclassifying renamings such as `frame → jframe`).

Finally, if the previous checks fail, REPENT checks if there is an *abbreviation/expansion*. It assumes abbreviation/expansion if the following two conditions hold: (i) one of the two terms does not exist in WordNet but the other does and (ii) all characters of one are contained in the other.

6.3 Part of Speech and Semantic Analyses

In natural language, a word carries a specific meaning. Words are often grouped into phrases which in turn can be combined to form sentences. The meaning carried by a phrase can narrow, generalize, or change the meaning of an individual term within the phrase. By analogy with natural language, to grasp the meaning of an identifier, one cannot rely only on the terms constituting the identifier in isolation. For example, the term `visible` (from the identifier `JavadocNotVisibleReference`) and the term `hidden` (from the identifier `JavadocHiddenReference`) have opposite meaning, whereas the identifiers have the same meaning.

Thus, after terms are mapped between the old and new name, REPENT explores the relations between the terms within the same identifier to classify identifier renamings. REPENT builds a synthetic sentence out of the identifier, then it performs a part of speech analysis.

As identifiers do not always follow well-formed grammatical structure, before applying part of speech analysis using natural language tools we apply a sentence template. Different templates have been proposed in the literature by Abebe *et al.* (Abebe et Tonella, 2010) and Binkley *et al.* (Binkley *et al.*, 2011). For all kinds of entities, except methods, REPENT applies the List Item Template by Binkley *et al.* (Binkley *et al.*, 2011). Indeed, they provided evidence that this template outperforms the other three templates they evaluated. For the identifier `inclusionPatterns` the template produces *inclusion patterns*. However, if the first term is a verb, as it is suggested according to Java standard for method names, REPENT uses a different template, *i.e.*, the verb template: "Try to <identifier terms>". A template is just an aid provided to the part of speech tagger to guide its analysis; thus for the method name `markAsDefinitelyUnknown`, REPENT applies the verb template on the term sequence, *i.e.*, *Try to mark as definitely unknown.*

REPENT part of speech analysis uses the Stanford Part-of-Speech Analyzer [1]. The Stanford NLP classifies terms using the Penn Treebank Tagset (Marcus *et al.*, 1993), thus not only distinguishing between nouns, verbs, adjectives, and adverbs, but also distinguishing between the different forms. From this step beyond, we use the part of speech of each term— *i.e.*, whether it is a noun, being it singular or plural, an adjective, an adverb, etc.—and the relations between terms. More precisely, we are interested in the following relations: negation modifier (*i.e.*, the relation between a negation word and the word it modifies, as in the identifier `ignoreNotFoundField`), adjectival modifier (*i.e.*, a modifier relation between an adjective and a noun, meaning that the adjective specifies the noun, as in the identifier `binaryField`), and noun compound modifier (*i.e.*, a modifier relation between two nouns, meaning that one noun specifies the other, as in the identifier `methodSignature`).

Given a renaming pair, old and new names, REPENT processes the two part of speech analyses and uses heuristics to assign a semantic label to the renaming. The heuristics work as follows:

Synonym phrase: when the following two conditions hold: (i) there exists a term mapping where the two terms hold an antonym relation and (ii) one of the two terms is involved in a negation modifier relation.

Opposite phrase: when one of the following two conditions holds: (i) a negation modifier relation is added/removed while the modified term exists in both identifiers or (ii) a term renaming towards a synonym is accompanied with an addition/removal of a negation modifier relation.

Specialization phrase: when the following two conditions hold: (i) a term is added and (ii) it

---

1. We will refer to it as the Stanford NLP.

participates in a modifier relation, either adjectival or noun, with an already existing term.

Generalization phrase: when the following two conditions hold: (i) a term is removed and (ii) a modifier relation between the removed term and a term existing in both identifiers is also removed.

Whole-part phrase: when more than one term mapping pair holds a whole-part relation.

Add meaning: when the following two conditions hold: (i) there exists a term mapping that is classified as term added, and (ii) the added term is not the modifier of another term in the new identifier.

Remove meaning: when the following two conditions hold: (i) there exists a term mapping that is classified as term removed, and (ii) the removed term is not the modifier of a term in the old identifier.

Unrelated: when a term mapping does not fall into any of the levels of *semantic change.*

Part of speech change: when the part of speech of two mapped terms are different.

6.4   Summary

REPENT classifies the detected renamings based on our proposed taxonomy. The first step of classification is to split the old and new identifiers based on Camel Case splitting algorithm. Then, term mapping is performed to identify added, removed and mapped terms between the old and new identifiers. We use WordNet to identify the semantic changes between the mapped terms. Finally, we use Stanford Part-of-Speech Analyze to extract more complex semantic changes such as synonym, specialization, and generalization phrases. Moreover, the part of speech tagger allow us to identify grammar changes between the mapped terms. In the next chapter, we provide the results of our empirical study for classifications of renamings.
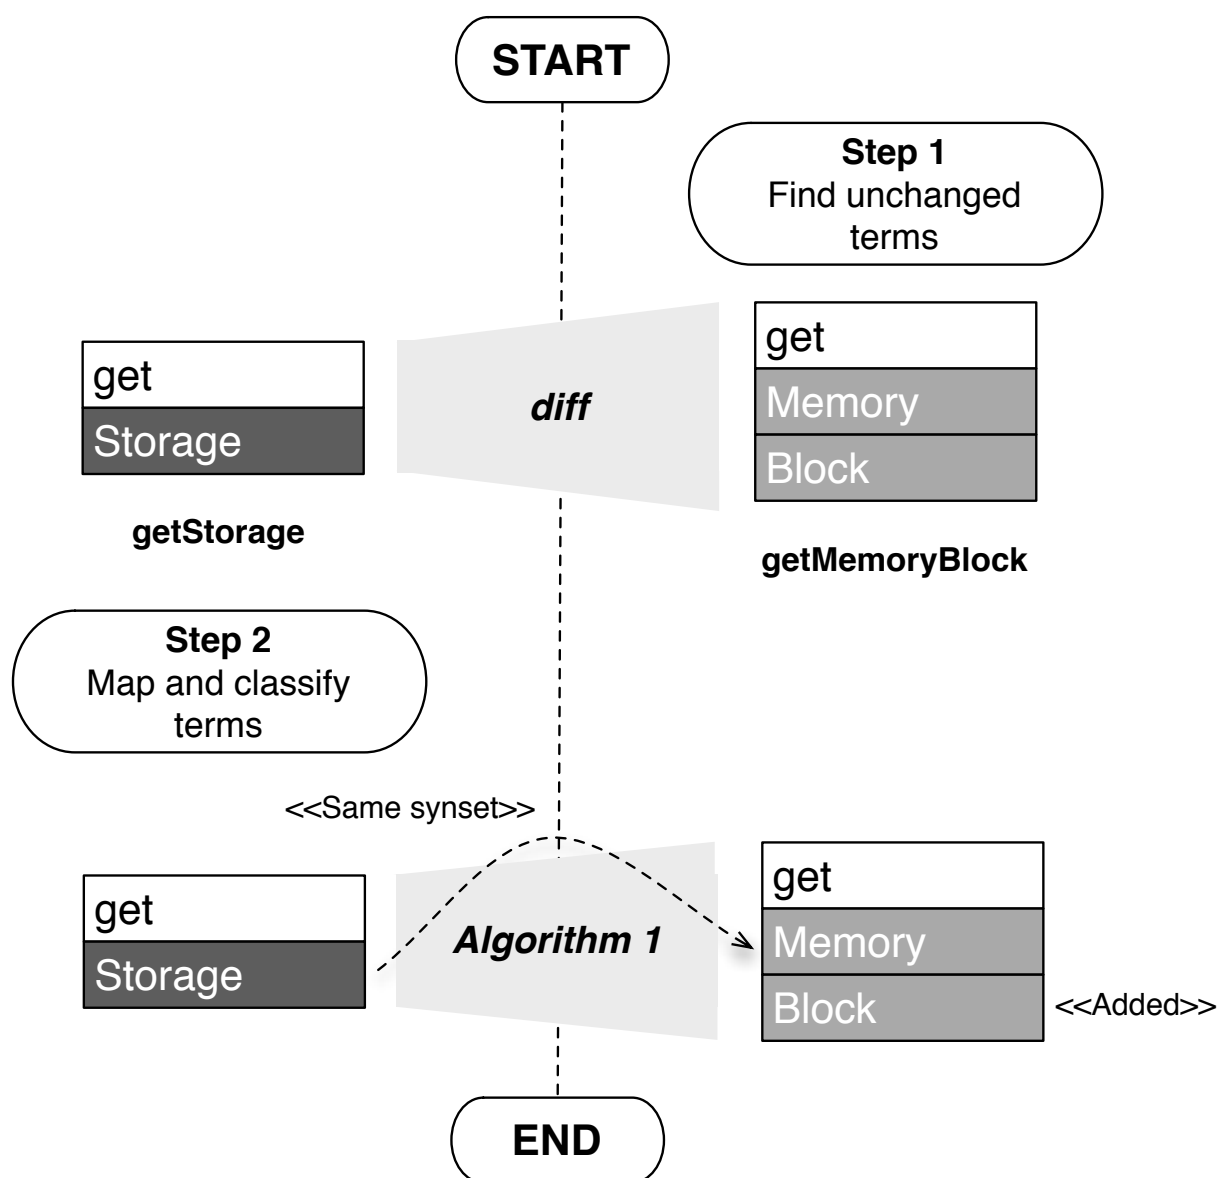
Figure 6.2 REPENT: Term mapping and classification process.

CHAPTER 7   RESULTS OF CLASSIFICATION

The *goal* of this study is to use REPENT to analyze renamings over the evolution history of software programs with the *purpose* of investigating to what extent such renamings fall into the dimensions defined in the taxonomy of Chapter 3. The *perspective* of the study is that of researchers who are interested in investigating how identifiers are renamed in the same *context* as the study reported in Chapter 5

7.1   Research Questions and Study Procedure

This empirical study aims at automatically detecting and classifying renamings in the five Java programs described in Table 5.1.

Since we use REPENT to identify renamings, we analyze the classification accuracy of REPENT to evaluate to what extent the classification of renamings with respect to our taxonomy is affected by the performance of REPENT, thus answering the following research question:

RQ-CP: *How accurate is the set of classified renamings?* This research question aims at providing an estimate of the accuracy of the classification, measured in terms of *precision*. Such an estimate indicates the accuracy of the results of this exploratory study, reported in Section 5.

While the focus of the previous research questions is to evaluate the reliability of REPENT as a tool to detect and classify identifier renamings, the following research questions are the core of this study, *i.e.*, they study the renaming phenomenon using the taxonomy defined in Chapter 3. For each dimension of the taxonomy, we investigate to what extent renamings of the programs fall into the different levels of the dimension.

RQ1: *To what extent do renamings occur with respect to the different kinds of entities?* Specifically, we compute the number and proportion of renamings occurring for package, type, constructor, method/getter/setter/function, field, parameter, and local variable names to investigate which entities are more prone to be renamed.

RQ2: *What kinds of changes occur to terms composing identifiers when these are renamed?* In other words, we compute the number and proportion of simple, complex, formatting only, and term reordering renamings to investigate which forms are more frequent.

RQ3: *What kinds of semantic changes occur in identifiers when they are renamed?* In other

words, we compute the number and proportion of renamings that preserve, change, narrow, broaden, add, and remove meaning to study how the renamings of the five programs are distributed over the different levels of semantic change.

RQ4: *What kinds of grammar changes occur in identifiers when they are renamed?* Specifically, we investigate to what extent the renamings imply changes to nouns (singular/plural), to verb conjugations, or other part of speech changes.

In order to find answers to our research questions, we investigate—from both a quantitative and qualitative point of view—how identifier renamings detected in the studied programs follow the taxonomy of Chapter 3.

7.1.1 Evaluating the Precision of the Classification Approach: Manual Validation

To evaluate the accuracy of the renaming classification on the Java programs we extract, for each level of each dimension of the taxonomy, a representative random sample ensuring a confidence interval of $\pm 10\%$ for a confidence level of 95%. This is different from the sampling in Chapter 5 where the sample is representative for each program stratified over the kinds of entities. Here, the confidence level and interval criteria are met for each level and each dimension of the taxonomy for the total population of classified renamings. For the *semantic change* dimension this means a representative number of expansions, a representative number of abbreviations, etc.

7.2 Results

In this section we first discuss how accurately REPENT classifies renamings (Section 7.2.1), then, in Sections 7.2.2 to 7.2.5, we discuss how the renamings of the programs that we studied follow the taxonomy defined in Section 3, *e.g.*, to what extent those renamings *preserve meaning* or consist of changes that are *formatting only*.

Table 7.1 Renamed entities identified by REPENT - Java programs.

	Package	Type	Field	Constructor	Method/Getter/Setter	Parameter	Local variable	Total
ArgoUML	0	18	2,156	16	391	690	712	3,983
dnsjava	0	67	58	159	219	448	144	1,095
Eclipse	4	180	1,942	139	3,205	3,218	3,845	12,533
JBoss	7	656	1,805	475	3,985	3,406	3,247	13,581
Tomcat	0	69	478	48	830	507	428	2,360
Total	**11**	**990**	**6,439**	**837**	**8,630**	**8,269**	**8,376**	**33,552**

7.2.1 RQ-CP: How accurate is the set of classified renamings?

We manually analyzed a sample of the classified renamings of the five Java programs to evaluate in how many cases REPENT correctly or wrongly classified the changes in the identifiers. In addition, when REPENT fails to correctly classify a change we further investigate the reason. The sample size and the number of correctly classified renamings for each dimension of taxonomy are reported in Tables .2 to .4 in Appendix .

With respect to the classification of *forms of renaming*, REPENT has an overall precision of 98% (see Table .2). The few misclassified cases are due to wrong term mapping.

For the classification of *semantic changes*, REPENT exhibits an accuracy of 80% (see Table .3). REPENT is very accurate in classifying renamings that add or remove meanings (82% and 91% respectively), and the miss classification is due to splitting of identifiers that are all in lower cases.

REPENT is also accurate in classifying renamings that preserve the meaning (overall precision of 93%). The lowest accuracy is achieved by REPENT when classifying renamings as *narrow* and *broaden meaning*, 62% and 69% respectively. Wrongly classified renamings in the category of *semantic changes* are due to wrong splitting, wrong term mapping, or wrong relations between terms.

We also observed cases where REPENT misclassified a renaming because of the ontological database. For example, WordNet infers a hyponym relation between "is" and "get" and an antonym relation between "long" and "short". Those relations are not valid in the context of Java where in many cases "is" and "get" are used for accessors of boolean attributes and where "long" and "short" are primitive types. Approaches by Yang and Tan Yang et Tan (2013) and Howard *et al.* Howard *et al.* (2013) can be used to improve the classification of *semantic changes.*

The classification accuracy of REPENT when classifying *grammar changes* is 74% (see Table .4). REPENT is accurate in classifying changes of nouns from singular to plural (and *vice versa*) and changes in verb conjugation, *i.e.*, the precision is 100% and 79%. Moreover, REPENT is very accurate where there is no grammar change (100%). By contrast, REPENT performances are fairly low (20% of precision) in classifying changes in other part of speech changes for Java programs.

Most of these cases are mainly due to the Stanford NLP not being accurate when parsing source code identifiers. In a very recent work, Gupta *et al.* Gupta *et al.* (2013) proposed an approach for part of speech tagging of source code identifiers and showed that the approach parses identifiers 10% to 20% more accurately. Unfortunately, at current date, the source

code identifier tagger is not publicly available. However, it could be integrated in REPENT to improve its performances. Other reasons for misclassification were incorrect splitting or incorrect term mapping.

> **REPENT classification accuracy:** REPENT almost perfectly classifies forms of renamings (98%) and classifies reasonably well *semantic changes* (with an accuracy of 80%). The lowest performance is on the classification of *grammar changes* (with an overall accuracy of 74%).

7.2.2 RQ1: To what extent do renamings occur with respect to the different kinds of entities?

Table 7.1, and Fig. 7.1 report the number and proportion, respectively, of renamings occurring for package, type, field, constructor, method/getter/setter, parameter, and local variable names [1].

The entities that are more prone to be renamed are methods, and local variables. In addition, we see that parameters tend to change more. Generally, such changes reflect the evolution of the programs. Indeed 89% of the surveyed developers confirm that they rename while changing functionality.

As an example in JBoss, REPENT identified that the parameter `webserviceClientDeployer` has been renamed to `webservicesClient`. The renaming was performed to reflect a change in the functionality, as confirmed by the log message: *decouple WebserviceClientDeployer from JSR109ClientService.*

There is a large number of field renamings in ArgoUML. In this program, REPENT identified 2,156 field renamings, representing about 54% of the renamings. About 67% of the field renamings consisted solely of underscore removal in the beginning of the field, *i.e.*, renaming away from the Hungarian notation. Finally, 11% of the field renamings were performed due to the third party library that ArgoUML uses for logging purposes, *i.e.*, Log4J. The class `org.apache.log4j.Category` was deprecated and users of the library were supposed to use class `org.apache.log4j.Logger` instead. As a result, fields were renamed from `cat` to `LOG` or `logger`.

Finally, 448 out of 1,095 of the renamings in dnsjava were performed on method parameters due to massive renaming activities that heavily changed the API and hence the method parameters. Also, a considerable number (20%) of the parameter renamings consisted solely

---

1. For the classification of renamings we only considered the TP renamings from the validated sample, thus the number of classified renamings is lower than the number of detected renamings.

Figure 7.1 Proportion of renamed entities identified by REPENT.

in removing a leading underscore. Conversely, in 52% of the parameter renamings (all being renamed in the same revision) a leading underscore was added. In addition, the majority of those parameters was renamed following the same pattern: names starting with the letter `r` were renamed to start with underscore. Examples of those include `rname` → `_name`, `rclass` → `_dclass`, and `rtype` → `_type`. All those renamings were performed as part of *"The big rewrite..."* in class `Record`, a generic resource record, or in one of its many subclasses. About 6% of the parameter renamings were performed on parameters of type `DataByteOutputStream` where the name changed from `dbs` to `out`.

> **RQ1 conclusion:** Renamings occur mostly for method, parameter, and local variable names, with 26%, 25%, and 25% of the renamings respectively. Field renamings represent also a large proportion of the renamings—close to 20%. Finally, type renamings, which in Java imply constructor renamings, as well as package renamings, represent a small proportion of the renamings (less than 3% each for type and constructor renamings, less than 1% for package renamings).

Table 7.2 Forms of renamings identified by REPENT in Java programs.

	Simple	Complex	Formatting only	Term reordering	Total
ArgoUML	1,787	493	1,702	1	3,983
dnsjava	894	85	116	0	1,095
Eclipse-JDT	7,910	4,456	132	35	12,533
JBoss	8,094	4,786	655	46	13,581
Tomcat	1,638	658	58	6	2,360
Total	**20,323**	**10,478**	**2,663**	**88**	**33,552**

Table 7.3 Semantic changes identified by REPENT in Java programs.

	Preserve meaning	Change in meaning	Narrow meaning	Broaden meaning	Add meaning	Remove meaning	None	Total
ArgoUML	77	1,311	97	59	789	441	1,661	**4,435**
dnsjava	12	576	62	24	79	312	112	**1,177**
Eclipse	413	5,522	1,297	1,104	4,481	3,750	122	**16,689**
JBoss	580	6,042	1,130	851	4,884	4,198	564	**18,249**
Tomcat	259	1,061	186	138	731	628	35	**3,038**
Total	**1,341**	**14,512**	**2,772**	**2,176**	**10,964**	**9,329**	**2,494**	**43,588**

7.2.3 RQ2: What kinds of changes occur to terms composing identifiers when these are renamed?

Table 7.2 reports the forms of identified renamings by REPENT, while Fig. 7.2 shows the proportion of the different forms. Most of the renamings were classified as *simple* renamings, where developers renamed a single term. In Eclipse-JDT, JBoss, and Tomcat, there is a considerable number of *complex* renamings while this form of renamings is not so frequent in the other two programs.

In ArgoUML, a substantial number of renamings is consist of *formatting only*—43% of the identified renamings. The analysis of this form of renamings indicates that in 77% (1,314 out of 1,702) of the cases, the renaming relates to the removal of leading underscores from identifiers, *i.e.*, towards Java naming conventions, which recommend not to start identifiers with underscore. However, in 2% (35 out of 1,702) of those renamings, a leading underscore was added to identifiers, hence, against Java naming conventions. These results go along with the opinion of surveyed developers—when the name of an entity does not follow the language naming conventions, 34% would definitely rename while 46% would probably rename.

REPENT identified only a few renamings (88 of the classified renamings) that were performed to change the order of terms. One may expect that *term reordering* involves entities with limited scope, thus limiting the impact of the renaming. However, in the analyzed renamings only 15% of such re-orderings involved local variables. We conjecture that developers tend to reorder terms to improve the comprehensibility of the identifier and avoid

Figure 7.2 Proportion of forms of renamings identified by REPENT.

misunderstanding. For example, in JBoss a developer changed a method parameter name from `serviceDest` to `destService`, to clarify that the parameter contains the address of the destination service.

> **RQ2 conclusion:** The majority of renamings (61%) are simple renamings, *i.e.*, only one term of the identifier is renamed. In a considerable number of the Java renamings (31%) multiple terms are changed simultaneously. Less often (8%), renamings consist only of formatting changes. Even less often are those renamings where the terms of the identifiers were reordered.

7.2.4 RQ3: What kinds of semantic changes occur in identifiers when they are renamed?

Table 7.3 report the number of semantic changes identified by REPENT while Fig. 7.3 shows the proportion of semantic changes.

Figure 7.3 Proportion of semantic changes identified by REPENT.

Renamings that *preserve meaning* are quite unusual in the analyzed programs. Table 7.4 shows detailed results for this category of renamings for Java programs. There is a number of renamings classified as *spelling error* as well. As expected, most renamings correct spelling errors while only a small number introduce spelling errors. 301 out of the 364 spelling error renamings are simple renamings indicating that spelling errors are corrected in isolation. Some renamings aim at correct spelling errors but even after multiple corrections, the identifiers still contain spelling errors (*e.g.*, `defferedSyntaxAllowedAsLitteral` $\rightarrow$ `deferedSyntaxAllowedAsLitteral` $\rightarrow$ `deferedSyntaxAllowedAsLiteral`) because only few available IDEs (*e.g.*, EMACS, ECLIPSE) provide support for spell-checking of identifiers.

We expected that renamings towards *expansions* would be performed for clarification purposes, *e.g.*, `getAlg` $\rightarrow$ `getAlgorithm`. 56% of such expansions concern Java local variables, which indicates that entities with limited scope are also important and developers take care of them. However, the overall number of renamings towards expansions is low (209): 49% of the surveyed developers would probably not undertake a renaming if the name of an entity contains an abbreviation or an acronym; 7% would definitely not rename; 30% were

undecided; and only 13% would probably rename. As for abbreviations, we expected that *abbreviation* renamings would occur when identifiers are long and are composed of many terms. Yet, in more than 75% of such renamings, the old names are composed of only one or two terms. For example, the parameter `parameters` in JBoss was renamed to `params`, while the local variable `association` in ArgoUML was renamed to `assoc`.

REPENT identified three cases of *synonym phrase* in Java programs only. Two fields were renamed from `NOT_CLOSED` to `OPEN`; the third renaming is a false positive. We also manually found an example in Eclipse-JDT, where `javadocNotVisibleReference` was renamed to `javadocHiddenReference`. REPENT failed to correctly classify this renaming because the Stanford NLP wrongly assigns the negation relation (due to the term `Not`) to the term `javadoc` instead of assigning it to the term `Visible`.

Renamings that *change meaning* are the most frequent in Java programs. In general, 33% of the Java renamings aim at changing the meaning of the identifiers. Such renamings are particularly frequent in dnsjava—48% of the semantic changes. As explained in Section 7.2.2, dnsjava underwent a massive renaming (*e.g.*, `rname` becomes `_name`). Those cases are classified as *change meaning* as REPENT fails to relate the meaning of the two terms due to the non-use of separator in the case of `rname`. Here, REPENT would benefit from a more sophisticated splitting technique that would split the identifier into two terms, *i.e.*, `r` and `name`. ArgoUML also underwent a massive renaming activity, due to the use of Log4J as explained in Section 7.2.2. While in the context of ArgoUML, REPENT classifies such renamings as *change in meaning*, we suspect that for the developers of the third-party library (Log4J) the terms `Category` and `Logger` have the same meaning, the former being a superclass of the latter. If this is indeed the case, a domain dictionary would improve the classification.

Table 7.5 shows the results of *change meaning* renamings at a fine-grained leve for Java programs—according to the proposed taxonomy. In general, there is no semantic relationship, *i.e.*, *unrelated* according to taxonomy, between the renamed terms in this category.

Table 7.4 Preserve meaning renamings as classified by REPENT.

	Synonym	Synonym phrase	Spelling error correction/introduction	Expansion	Abbreviation	Total
ArgoUML	10	0	12	44	11	77
dnsjava	0	0	1	9	2	12
Eclipse	163	1	137	61	51	413
JBoss	195	2	180	84	119	580
Tomcat	185	0	34	11	29	259
Total	**553**	**3**	**364**	**209**	**212**	**1,341**

As an example, in ArgoUML, REPENT identified that a parameter name was changed from `eventNames` to `propertyNames` to reflect the new semantics of the parameter. However, although quite rare, there are cases where the developers inverted the responsibility of an entity, *e.g.*, in JBoss REPENT identified that a method name was changed from `isInvisibleAnnotationPresent` to `getVisibleAnnotation` to reflect the new behavior of the method. REPENT identifies only two identifiers where the names changed from `body` to `node`, *i.e.*, renamings where the semantic change is a whole-part phrase. This type of renaming is less likely to occur.

Particularly interesting are renamings that involve identifiers that contain a negation. Such identifiers are usually renamed towards positive names; this is a particular example of *opposite phrase* renamings identified by REPENT. For example, in Eclipse-JDT the method `isNotPrimitiveType` was renamed to `isPrimitiveType` and the local variable `dontSetFigs` of ArgoUML was renamed to `setFigs`. From the analysis of the entities involved in such renamings, we observed that they are usually used with the negation operator. In such cases it is more difficult to interpret an expression containing such entities, especially if the expression contains a Boolean negation operator. However, among the surveyed developers only 30% would rename an entity if the name contains a negation.

There is a substantial number of renamings classified as *narrow* and *broaden meaning* in Java programs. For example, the method `testEJB3RemoteAccess` of JBoss was renamed to `testRemoteAccess` to emphasize a more general behavior of the involved entity. A similar example is represented by the method `getServletRequest` of Tomcat that was renamed to `getRequest`. There are also cases where the identifier was made more specific. For example, the method `isRemoteInvocationExecutedInNewThread` of JBoss was renamed to `isRemoteAsyncInvocation ExecutedInNewThread` to highlight that the remote invocation is asynchronous. A similar example is represented by the renaming `type` → `authType` in Tomcat.

There is also a high number of renamings that *add* or *remove meaning.* An example of

Table 7.5 Change in meaning renamings as classified by REPENT.

	Opposite	Opposite phrase	Whole-part	Whole-part phrase	Unrelated	Total
ArgoUML	0	2	0	0	1309	1,311
dnsjava	0	0	0	0	576	576
Eclipse	44	29	0	0	5449	5,522
JBoss	29	38	0	0	5975	6,042
Tomcat	16	7	2	0	1036	1,061
Total	**89**	**76**	**2**	**0**	**14,345**	**14,512**

adding a meaning is `delete → removeFromDiagram`, whereas an example of remove meaning is `addRecord → add`. Although these two kinds of renamings cover about half of the renamings identified by REPENT, it is worthwhile to point out that the interpretation can be subjective. Some of the examples may be classified differently by different people, *e.g.*, *narrow meaning* rather than *add meaning*.

Finally, 5% of the renamings contain no semantic change, *i.e.*, are classified as *none*.

Our qualitative analysis confirms that in many renamings the goal of developers when performing renamings is to improve the comprehensibility of identifiers. We observed that most of these renamings are performed to increase the consistency between the name of an entity and its functionality, or between an identifier and other identifiers. This goes along with the high number of survey participants who would definitely rename an entity when the name and functionality are inconsistent (66%). Specifically, analyzed renamings aimed at improving the consistency with the existing code. For example, the method `isChildOf` in Eclipse-JDT was renamed to `isDescendantOf` as its functionality considers all super types, rather than the direct parent only. Sometimes developers rename identifiers to reflect new functionality represented by an entity. For example, field `typeMapping` in JBoss was renamed to `datasourceMapping`. The analysis of the log message confirmed that name changed to reflect the new functionality: *"Changed type-mapping to datasource-mapping as is required by new dtd."* Another example is the parameter `principal` in JBoss, renamed to `authPrincipal`. Here the renaming was a result of a bug fixing (*"incorrect principal used"*).

> **RQ3 conclusion:** Renamings rarely preserve the meaning of identifiers (less than 3%). Slightly more often, the meaning is narrowed (6%), or broadened (5%). Moreover, renamings with no semantic changes are rare (5%). Most often, renamings change (33%), add (25%), or remove (21%) a meaning.

7.2.5 RQ4: What kinds of grammar changes occur in identifiers when they are renamed?

Table 7.6 and Fig. 7.4 show the proportion of the grammar changes in the five programs.

76% of the classified renamings do not involve a part of speech change, *i.e.*, are classified as *none* in the *grammar change* dimension. When there is a part of speech change however, only 5% of the changes involve a change in *verb conjugation*; 13% involve a *singular/plural* change. The majority of the renamings, *i.e.*, 83%, involve *other part of speech* changes.

One good reason for developers to change *singular* to *plural* and vice versa is to align an identifier with the entity (or collection of entities) to which it refers. For example, a field

Figure 7.4 Proportion of grammar changes identified by REPENT.

or a local variable of a collection type (*e.g.*, `ArrayList`) should have a plural name, whereas atomic entities should have, in general, a singular name. Such inconsistencies have been previously studied and denoted as *linguistic antipatterns* (LAs) Arnaoudova *et al.* (2013) *i.e.*, *recurring poor practices in the naming, documentation, and choice of identifiers in the implementation of an entity, thus possibly impairing program understanding.* That is, a method name containing only singular nouns but returning a collection is a LA of kind "Expecting but not getting a single instance". Similarly, method names containing plural nouns but returning a single object are LA of kind "Expecting but not getting a collection". Similar considerations can be made for fields. In our study, one example of renaming aimed at removing a LA occurred in JBoss, class `BasicMBeanRegistry`, where a local variable named `descriptors` was renamed to `descriptor` because its type changed from `Descriptor[]` to `Descriptor`. A similar case occurred in ArgoUML (class `LabelledLayout`) where a local variable of type `int` named `unknownHeights` was renamed to `unknownHeightCount`, the new name being consistent with the type. In this case, the program only keeps the count of heights rather than the list of heights, as suggested by the old name. In Tomcat, a method of

Table 7.6 Grammar change renamings as classified by REPENT.

	Singular/ Plural	Verb conjugation	Other part of speech change	None	Total
ArgoUML	31	22	608	3,322	3,983
dnsjava	8	0	241	846	1,095
Eclipse	602	198	2,785	8,951	12,536
JBoss	337	125	2,531	10,589	13,582
Tomcat	51	25	479	1,805	2,360
Total	**1,029**	**370**	**6,644**	**25,513**	**33,556**

class `RealmBase` named `findSecurityConstraint` was renamed to `findSecurityConstraints` and its return type changed from `SecurityConstraint` to `SecurityConstraint[]`.

Examples of *other part of speech changes* include the parameter renaming `localDeclaration` → `location` and the method renaming `deployOnMember` → `doDeployment`. There were renamings where although the part of speech changes, the role played by the renamed term remains the same. Examples are the method renaming `getTreeCache` → `getClusteredCache` and the field renaming `_multiPane` → `_editorPane` where although the part of speech of the renamed terms changed from noun to adjective and from adjective to noun respectively, the role played by the terms before and after the renaming is the same, *i.e.*, modifier of the term after, *i.e.*, `Cache` and `Pane` respectively.

REPENT reports 370 *verb conjugation* changes. In 32% of such changes the term `is` is renamed to `get` or `has` (or vice versa). The part of speech that the tagger assigns to the verb `is` and `has` is different from the part of speech of verb `get`. When using automatic code generator tools (such as the Eclipse IDE) for generating getters and setters, the names of getter methods returning a boolean value start with `is`. This is also the recommendation of JAVA/J2EE naming conventions. This could be a possible reason to rename getter methods of Boolean fields to start with `is`. One such example is the renaming `getValidProject` → `isValidProject`, where the return type of both methods is Boolean. The rest of verb conjugation changes are change in the verb tense (past to present or vice versa) or changes of a verb to gerund. Examples include method renamings such as `methodNeedingAbstractModifier` → `methodNeedBody` and `isOverridden Method` → `areOverriddenMethods`.

> **RQ4 conclusion:** With respect to the grammar changes, 76% of the classified renamings did not involve a part of speech change. Of the 24% of the renamings involving a part of speech change, a small proportion involved a verb conjugation change (5%); 13% involved changes in nouns (singular/plural); 83% involved other part of speech changes.

7.3 Threats to Validity

This section discusses the threats to validity that can affect classification.

Construct validity threats concern the relationship between theory and observation. The construct validity threats are related to (i) the precision and recall of the set of detected renamings on which the classification is performed, and (ii) the accuracy of the automatic classification. Concerning the former, results of the study reported in Section 5.1 provide an indication of such precision and recall. Concerning the latter, we performed—using a process similar to the one described above—a manual validation of a sample of the classified renamings, to provide an idea of how accurate such a classification is.

Conclusion validity threats concern the relationship between the experimentation and the outcome. As mention before, our study is an exploratory study in which we do not make use of statistical tests to reject specific hypotheses. The only issue related to conclusion validity is the representativeness of the sample used to validate the classification accuracy. We performed a random sampling for all dimensions and levels of the taxonomy considering a confidence level of 95% and a confidence interval of at least ±10%.

External validity threats concern the generalizability of our results. Although we have chosen a pretty variegated set of projects—belonging to different domains and development organization, and having different size—it could happen that replicating the study on other projects could lead to different results, different distribution of the classified renamings with respect to the proposed taxonomy. A different matter is the application of the proposed approach on programming languages different from Java. For example, differently from Java identifiers, C/C++ identifiers are more likely to contain abbreviations Madani *et al.* (2010); Guerrouj *et al.* (2011), and would rarely use camel case or other explicit term separators. As explained in Section 6.1, this would require the use of appropriate identifier splitting or normalization approaches Enslen *et al.* (2009); Guerrouj *et al.* (2011); Lawrie et Binkley (2011); Corazza *et al.* (2012).

7.4 Summary

In this chapter we provide the results of renaming classification for the five Java programs. The results show all dimensions of the taxonomy are covered by renaming instances some with more occurrences than the others and thus answer our thesis on how identifiers are renamed. Moreover, we see that methods and local variables are renamed more than the other entities. Though we found some rare case of renamings that introduce spelling error, we believe that spelling error correction, and term reordering, are evidence that renamings' purpose is to improve the quality of identifiers. We see that REPENT has high precision in classifying the renamings in *forms of renaming* as well ad *semantic changes*. For *grammar changes*, REPENT performs fairly accurate and the miss classifications are due to application of the tool designed to tag natural language text. In the following chapter we provide a preliminary study of extending REPENT for detection and classification of renamings in PHP programs.

CHAPTER 8   CHALLENGES WITH DYNAMIC LANGUAGES

Java shares principles of object-oriented programming with other languages such as C++, C# and Delphi. Moreover Java is statically typed language. It is interesting to further investigate the applicability of the proposed methodology for other programming languages that are very different from Java. Popularity of dynamic languages are increasing [1], (Paulson, 2007). Perl, Python, JavaScript, PHP, and Ruby are among the popular dynamic languages. We choose PHP as the vast majority of websites use PHP [2]. The rest of this chapter is organized as followings: First we briefly discus the related works on PHP programs. Next, we explain challenges for adapting the propose methodology for detection and classification of renamings in PHP programs. Then we explain how we addressed some of the challenges and finally we report the preliminary results on detection and classification of renamings in PHP programs.

8.1   Related Work

The following covers a number of areas on: PHP analysis, and Web application reverse- and re-engineering.

8.1.1   Web Application Reverse Engineering

While our overarching goal is clearly different, certain commonality can be found with the reverse engineering of WEB applications (WAs), in particular static and dynamic analysis. The first significant contribution was given by Ricca and Tonella, who developed the REWEB tool to perform analyses on web sites (Ricca et Tonella, 2001a,b). In particular, Ricca and Tonella introduced a graphical representation of the web site to allow for traditional static flow analyses such as reachability, dominance, and data flow analysis on WAs. The same authors also proposed to enhance static analysis by using dynamic information (Tonella et Ricca, 2002). Clearly, REWEB does not need page instrumentation; on the other hand, web server logs, do not allow fine-grained analyses such as needed to detect of a variable is being assigned a different types in two dfferent execution paths.

Di Lucca *et al.* (2002, 2003a,b) proposed an approach and a tool to extract Conallen's UML documentation, use cases and business object from Web applications. Their approach uses

---

static analysis, however they pointed out that diagrams can be refined using dynamic information. WARE performs static analyses on WAs, stores the extracted information into a database and then uses such an information for the reverse-engineering of UML diagrams.

Architectural recovery was also the goal of the works of Hassan et Holt (2002) and Antoniol *et al.* (2004). Both teams reverse engineered high level views of the WEB application.

8.1.2 Analysis on PHP application

Nguyen *et al.* (2011) propose an automated approach, PHPSYNC, for detecting validation errors in HTML files of PHP application. PHPSYNC uses and HTML validator tool, Tidy, to detect errors in the HTML page and map the error-some HTML fragments to PHP page. PHPSYNC builds a tree model, D-model, that represents the symbolic execution of the server-side PHP code. Then it maps text in a given HTML page to the D-model. Using Tidy, PHPSYNC is able to map the fragments contains error to the D-model. If Tidy fixes the validation error in HTML page, PHPSYNC replaces the corrected HTML into D-model.

Nguyen *et al.* (2013a) developed a tool, WEBDYN, for dynamic refactoring of PHP Web applications. They manually analyzed 2,664 revisions of four open-source PHP-based Web applications, and found that there exists an special form of refactoring that is specific to dynamic Web applications. Next, they categorized these refactorings (which they called output-oriented refactoring operations) in five groups: 1) dynamicalization (*e.g.*, replacing inline HTML/Javascript code with a PHP fragment or function), 2) re-structuring server and client code, 3) renaming embedded HTML/Javascript elements, 4) standardizing embedded HTML code, and 5) refactoring for separation of concerns. They use dynamic analysis coupled with symbolic execution to identify variable declarations, references as well as dangling references.

Merlo *et al.* (Gauthier et Merlo, 2013) detect semantic smells and errors in access control of PHP application using static analysis, model checking, and information retrieval technique. They defined semantic smell as poor implementations of the semantic of an access control model and semantic errors as wrong implementations that need to be corrected. They assume that semantically related sections of source code should be protected by similar privileges. Using model-checking technique and static analysis, they first extract mappings between source code fragments and the privileges. Next, they applied Latent Dirichlet Allocation to extract topics hidden in code. Using logistic regression, they find related topics to the extracted privileges. Finally, they infer privileges for blocks of code. Semantic smells are identified when a semantically inferred privileges are different form the enforced privileges. They applied the proposed technique on a medium size open source application (Moodle

2.3.2) with 307 privileges. They find 31 semantic smells and two errors. The found errors were confirmed by the developers of Moodle, and action were taken to resolve the errors.

Nguyen *et al.* (2013b) proposed an automatic approach, DRC, to identify dangling reference errors in PHP programs using static analysis of source code. DRC applies symbolic execution of PHP programs to identify variable declarations and references. For each detected declaration or reference, DRC associates it with the current path constraint of the symbolic execution. To identify the declaration and references of entities embedded in HTML or SQL script within the PHP code, DRC uses the tree-based representation, called D-model (described above). Next, for all variable references, DRC identifies a declaration that matches the reference.

Gauthier *et al.* propose inter-procedural and intra-procedural algorithms in Datalog for propagation of pattern-based properties such as permissions in access controls of PHP programs (Gauthier et Merlo, 2012). They extract relation and rules from AST of the PHP program. Assignment of patterns and variables to variables are translated to Datalog relations. Rules are defined to infer "must hold" pattern relation recursively. The results show that security checks are detected with high precision in eight open-source PHP programs.

8.2   Challenges and adaptations

Since PHP is very different from Java we need to adopt our approach to be able to detect and classify renamings. In the followings we briefly discuss the challenges for such adaptation.

Taxonomy: The first dimension of taxonomy is the programming entity. PHP is scripting, procedural, and object-oriented language, with the following programming entities: Namespace, class/interface/trait, function/method, parameter, field, and variables (local, global, constants). Though namespace in PHP and package declaration in Java are different we map them as both provide a mechanism to support organization of source code as well as enabling accessing controls. Class, interface and method share same concepts and principles in both programming languages. PHP is a dynamic type language and thus fields, parameters and variables are not bound to types. While parameters and fields need to be declared to be used, variables in PHP does not require a declaration. We consider all above mentioned entities for detection of renamings in PHP programs.

Our taxonomy is defined at first place through manual analysis of several identifiers renamings in Java programs. We believe that the other three dimensions of taxonomy are language independent; however, we need empirical data to support our claim.

Renamings detection: The core of detection component is extraction of programming entities and their def/use statements. We need to parse PHP code to exact entity declarations, and for variables we consider the assignment statement as declaration statements. To extract def/use of the entities we need to perform data flow analysis. One crucial difference between Java and PHP, is the way to access entities defined in other files. In Java entities defined in other files or packages are accessible through `import` mechanism. In PHP it is done through `include` statement that accepts as a parameter an expression. There are no constraints on the include expression which can contain variables and calls to functions as well as string operators. In other words, often, the file name path is dynamically computed and built at run-time. Unlike Java, there is no restriction on where to include a file which makes it more challenging to track variable's def/use statement.

For comparing def/use statements of mapped entities we calibrated the thresholds on one Java program and applied it for renamings detection of other programs. We may need to re-calibrated the thresholds to achieve desired accuracy for detection of renamings in PHP programs.

Renamings classification: Identifier splitting is the first and crucial step in renamings classification. Unlike Java, there is no consensus among PHP developer for naming programming entities, and rather it is framework and project dependent *e.g.*, Zend[3] does not permit underscores, and Wordpress[4] encourages underscores instead of camelCase. Moreover, PHP is case insensitive and thus both names `getSessionID` and `getSessionid` can be used to invoke function `getSessionId`, and thus we expect less renamings due to change of letter cases in PHP programs. The project/framework dependent rules on naming may impact the performance of renaming classifier. Moreover, we may need to re-calibrate the threshold used for classification of renamings under spelling correction as well. The following sections explain how we addressed some of the challenges discussed above for applying REPENT on PHP programs.

8.3   Resolving file inclusion in PHP programs

Before performing deeper analyses such as data flow analysis, we need to resolve includes. In PHP the inclusion is done with the `include` statement that accepts as a parameter an expression. As mentioned before, there are no constraints on the include expression which can contain variables and calls to functions as well as string operators. In other words,

---

3. http://framework.zend.com/manual/1.12/en/coding-standard.naming-conventions.html
4. https://make.wordpress.org/core/handbook/best-practices/coding-standards/php/

often, the file name path is dynamically computed and built at run-time. To avoid circular inclusions, PHP provides different include statements, namely `include`, `include_once`, `require` and `require_once`. `include` and `require` always include the file passed as parameter. The difference is that `require` produces a compiler error upon failure. `include_once` and `require_once` work similarly to `include` and `require`, however do not include a file if it has been included already, *i.e.*, avoiding a multiple inclusion.

Consider the example in Figure 8.1. The PHP `define` statement takes two arguments: a string identifying the constant to be defined, and an expression to be evaluated and assigned to the constant (the right hand side). In Figure 8.1 the constant `CWD` is defined in f_1 by the statement `define( 'CWD', '/' )` and `'/'` is its value. In the same example we see file f_1 include file f_2 through the include statement `include (CWD."f2".".php")` where the path to the file to be included is the result of the concatenation of three strings *CWD*, *"f2"*, and *".php"*, where the constant `CWD` is already defined in previous line.

Moreover, If included files are not resolved, one would think that variable *pos* is first encountered and thus defined in file f_1, while actually it is declared/defined in file f_2.

8.3.1   A Fixed-Point algorithm to resolve include

The approach for resolving include statements is based primarily on a static analysis, complemented by dynamic analysis.

To analyze PHP source code, we rely on two widely-used infrastructures, the PHP parser from the Eclipse PHP Development Environment (PDT) [5], and TXL (Cordy, 2006). Specifically, we use the Eclipse PDT parser to assist in the static analysis, and the TXL source transformation engine to add the code instrumentation needed to collect information at run-time.

First, we use the PDT parser to create an Abstract Syntax Tree (AST) for each PHP file. We build a symbol table to collect information about PHP entities namely files, classes/interfaces, methods/functions, and constants. Algorithm 2 presents a high-level view of the computation performed to resolve includes and constants. First, all files in the system are traversed and entities of interest are added to the symbol table. For constant entities we collect their values (*i.e.*, right end side) if not dynamically built at run-time. If the value is not a scalar but it is a concatenation of strings and/or return value of functions, we collect the whole statements and mark the entity as non resolved. In the same way we process the include statements (`include`, `include_once`, `require`, `require_once`) as well. We store in the symbol table

---

5. http://projects.eclipse.org/projects/tools.pdt
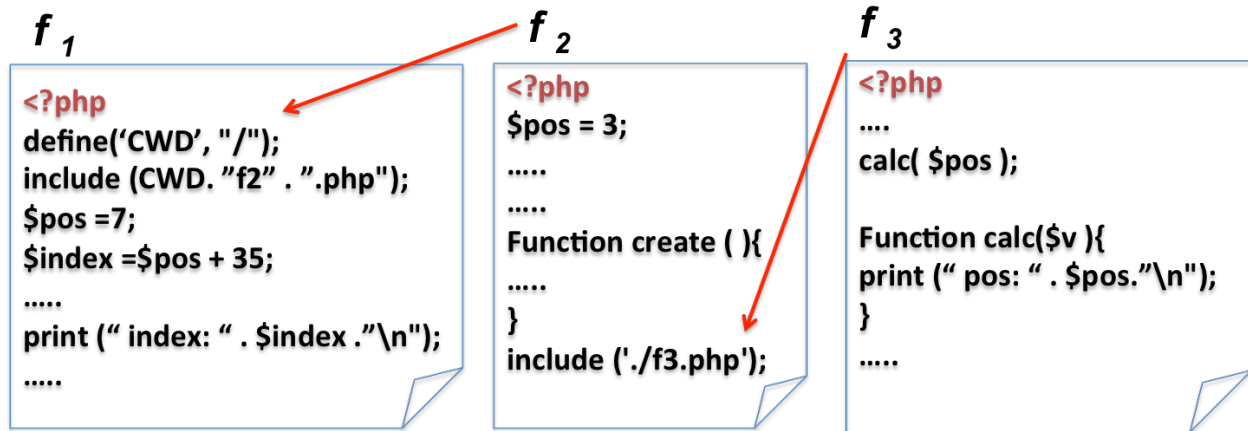
Figure 8.1 Example of include in PHP.

structure also the statement parse tree, later used by the symbolic execution. To statically resolve constant and include statements, we implemented a simplified symbolic execution. The simplified symbolic execution implements the behaviour of two PHP magic variables and two functions as well as the string concatenation operator (See Table 8.1). We rely on API function provided through `FileNetworkUtility` in PDT to handle the access to the parent (current) directory (*e.g.*, "../" and "./").

Once files are traversed and the symbol table initially populated, the fixed-point algorithm attempts to resolve as many include file paths as possible. Before this step, we apply a simple heuristic using the function
ScanAndAssignNameIfUnique. It may happen that the include path ends with a file name (*e.g.*, `'ABSPATH' . ngg-config.php`) and it may happen that this file name is actually unique and not the suffix of any other file. In such cases, there is no need to perform any

Table 8.1 Implemented functions, operators and Magic Variables

Function/ Operator/ Variable	Description
dirname	Extract directory name
basename	Name with any leading directory components removed
.	Concatenate two strings
___FILE___	Contains the current file absolute path
___DIR___	Equivalent to dirname(___FILE___)

Algorithm 2 Include resolution and symbol table construction.

$update$=false
for $file$ in $System$ **do**
    $entities = extractEntities()$
        $addToSymbolTable(entities)$
end for
for ent in $SymbolTable$ **do**
    if $ent == Const|ent == Include\&\neg ent.resolved$ **then**
        $ScanAndAssignNameIfUnique(ent)$
    end if
end for
repeat
    for ent in $SymbolTable$ **do**
        if $ent == Const|ent == Include\&\neg ent.resolved$ **then**
        $update| = SymbExecAndUpdate(ent.stmnt, SymbolTable)$
        end if
    end for
until $update$

---

complicated calculations and the include relation is resolved based on the file name identity.

Next, as shown in Algorithm 2 if the current entity is an include or a constant, and (so far) it has not been resolved *i.e.,* the value is not known, the function *SymbExecAndUpdate* attempts to compute the current define right end side or include parameter value. If *SymbExecAndUpdate* succeeds, it returns true so that the collected new information can percolate and improve collected information, as well as updating the values of other entities depending to the current entity. If *SymbExecAndUpdate* fails, the entity will be marked as delayed, and its value may be resolved in following passes. In fact, if a new include file path (or constant) is resolved this may impact on variable definition as well as on other constants or includes.

Algorithm 2 terminates when it is not able to further update any information (*i.e.,* when it has reached a fixed-point). In general, it may happen that, at this stage, some constants or includes are still not resolved, because the path of some files is still unknown. This typically happens for two reasons: (i) the right-hand side of the string constant definition or the include parameter contains one or more PHP variables; or (ii) the string value is obtained from a user defined function.

The remaining unknown includes make the analysis imprecise. In general, the way the analysis is performed avoids the presence of false positives, but cannot exclude false negatives. Typically—as it also happened in our study (see Section 8.3.2)—most includes are statically resolved even before the fixed-point algorithm, and thus in principle, a good quality analysis

can be done without the need for dynamic analysis. However, if a complete analysis is needed, the only possible solution is to resort on dynamic analysis to track the remaining unresolved relations. Clearly, dynamic analysis requires deploying and executing the application, and the design of scenarios that exercise the unresolved includes.

Once the includes have been resolved, we build a directed graph for the whole application.

8.3.2 Case study on resolving include

The *context* of the study consists of Wordpress (WP) itself—in two releases, namely 3.6 and 3.7—along with a set of 10 installed plugins. We did not consider the most recent versions of WP (*e.g.*, 3.8) because many plugins were not tested with that release. For what concerns the selection of plugins, we focused on 10 popular ones. There is no general consensus on the top ten must-have plugins; depending from the Web site's application domain, user communication goals and project constraints, different plugins may better serve the Web site's objectives. Indeed, there are several different lists of the top ten must-have WP plugins on the Web. Based on the WP most popular plugins [6] and other three most popular lists, namely WeDesignPixel [7], TreeHouse Blog [8] and Selz Blog [9], we have chosen the set of plugins reported in Table 8.2. As the table shows, we performed our study on two configurations of WP: (i) an old one, consisting of WP 3.6 plus some old versions of the 10 plugins, and (ii) a more recent one, consisting of WP 3.7 plus some newer versions of the plugins.

Table 8.2 reports a summary of the include relations in both releases of WP and their corresponding plugins.

To simplify the computations we assume that we know the values of the constants: `ABSPATH` (the path where WP is installed), `WPINC` (location of include directory), and `WP_PLUGIN_DIR` (location of plugin directory). As Table 8.2 shows, most of WP includes are resolved before the fixed point step; this points to a disciplined and not overly complex include file regimen in the core framework. However, for plugins the situation changes. Consider the plugin `W3 total cache`. This plugin has almost the same number of include relations as WP itself, but almost half of the included file relations are resolved only by the fixed point step. On summary, only 8.2% (number of unknown after fix point over the total number of include relations) cannot be identified statically.

When looking at release 3.7 of WP and its related plugins, the situation does not change

---

6. http://wordpress.org/plugins/browse/popular
7. http://wedesignpixel.com/top-must-have-wordpress-plugins
8. http://blog.teamtreehouse.com/best-free-wordpress-plugins-for-common-website-functionality
9. https://selz.com/blog/10-must-wordpress-free-plugins-starting-online-business-selling-digital-downloads/

dramatically, and overall only about 9% of the include relations are not statically computed.

Table 8.2 Analyzed releases of WP and its plugins, with details about include relations.

	Old Release						New Release					
	Rel.	Includes	Unknown Before Fix Point	Unknown After Fix Point	Edges	Nodes	Rel.	Includes	Unknown Before Fix Point	Unknown After Fix Point	Edges	Nodes
WordPress	3.6	629	37 (6%)	35 (5%)	627	366	3.7	647	37 (6%)	35 (5%)	645	368
NextGen Gallery	1.9.3	114	26 (23%)	12 (10%)	114	71	2.0.40	144	37 (26%)	22 (15%)	144	149
Google XML Sitemap	3.2.7	5	2 (40%)	2 (40%)	5	7	3.3.1	5	2 (40%)	2 (40%)	5	7
Contact Form 7	3.2	16	15 (94%)	4 (25%)	16	19	3.6	19	18 (95%)	5 (26%)	19	20
Akismet	2.5.6	3	0 (0%)	0 (0%)	3	4	2.5.9	3	0 (0%)	0 (0%)	3	4
SEO by YOAST	1.1.7	22	18 (82%)	2 (9%)	22	22	1.4.22	42	35(83%)	2 (5%)	42	40
WP Sitemap Page	1.0.12	1	1 (100%)	1(100%)	1	2	1.0.12	1	1 (100%)	1 (100%)	1	2
Google XML Sitemaps for qTranslate	3.2.7.1	6	2 (33%)	2 (33%)	6	8	3.3.1	6	2 (33%)	2 (33%)	6	8
YARPP	3.5	17	16 (94%)	5 (29%)	17	17	4.1.1	26	23 (88%)	4 (15%)	26	25
Jetpack	2.7	95	37 (39%)	15 (16%)	95	108	2.3.5	126	63 (50%)	20 (16%)	126	143
W3 Total Cache	0.9.2.4	592	335 (56%)	45 (8%)	593	332	0.9.3	436	168 (38%)	33 (7%)	436	299
Total		1,500	489 (33%)	123 (8%)	1,499	956		1,455	386 (26%)	126 (9%)	1,453	1,065

Limitation of approach: Although, as explained, we used static analysis to resolve includes, we have validated it by means of dynamic analysis. Clearly, a thorough dynamic analysis would have required to exercise all possible paths of the Web application that alter the values of include file names. In our analysis, we focus on five execution scenario, to see how we could have discovered some includes not resolved statically.

By comparing the log of dynamic analysis and the include relations that are extracted statically we identify a subset of unknown relations. Table 8.3 indicates the number of unknown includes that were discovered by these five scenarios in WP 3.6 and its corresponding plugins. The second column of the table shows the number of unknown files that are resolved dynamically. Overall, by executing these scenarios 26% of all unknown include relation were identified.

Table 8.4 summarizes all cases in which the static analysis analysis fails. In some cases the include could not resolved because the string passed as a parameter to the `include` function was produced as output of a function call, and the simplified symbolic execution could not resolve that. In other cases, the file name could not be produced because this would have required a symbolic execution able to support a context- and flow-sensitive data flow analysis. In summary, very likely a more sophisticated symbolic executor could have further improved the completeness of the include resolution. However, since the unknown includes represent only about 9% of the total include relations, a lightweight analysis like the one we proposed is appropriate and able to scale up to the size and complexity of WP with its plugins. The study is limited to two releases of WP, and to a subset of its plugins. Although we expect that similar problems can occur with other plugins and, possibly, with other PHP frameworks besides WP, further, larger studies need to be conducted to verify such a conjecture.

Table 8.3 Unknown includes resolved in WP 3.6 and its plugins by means of some dynamic analysis.

Scenario	Discovered Unknown	%
1	17	(13%)
2	9	(7.3)%
3	31	(25%)
4	13	(10%)
5	12	(9.7%)
Overall	33	(26%)

Table 8.4 Where static analysis fails.

Program	Path contains	
	variables	function calls
WP 3.6	73	12
WP 3.7	89	8

8.4 preliminary study of renamings in PHP aplication

In this section we report the application of REPENT for detection and classification of renamings in PHP programs. We extended REPENT's detection component specifically the pre-processing which relies on Java parser (see Figure 4.1). We use classification component of REPENT in the same way as it was used for classification of Java programs.

Figure 8.2 illustrates the detection and the classification for PHP application. In the following we explain the steps that are different from Java renaming detection.

8.4.1 Renaming Detection

The first step of the process is to build the line mappings between two source files. For file renaming, REPENT takes the same approach as for Java file renamings. For PHP systems, the change set is computed by files deleted and added in the same day to the Git repository. Then each couple is compared using the same technique with the same threshold.

Line differencing algorithm is then used to generate line mappings. We use Eclipse PHP Development Tools (PDT) to parse the code and extract the abstract syntax tree (AST). We traverse the AST to build control flow graph of PHP files. Using Eclipse PDT we resolve the method and type binding. The binding information will be used in data flow analysis. The AST of PHP file does not contains variable declaration, thus we take every assignment
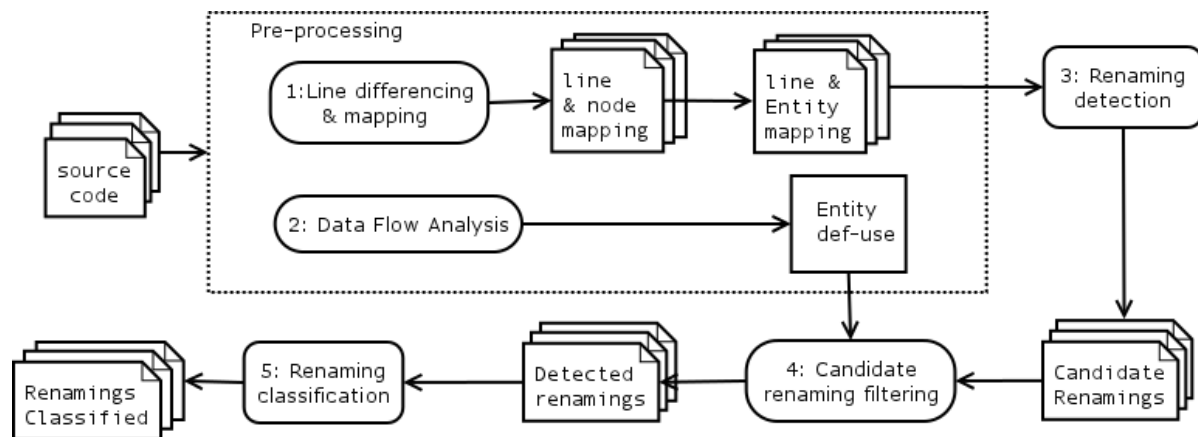
Figure 8.2 PHP renaming detection and classification.

where the left hand side (LHS) is a variable as variable declaration. Using the line mapping we map nodes of the same type. Next we use the node mapping to extract entity mapping. Figure 8.3 shows an example of line mapping.
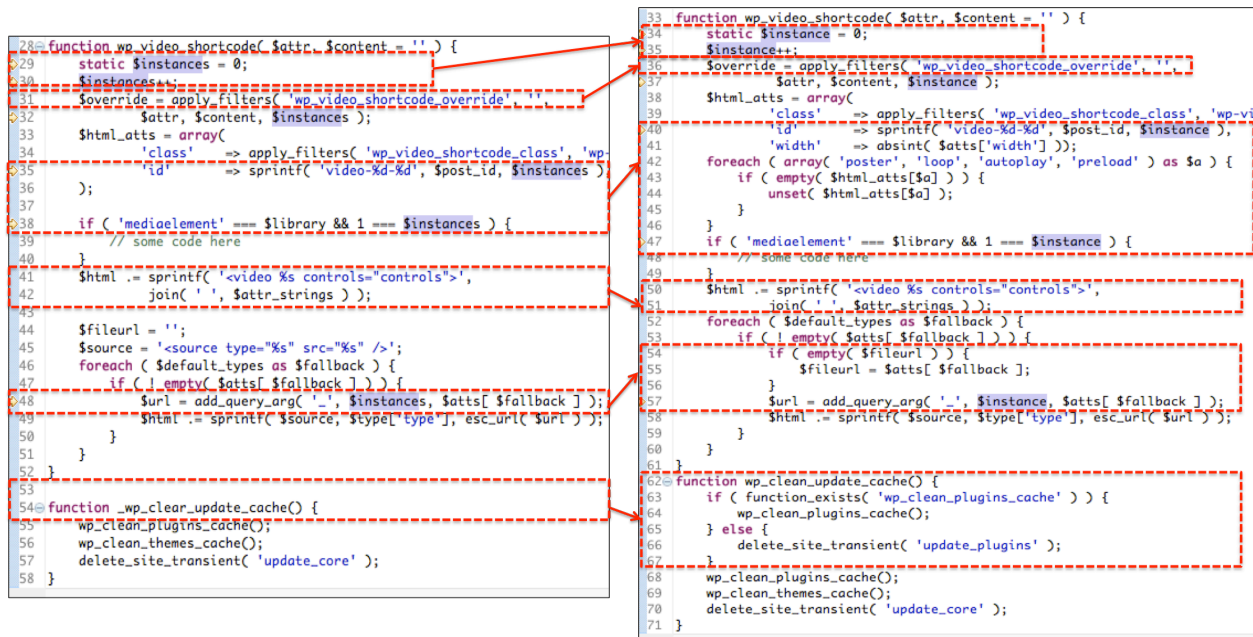
Figure 8.3 Example of line mapping.

For example line 48 in old version is mapped to lines 54-57 which corresponds to a one-to-many line mappings. Using the line mapping we map the cfg nodes within those mapped lines. $\{assignment\_statement\} -> \{if\_statement, assignment\_statement, assignment\_statement\}$. We first group the nodes with same type in both old and new version and then we extract assignments from each group to map the variable declarations. For this example we will have $\{\$url\} -> \{\$fileurl, \$url\}$. As the mapping in Figure 8.3 shows, lines 53-54 in old version is mapped to line 62-67 in new version, that gives us the node mapping: $\{function\_declaration\} -> \{function\_declaration, if\_statament, function\_invocation, function\_invocation\}$. Form these node mapping we can map the function declarations only (one-to-one-mapping) that corresponds to mapping $\{\_wp\_clear\_update\_cache\} ->$

{*wp_clean_update_cache*}.

At the end of this step we will have a list of candidate renaming that contains one-to-one, one-to-many, many-to-one and many-to many mappings.

Candidate Renaming Filtering

As in our previous analysis we need the data flow analysis to filter false positives from the candidate renaming list. For each entity we extract statements in which the entity is used and redefined (for parameter, local and global variable, and constants). Using the type, method, and function bindings we extract statements in which entities of type name-space, class/interface method/function and field are used. For local variable we use intra-procedural control flow analysis on the cfg of each method and function to extract the use and redefinition of method/function parameters, and the local variables. For the variables at file level we perform flow sensitive context insensitive analysis on the cfg of each PHP program. That is, if a variable is being defined global in a function or method, we collect its uses and redefinition from the function/method bodies as well. By connecting cfg of PHP files through resolved include statements and resolved function/method calls we build the cfg of the whole program. Then we collect the def/use statements for global variables through flow sensitive context insensitive data flow analysis.

Computing the Score between Entities

Regarding comparison of mapped entities' def/use for calculating the scores, REPENT applies the same steps and thresholds used for Java programs. The only difference here is that for parameters and fields with no def/use, we use their names, as there is no declaration statement for such entities. In Java systems if we have a mapping between two fields with no def/uses, we compare the field declaration statements which include: the modifiers, access visibility, and the type declaration of the fields being mapped. In PHP, we do not have type declaration and in majority of the cases fields (and parameters) do not have modifiers. Thus if we follow the same approach as for Java systems, we would end up comparing the access visibility of the fields and for parameter we have nothing to compare. Thus, we consider the names of fields and parameters together with the modifiers (if exists) in cases where there is no def/use for such entities.

8.4.2 Results of Renamings Detection

We have applied REPENT on three PHP open-source programs (Wordpress, Drupal, and phpBB). We analyzed the daily file commits for a period of 30 consecutive days. Wordpress [10] and Drupal [11] both are content management systems that allows users to easily build and customize complex Web applications. phpBB is an internet forum application that allows the user to create a very unique forum in minutes, it has variety of styles and support multiple database engines [12].

Table 8.5 reports the main characteristics of the analyzed programs: the analyzed time periods, size ranges in KLOCs, numbers of files, numbers of analyzed revisions, and numbers of committers. All PHP programs are versioned under Git repository.

Table 8.6 shows the detected renamings in PHP programs for different entities. We see that all except one function renamings are in Drupal, and it seems that changes are due to upgrades from version seven to eight with the log message: *bfroehle: remove all 7xxx update functions and tests (D6 to D7 upgrade path).*

Precision of PHP programs: the number of detected renaming for PHP application is very small as only 30 days of commits were analyzed. Using the *gitdiff*, we can extract the added and deleted lines for each committed file, and an evaluator manually verified the changes for all 1,067 files to build our oracle. In case of doubt, second opinion was sought. Then by comparing the oracle and the detected renamins, we identify the TPS as the true positive set and FPS as false positive set. The precision is computed using the same formula for computing precision of detected renamings of Java programs:

---

10. `https://wordpress.org/`
11. `https://www.drupal.org/`
12. `https://www.phpbb.com/`

Table 8.5 Characteristics of the analyzed PHP programs.

Program	Analyzed period	KLOCs (range)	Files (total)	File revisions	Committers
Wordpress	6-Mar-2015 to 6-Apr-2015	263-265	547	386	15
Drupal	19-Mar-2011 to 30-May-2011	227-242	492	322	2
phpBB	18-Jun-2006 to 21-Jul-2006	90-97	143	368	5

Table 8.6 Renamed entities identified by REPENT - PHP programs.

	Type	Field	Constructor	Function/Method	Parameter	Local variable	Total
Wordpress	1	0	0	0	1	8	10
Drupal	0	0	0	24	0	7	31
phpBB	0	0	0	1	0	17	18
Total	**1**	**0**	**0**	**25**	**1**	**32**	**59**

$$Pr = \frac{|TPS|}{|TPS| + |FPS|}$$

The overall precision is 85% and we can see that the entity renamings are not so divers as in Java programs. We did not find any name space declarations and thus we don not expect name space renamings. The lowest precision is for variable renmaings in Drupal. We observed that the impressions comes from line mapping and comparison of declarations when variables have no uses.

Recall for PHP programs: Using the oracle we compute the recall for each entity in PHP programs. Table 8.8 reports—for each kind of entity—the true detected renamings and the recall. We used the same thresholds used for detecting renamings in Java programs. This could be a reason for low recall for local variables.

8.4.3 Results of Renamings Classification

For PHP programs, we manually analyzed all the classified renamings. With respect to the classification of *forms of renaming*, REPENT has precision of 100% (see Table 8.9). We

Table 8.7 Precision *Pr* for renaming detection of different entities

	Wordpress	Drupal	phpBB	Overall
Type	100%	-	-	100%
Constructor	-	-	-	-
Field	-	-	-	-
Method /Function	-	96%	100%	96%
Getter/Setter	-	-	-	-
Parameter	100%	-	-	100%
Variables	100%	76%	43%	75%
Overall	**100%**	**78%**	**84%**	**85%**

Table 8.8 Detected renamings and recall *Rc* of different entities.

	Wordpress		Drupal		phpBB		Overall	
Type	100%	(1/1)	-		-		100%	(1/1)
Constructor	-		-		-		-	
Field	-		-		-		-	
Method/Function	0%	(0/2)	100%	(23/23)	100%	(1/1)	92%	(24/26)
Getter/Setter	-		-		-		-	
Parameter	100%	(1/1)	-		0%	(0/1)	50%	(1/2)
Variable	67%	(8/12)	33%	(3/9)	93	(13/14)	68%	(24/35)
Overall	62%	(10/16)	81%	(26/32)	87%	(14/16)	78%	(50/64)

see that most of the renamings are classified as *simple* renamings. We did not find any term reordering instances and the term formatting instances are due to addition of underscore to the beginning of the identifiers.

For the classification of *semantic changes*, REPENT exhibits an accuracy 81% (see Table 8.10). It has low performance in classifying renamings that add or remove meanings. The miss classification is due to splitting of identifiers that are all in lower cases. The only case of abbreviation is in identifier `admin_ldap` $\rightarrow$ `acp_ldap` in phpBB and REPENT failed to classify it correctly. The term `acp` is the abbreviation of `administrator control panel` and only recognizable if one has the domain knowledge. REPENT failed to classify two identifier as whole-part in Wordpress. One case is due to use of abbreviation `src` instead of `source` in the following renaming: `new_content` $\rightarrow$ `new_src`. The other missed instance is `title` $\rightarrow$ `link_text` which is due to wrong term mappings. No instance of *broaden meaning* is found for PHP programs. The three instances of PHP renamings classified as *narrow mean-*

Table 8.9 Evaluation of classification for "Forms of renaming"- PHP programs.

	Wordpress		Drupal		phpBB		Overall	
Simple	100%	(9/9)	100%	(26/26)	100%	(11/11)	**100%**	(46/46)
Complex	100%	(1/1)	-		100%	(1/1)	**100%**	(2/2)
Formatting only	-		-		100%	(2/2)	**100%**	(2/2)
Term reordering	-		-		-		-	
Overall	100%	(10/10)	100%	(26/26)	100%	(14/14)	100%	(50/50)

ing are all specialization phrases. REPENT failed to classify three cases of specialization phrases in phpBB. The reason is wrong splitting of identifier with lower case terms and use of abbreviation in the following cases: `module_name` → `module_basename` , `column_type` → `orig_column_type` and `admin_ldap` → `acp_ldap`.

Table 8.11 shows the results of and the precision of classification of the grammar changes in the three programs.

Unlike for Java programs the accuracy of REPENT is 100% for PHP programs we analyzed. Its high performance is due to higher number of renamings classified as *None*.

8.5   Discussion

In this chapter we explain applicability of proposed approach for the detection and classification of renamings in PHP programs. We resolved file inclusion mechanism in PHP using a simplified symbolic execution and a fix point algorithm. For Wordpress, in average 91% of the include statement are resolved however in Drupla and phpBB in average 80% and 78% of include statements are resolved. This is due to use of variables and function calls in include statements of these two systems. Improvement requires propagation of variables as proposed by Gauthier *et al.* (Gauthier et Merlo, 2012).

We used the same thresholds calibrated for the detection and classification of Java programs. To re-calibrate the thresholds for PHP programs we need more data. We did not consider dynamic features such as variable variables, function pointers and references passed to function/method while performing data flow analysis.

One interesting observation in the three PHP programs is that all the renamed variables (except for two constants) are local variables. PHP has three scope levels: the (global) file-level, the class-level, and the method/function level. Variables at file-level are visible in other files trough include mechanism, and in other functions through use of global modifier. Renaming of a variable at file-level requires more efforts than renaming of a variable in a local scope specially if renaming is performed manually. Rename refactoring plugins embedded in IDEs such as Eclipse and NetBean provide a fully automatic renaming for Java programs. *Netbeans for PHP*[13] and *phpStorm*[14] provide a semi automatic renamings in PHP by identifying all occurrences of a name in the program. However, it is up to the developer to manually inspect and evaluate if renmaing is safe and would not alter the desired behavior. This could be one factor contributes to less renamings in PHP programs. Some dynamic features of PHP such as variable variables, and function variables make it difficult to provide a fully automatic re-

---

13. https://netbeans.org/features/php/
14. https://www.jetbrains.com/phpstorm/

Table 8.10 Evaluation of classification for "Semantic changes" - PHP programs.

		Wordpress		Drupal		phpBB		Overall			
Preserve	Synonym	-		-		-		-			
	Synonym phrase	-		-		-		-			
	Spelling error	-		-		100%	(3/3)	100%	(3/3)		
	Expansion	-		-		-		-			
	Abbreviation	-		-		-		-			
	Overall	-		-		100%	(3/3)	100%	(3/3)		
Change	Opposite	-		-		-		-			
	Opposite phrase	-		-		-		-			
	Whole-part	-		-		-		-			
	Whole-part phrase	-		-		-		-		-	-
	Unrelated	60%	(3/5)	57%	(4/7)	-		58%	(7/12)		
	Overall	60%	(3/5)	44%	(4/9)	-		40%	(7/14)		
Narrow	Specialization	-		-		-		-			
	Specialization phrase	100%	(3/3)	-		-		100%	(3/3)		
	Overall	100%	(3/3)	-		-		100%	(3/3)		
Broaden	Generalization	-		-		-		-			
	Generalization phrase	0%	(0/1)	-		-		0%			
	Overall	0%	(0/1)	-		-		0%			
Add		100%	(1/1)	50%	(1/2)	-		67%	(2/3)		
Remove		0%	(0/1)	100%	(1/1)	-		50%	(1/2)		
None		-		100%	(2/2)	100%	(26/26)	100%	(28/28)		
Overall		64%	(7/11)	65%	(11/17)	100%	(26/26)	81%	(44/54)		

Table 8.11 Evaluation of classification for "Grammar changes"- PHP programs.

			Wordpress		Drupal		phpBB		Overall		
Grammar change	Part of speech change	Singular/Plural	100%	(1/1)	-		-		100%	(1/1)	100%
		Verb conjugation	-		-		-		-		
		Other	100%	(1/1)	-		100%	(2/2)	100%	(3/3)	
	None		100%	(8/8)	100%	(26/26)	100%	(12/12)	100%	(46/46)	
Overall			100%	(10/10)	100%	(26/26)	100%	(14/14)	100%	(50/50)	

name refracting because aside from the occurrence of name of a function or variable, strings containing such name need to be evaluated. The following code snippet shows example of the above mentioned dynamic features. Renaming variable $*bar* and function *toto* requires renaming of both strings value "*bar*" and '*toto*' respectively.

```php
// Variable variable
$foo = "bar";
$$foo=35;
echo $bar;

// Function variable
function toto($arg = ''){
    echo "In toto(); argument was '$arg'.<br />\n";
}
$func = 'toto';
$func();
```

Another challenge to provide a fully automatic renamings is due to use of *hook* in plugin or module based frameworks. For example, Wordpress and Drupal provide a callback mechanism which both community refer to it as *hook*. The following code snippet show an example of '*helloworld*" *hook* in Drupal [15]. Call to function *module_invoke_all* initiates the call to function *module_helloworld*, and thus renaming the function name requires renaming of the argument passed to *module_invoke_all*.

```php
define('DRUPAL_ROOT', getcwd());
require_once DRUPAL_ROOT . '/includes/bootstrap.inc';
drupal_bootstrap(DRUPAL_BOOTSTRAP_FULL);
module_invoke_all('helloworld');
echo '<p>Done</p>';

#File: sites/all/modules/moduleb/moduleb.module
<?php
function module_helloworld()
{
    echo "<p>Our friends at ".__FUNCTION__." want to say Hello World</p>";
}
```

---

15. http://alanstorm.com/drupal_module_hooks

It is a preliminary study and to make any conclusion we need to study more file commits. However, we see that REPENT has an overall good accuracy in terms of precision and recall for both detection and classification of identifiers in PHP programs.

CHAPTER 9   CONCLUSION

In this chapter we summarize the findings and extract lessons learned from the classified renamings. Identifier renaming is considered by developers as an important and non-trivial task in the context of software evolution. We conducted a survey with 71 industrial and open-source developers and show that, as part of the evolution of software programs, developers rename identifiers to improve the quality of the source code lexicon and its consistency with the program functionality. Despite the importance of renaming and the associated with it cost (92% of the surveyed developers do not consider renaming as straightforward), renamings are hardly ever documented (1% of the renamings in the five programs that we studied) hence the need for an automatic documentation of renamings that 52% of the surveyed developers consider useful.

We propose REPENT (REnaming Program ENTities)—an approach to automatically document, *i.e.*, detect and classify, identifier renamings between different versions of a Java program. For detecting renamings, REPENT first reduces the search space by identifying changes from mapped source code lines between versions of a program resulting in candidate renamings, after which it performs data flow analysis on the entities involved in the candidate renamings to filter out false positives. We analyzed renamings detected by RE-PENT in the evolution history of five open-source programs (ArgoUML, dnsjava, Eclipse-JDT, JBoss, and Tomcat), and reported a precision of 88% and a recall—with respect to the documented renamings—of 92%. By combining an ontological database (WordNet) with a natural language parser (Stanford NLP) REPENT classifies renamings according to the different dimensions of our taxonomy, and specifically (i) the kind of entity being renamed, (ii) the form of the renaming, (iii) semantic change, and (iv) grammatical change. By relying on REPENT, we conducted an exploratory study—on the five Java programs used to evaluate the performances—aimed at determining how developers rename identifiers according to the proposed taxonomy. We assess feasibility of extending REPENT for detection and classification of renamings in PHP programs. We analyzed three open-source PHP programs (Wordpress, Drupal, and phpBB) and the results are motivating.

We conjecture that renamings documented and classified by REPENT can be used as a base towards building a renaming recommender system. This also reflect results of our survey, where 68% of the surveyed developers indicated the usefulness of automatic recommendations for renaming, provided that such recommendations are non-intrusive and offer reliable suggestions.

9.1   Lessons Learned

These lessons, being them common to all studied programs, or specific to only some of them, have the purpose to make developers aware about situations in which renamings occur where some improve the quality of the identifier and therefore should be promoted, while other decrease the quality and thus should be avoided. Lessons are extracted form detected and classified renamings but they are neither based on how often such situations occur nor on how important/crucial they are for the analyzed programs.

In **RQ1** we concluded that developers tend to rename methods and parameters more than other entities. In OO programming, methods express objects behavior and thus changing constantly due to changing requirements. Renaming methods and parameters is a must when one is concerned about consistency, as they carry a summary of the functionality and are used to communicate with others. We thus derive the following lesson:

> **L$_1$**. *Methods and parameters renamings are unavoidable due to evolution, i.e., constant changes in requirements.*

In **RQ1** we also observed, based on results from ArgoUML, that changes in used API can have a noticeable impact on the client lexicon, thus allowing us to conclude the next lesson:

> **L$_2$**. *Using APIs without planning for change can cause ripple effect on the client lexicon.*

While answering **RQ2**, we observed that when formatting renamings occur, the majority of the formatting changes tend to be renamings away from Hungarian notation. This allows us to derive the following lesson:

> **L$_3$**. *It is important to choose the naming conventions for each specific project in an early stage of the development process and following it consistently.*

We also observed a small fraction of renamings being about term reordering, which occurred on entities with different scope/visibility, and thus with different impact of the renaming. We can derive the next lesson:

> **L$_4$**. *It is worth taking the effort to identify the right order of terms constituting an identifier to clarify its meaning and avoid possible misunderstandings.*

While answering **RQ3**, we identified a small set of renamings (1%) related to spelling errors correction but also to spelling error introduction. Some identifiers were also involved into a sequence of renamings where each renaming is only partially correcting the spelling error. We deduce the next lesson, as follows:

> **L$_5$**. *To avoid the need for a sequence of renamings towards spelling error correction, it is worth taking the time to spellcheck the identifier name when creating*

or modifying an entity.

When answering **RQ3** we also noted that about 1% of the renamings were performed towards expansion or abbreviation and the distribution is relatively equal. We thus conclude the following lesson:

$\mathbf{L}_6$. *It is worth investigating which one of the two, an abbreviation or its English alternative, is more common and thus should be used.*

The qualitative analysis of the classified renamings while answering **RQ3** resulted in identifying a set of renamings where the old identifier name contains a negation, whereas the new identifier does not. Such analysis allows us to derive the next lesson:

$\mathbf{L}_7$. *Identifiers that contain negation tend to be renamed towards positive names.*

In the discussion of **RQ3**, we also bring evidence that, with respect to semantic changes, the majority of classified renamings are **not** about preserving the semantic meaning of identifiers, but rather changing, narrowing, broadening, adding, or removing a meaning. We thus derive the following lesson:

$\mathbf{L}_8$. *The majority of semantic changes during renamings change, narrow, broaden, add, or remove a meaning to the identifier, as part of the evolution process and thus cannot be avoided.*

In the discussion of **RQ3** and **RQ4**, we encountered renamings that were performed to increase the consistency between an entity and its functionality, between the name of an entity and names of other entities, and between an entity name and its type. From all those examples we derive the last lesson as follows:

$\mathbf{L}_9$. *It is worth the effort to assure consistency between, on the one hand, the name of an entity, and, on the other hand, its functionality, type, or other entities.*

Table 9.1 Actionable knowledge.

AK$_1$	Plan ahead to avoid an avalanche of renamings: design for change when using APIs, choose naming conventions at an early stage for the programming language you are using.	L$_2$, L$_3$
AK$_2$	Provide a small amount of effort regularly to avoid unnecessary renamings: check identifiers for spelling errors at the time of creation/renaming.	L$_5$
AK$_3$	Name identifiers for others: identify the right order of terms, choose between abbreviation or full English word, avoid using negation.	L$_4$, L$_6$, L$_7$
AK$_4$	Strive for consistency: when changing behavior/type, change the name accordingly, choose identifiers to be consistent with the rest of the programming lexicon.	L$_9$
AK$_5$	Do not try to prevent all renamings – renaming is part of evolution: program behavior evolves, method names and parameters change accordingly.	L$_1$, L$_8$

Based on the above lessons, we further derive the so called *actionable knowledge* reported in Table 9.1. Each *actionable knowledge* derives from one or more lessons and intend to provide practical advice that the authors believe worth highlighting.

The study allowed us to distill nine lessons and a set of five guidelines, *i.e.*, pieces of *actionable knowledge*, summarized in Table 9.1. Such knowledge can be used by for developers, to drive then towards using appropriate naming conventions—*e.g.*, using singular/plurals properly, making a wise usage of abbreviations, avoiding negated forms and Hungarian notation, and also keeping in mind that API renaming can have quite an impact on the rest of the source code if one wants to keep the lexicon consistent.

9.2 Future Work

Based on the results of our studies, the following extensions are possible future work. We present them from short-term and long-term perspectives.

9.2.1 Short-term

As pointed out above, REPENT relies on external tools such as the Unix diff, WordNet, and Stanford NLP, which turned out to suffice for our needs. However, future development activities could easily extend REPENT by replacing such tools with alternative and possibly more accurate ones.

9.2.2 Long-term

Base on the literature, our studies and observations, we argue that the following changes will happen to frameworks and their APIs.

More specifically, REPENT can be the core component of renaming recommenders that, by learning from past renamings, could automatically point out inconsistencies to developers— *e.g.*, linguistic antipatterns Arnaoudova *et al.* (2013)—to proactively suggest how to rename identifiers to improve source code comprehensibility as well as pointing to past renamings conflicting with the ongoing renaming activity. Such foreseeable recommenders would for example be useful in the following situation. In revision 67429 of JBoss, method `deploy` was renamed to `internalDeploy` in five different classes. In the same revision, in three of those classes method `unDeploy` was renamed to `internalUnDepoly`. The same set of renamings occurred at a later stage (revision 79147) for a different class. Hence, documenting the renamings in revision 67429 and learning from them would have facilitated the work of developers in later revisions, when creating an entity or renaming it, by pointing to names

used in a similar context. If such a recommendation is not accepted by the developer, it will still be beneficial as it will be clear that such contrast is deliberate and it is performed with the developer's full awareness, thus the rationale behind it must be explicitly documented for future evolution.

Besides building a renaming recommender, work-in-progress also aims at supporting programming languages other than Java, as for example scripting languages like PHP. Indeed our preliminary results on detection and classification of identifiers in PHP programs show the applicability of proposed methodology for this language as well. Although the results are motivating we need to improve the detection component of REPENT by considering dynamic features of PHP such as variable variable, function pointers and reference passing.

REFERENCES

Russell J. Abbott (1983). Program design by informal english descriptions. *Commun. ACM*, *26*(11), 882–894.

Surafel Lemma Abebe and Sonia Haiduc and Andrian Marcus and Paolo Tonella and Giuliano Antoniol (2009). Analyzing the evolution of the source code vocabulary. *Proceedings of the European Conference on Software Maintenance and Reengineering.* 189–198.

Surafel Lemma Abebe and Paolo Tonella (2010). Natural language parsing of program element names for concept extraction. *Proceedings of the International Conference on Program Comprehension (ICPC).* 156–159.

Giuliano Antoniol and Gerardo Canfora and Gerardo Casazza and Andrea De Lucia and Ettore Merlo (2002). Recovering traceability links between code and documentation. *IEEE Transactions on Software Engineering, 28*(10), 970–983.

Giuliano Antoniol and Massimiliano Di Penta and Michele Zazzara (2004). Understanding web applications through dynamic analysis. *Proceedings of 12th IEEE International Workshop on Program Comprehension.* 120–129.

Venera Arnaoudova and Massimiliano Di Penta and Giuliano Antoniol and Yann-Gaël Guéhéneuc (2013). A new family of software anti-patterns: Linguistic anti-patterns. *Proceedings of the European Conference on Software Maintenance and Reengineering (CSMR).*

Bavota, Gabriele and De Lucia, Andrea and Oliveto, Rocco (2011). Identifying extract class refactoring opportunities using structural and semantic cohesion measures. *Journal of Systems and Software, 84*, 397–414.

David Binkley and Matthew Hearn and Dawn Lawrie (2011). Improving identifier informativeness using part of speech information. *Proceedings of the International Working Conference on Mining Software Repositories.*

Black, Thomas R. (1999). *Doing Quantitative Research in the Social Sciences: An Integrated Approach to Research Design, Measurement and Statistics.* Statistics Series. SAGE Publications.

Broder, A.Z. (1997). On the resemblance and containment of documents. *Proceedings of the Compression and Complexity of Sequences.* 21–29.

B. Bruegge and Dutoit, A. H. (2003). *Object-Oriented Software Engineering: Using UML, Patterns, and Java.* Prentice Hall.

G. Capobianco and De Lucia, A. and R. Oliveto and A. Panichella and S. Panichella (2013). Improving ir-based traceability recovery via noun-based indexing of software artifacts. *Journal of Software: Evolution and Process, 25*(7), 743–762.

Bruno Caprile and Paolo Tonella (2000). Restructuring program identifier names. *Proceedings of the International Conference on Software Maintenance.* 97–107.

Anna Corazza and Sergio Di Martino and Valerio Maggio (2012). Linsen: An approach to split identifiers and expand abbreviations with linear complexity. *Proceedings of the International Conference on Software Maintenance, (ICSM).*

James R. Cordy (2006). The TXL source transformation language. *Science of Computer Programming, 61*(3), 190–210.

Thomas H. Cormen and Charles E. Leiserson and Ronald L. Rivest (1990). *Introductions to Algorithms.* MIT Press.

Andrea De Lucia and Massimiliano Di Penta and Rocco Oliveto (2011). Improving source code lexicon via traceability and information retrieval. *IEEE Transactions on Software Engineering, 37*(2), 205–227.

Florian Deissenbock and Markus Pizka (2005). Concise and consistent naming. *Proceedings of the International Workshop on Program Comprehension.*

Serge Demeyer and Stéphane Ducasse and Oscar Nierstrasz (2000). Finding refactorings via change metrics. *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications.* 166–177.

Di Lucca, G.A. and Fasolino, A.R. and Pace, F. and Tramontana, P. and De Carlini, U. (2002). WARE: A tool for the reverse engineering of web applications. *Proceedings of the European Conference on Software Maintenance and Reengineering.* Budapest, Hungary, 241–250.

Di Lucca, G.A. and Fasolino, A.R. and Tramontana, P. and De Carlini, U. (2003a). Abstracting business level UML diagrams from web applications. Amsterdam, The Netherlands, 12–19.

Di Lucca, G.A. and Fasolino, A.R. and Tramontana, P. and De Carlini, U. (2003b). Recovering a business object model from web applications. *Proceedings of 27th International Computer Software and Applications Conference (COMPSAC).* 348–353.

Danny Dig and Can Comertoglu and Darko Marinov and Ralph E. Johnson (2006). Automated detection of refactorings in evolving components. *Proceedings of the European Conference on Object-Oriented Programming.* 404–428.

Eric Enslen and Emily Hill and Lori L. Pollock and K. Vijay-Shanker (2009). Mining source code to automatically split identifiers for software analysis. *Proceedings of the International Working Conference on Mining Software Repositories.* 71–80.

Laleh Mousavi Eshkevari and Venera Arnaoudova and Massimiliano Di Penta and Rocco Oliveto and Yann-Gaël Guéhéneuc and Giuliano Antoniol (2011). An exploratory study of identifier renamings. *Proceedings of the 8th International Working Conference on Mining Software Repositories (MSR 2011).* 33–42.

Fluri, Beat and Wuersch, Michael and PInzger, Martin and Gall, Harald (2007). Change distilling: Tree differencing for fine-grained source code change extraction. *IEEE Transactions on Software Engineering, 33*(11), 725–743.

François Gauthier and Ettore Merlo (2012). Alias-aware propagation of simple pattern-based properties in PHP applications. *Proceedings of 12th International Working Conference on Source Code Analysis and Manipulation, (SCAM).* 44–53.

François Gauthier and Ettore Merlo (2013). Semantic smells and errors in access control models: a case study in php. *Proceedings of 35th International Conference on Software Engineering (ICSE).* 1169–1172.

Barney G. Glaser (1992). *Basics of grounded theory analysis.* Sociology Press.

Robert M. Groves and Floyd J. Fowler Jr. and Mick P. Couper and James M. Lepkowski and Eleanor Singer and Roger Tourangeau (2009). *Survey Methodology, 2nd edition.* Wiley.

L. Guerrouj and M. Di Penta and G. Antoniol and Y. Gaël Guéhéneuc (2011). Tidier: An identifier splitting approach using speech recognition techniques. *Journal of Software Maintenance - Research and Practice, 25*(6), 575–599.

Gupta, Samir and Malik, Sana and Pollock, Lori and Vijay-Shanker, K. (2013). Part-of-speech tagging of program identifiers for improved text-based software engineering tool. *Proceedings of the International Conference on Program Comprehension (ICPC).* 3–12.

Haiduc, Sonia and Marcus, Andrian (2008). On the use of domain terms in source code. *Proceedings of the International Conference on Program Comprehension (ICPC).* 113–122.

Ahmed E. Hassan and Richard C. Holt (2002). Architecture recovery of web applications. *Proceedings of the 22rd International Conference on Software Engineering, (ICSE).* 349–359.

Abram Hindle and Neil A. Ernst and Michael W. Godfrey and John Mylopoulos (2011). Automated topic naming to support cross-project analysis of software maintenance activities. *Proceedings of the International Working Conference on Mining Software Repositories (MSR).* 163–172.

Howard, Matthew J. and Gupta, Samir and Pollock, Lori and Vijay-Shanker, K. (2013). Automatically mining software-based, semantically-similar words from comment-code mappings. *Proceedings of the Working Conference on Mining Software Repositories (MSR)*.

Toshihiro Kamiya and Shinji Kusumoto and Katsuro Inoue (2002). CCFinder: A multilinguistic token-based code clone detection system for large scale source code. *IEEE Transactions on Software Engineering*, *28*(7), 654–670.

Kawrykow, David and Robillard, Martin P. (2011). Non-essential changes in version histories. *Proceedings of the International Conference on Software Engineering.* 351–360.

Rainer Koschke and Gerardo Canfora and Jörg Czeranski (2006). Revisiting the delta ic approach to component recovery. *Science of Computer Programming*, *60*(2), 171–188.

Harold W. Kuhn (1955). The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*, 83–97.

Dawn Lawrie and David Binkley (2011). Expanding identifiers to normalize source code vocabulary. *Proceedings of the International Conference on Software Maintenance, (ICSM).* 113–122.

Dawn Lawrie and Henry Feild and David Binkley (2006a). Syntactic identifier conciseness and consistency. *Proceedings of the International Workshop on Source Code Analysis and Manipulation.* 139–148.

Dawn Lawrie and Christopher Morrell and Henry Feild and David Binkley (2006b). What's in a name? a study of identifiers. *Proceedings of the International Conference on Program Comprehension.* 3–12.

Dawn Lawrie and Christopher Morrell and Henry Feild and David Binkley (2007). Effective identifier names for comprehension and memory. *Innovations in Systems and Software Engineering*, *3*(4), 303–318.

Vladimir I. Levenshtein (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Cybernetics and Control Theory*, *10*(8), 707–710.

Nioosha Madani and Latifa Guerrouj and Massimiliano Di Penta and Yann-Gaël Guéhéneuc and Giuliano Antoniol (2010). Recognizing words from source code identifiers using speech recognition techniques. *Proceedings of the European Conference on Software Maintenance and Reengineering.* 68–77.

Jonathan I. Maletic and Giuliano Antoniol and Jane Cleland-Huang and Jane Huffman Hayes (2005). *International Workshop on Traceability in Emerging Forms of Software Engineering.* 462.

Malpohl, Guido and Hunt, James J. and Tichy, Walter F. (2000). Renaming detection. *Proceedings of the International Conference Automated Software Engineering (ASE)*. 73–80.

Marcus, Andrian and Poshyvanyk, Denys and Ferenc, Rudolf (2008). Using the conceptual cohesion of classes for fault prediction in object-oriented systems. *IEEE Transactions on Software Engineering*, *34*(2), 287–300.

Marcus, Mitchell P. and Marcinkiewicz, Mary Ann and Santorini, Beatrice (1993). Building a large annotated corpus of english: the penn treebank. *Journal of Computational Linguistics - Special issue on using large corpora*, *19*(2), 313–330.

George A. Miller (1995). Wordnet: A lexical database for english. *Communications of the ACM*, *38*(11), 39–41.

Neamtiu, Iulian and Foster, Jeffrey S. and Hicks, Michael (2005). Understanding source code evolution using abstract syntax tree matching. *Software Engineering Notes*, *30*(4), 1–5.

Hoan Anh Nguyen and Hung Viet Nguyen and Tung Thanh Nguyen and Tien N. Nguyen (2013a). Output-oriented refactoring in php-based dynamic web applications. *Proceedings of International Conference on Software Maintenance (ICSM)*. 150–159.

Hung Viet Nguyen and Hoan Anh Nguyen and Tung Thanh Nguyen and Tien N. Nguyen (2011). Auto-locating and fix-propagating for HTML validation errors to PHP server-side code. *Proceedings of 26th International Conference on Automated Software Engineering (ASE)*. 13–22.

Hung Viet Nguyen and Hoan Anh Nguyen and Tung Thanh Nguyen and Tien N. Nguyen (2013b). Drc: a detection tool for dangling references in php-based web applications. *Proceedings of the International Conference on Software Engineering, (ICSE)*. 1299–1302.

Linda Dailey Paulson (2007). Developers shift to dynamic programming languages. *IEEE Computer*, *40*(2), 12–15.

M. F. Porter (1980). An algorithm for suffix stripping. *Program*, *14*(3), 130–137.

Denys Poshyvanyk and Andrian Marcus (2006). The conceptual coupling metrics for object-oriented systems. *Proceedings of the International Conference on Software Maintenance*. 469–478.

Kyle Prete and Napol Rachatasumrit and Nikita Sudan and Miryung Kim (2010). Template-based reconstruction of complex refactorings. *Proceedings of the International Conference on Software Maintenance (ICSM)*. 1–10.

Ricca, F. and Tonella, P. (2001a). Analysis and testing of web applications. *Proceedings of the International Conference on Software Engineering, (ICSE)*. 25–34.

Ricca, F. and Tonella, P. (2001b). Understanding and restructuring web sites with ReWeb. *IEEE Multimedia*, *8*(2), 40–51.

George Santayana (1905). *The Life of Reason: Introduction and Reason in Common Sense*, vol. 1. Charles Scribner's Sons.

David J. Sheskin (2007). *Handbook of Parametric and Nonparametric Statistical Procedures (fourth edition)*. Chapman & All.

Standish, Thomas A. (1984). An essay on software reuse. *IEEE Transactions on Software Engineering (TSE)*, *10*(5), 494–497.

Anselm L. Strauss (1987). *Qualitative analysis for social scientists*. Cambridge Univsersity Press.

Armstrong Takang and Penny A. Grubb and Robert D. Macredie (1996). The effects of comments and identifier names on program comprehensibility: an experiential study. *Journal of Program Languages*, *4*(3), 143–167.

Suresh Thummalapenta and Luigi Cerulo and Lerina Aversano and Massimiliano Di Penta (2010). An empirical study on the maintenance of source code clones. *Empirical Software Engineering*, *15*(1), 1–34.

Rebecca Tiarks (2011). What maintenance programmers really do: An observational study. *Proceedings of the Workshop Software Reengineering (WSR)*. 36–37.

Tonella, P. and Ricca, F. (2002). Dynamic model extraction and statistical analysis of web applications. 43–52.

Toutanova, Kristina and Manning, Christopher D. (2000). Enriching the knowledge sources used in a maximum entropy part-of-speech tagger. *Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*. 63–70.

Anneliese von Mayrhauser and Marie A. Vans and Adele E. Howe (1997). Program understanding behaviour during enhancement of large-scale software. *Journal of Software Maintenance: Research and Practice*, *9*(5), 299–327.

Weissgerber, Peter and Diehl, Stephan (2006). Identifying refactorings from source-code changes. *Proceedings of the International Conference on Automated Software Engineering (ASE)*. 231–240.

Xing, Zhenchang and Stroulia, Eleni (2006). Refactoring detection based on UMLDiff change-facts queries. *Proceedings of Working Conference on Reverse Engineering*. 263–274.

Jinqiu Yang and Lin Tan (2013). SWordNet: Inferring semantically related words from software context. *Empirical Software Engineering*.

Thomas Zimmermann and Peter Weisgerber (2004). Preprocessing CVS data for fine-grained analysis. *Proceedings of the International Workshop on Mining Software Repositories.* 2–6.

Thomas Zimmermann and Peter Weissgerber and Stephan Diehl and Andreas Zeller (2004). Mining version histories to guide software changes. *Proceedings of the International Conference on Software Engineering.* 563–572.

Survey details

This appendix reports detailed results of the survey

First, we report information about participants' background. In particular, Fig. .1 shows statistics regarding the native language of the participants, whereas Fig. .2 reports their years of experience in industrial and open-source software development.

Fig. .3 reports how often developers rename. Only 14% of participants rename rarely (up to once per month) : 46% rename occasionally (a few times per month) while 18% rename frequently (a few times per week) and 21% rename very frequently (almost every day).

Fig. .4 indicates activities during which developers rename. Note that a participant may select more than one activity, thus the sum of the percentages is above 100%. Participants rarely perform renaming as a standalone activity (17%). Often, they rename when performing other refactorings (90%), changing the functionality (89%), adding new functionality (65%), understanding code (51%), or fixing a bug (42%).

Fig. .5 provides insights about the opinion of participants about the cost of renaming. 35% of participants consider that renaming requires time and effort (at least in most cases) ; 32% consider that the cost of renaming depends on the particular case ; 32% consider renaming to be straightforward (at least in most cases). Note that the sum of the above is 99% due to rounding errors.

Fig. .6 reports results on the use of tool support for renaming. The majority of the participants (72%) use automatic tool support to perform renaming. There are however participants that rename manually (20%) and participants that perform both, manual and automatic renaming (8%).

We asked participants to share the reasons for which they recall having decided not to rename an entity ; results are shown in Fig. .7. 52% of the participants recall the reason to be the potential impact on other systems. 35% recall that the renaming was too risky, *i.e.*, it might have introduced a bug. 25% of the participants answered that the high impact of the renaming


Figure .1 Native language of the participants.

Figure .2 Experience of the participants in software development.


Figure .3 How often do developers rename?


Figure .4 Activities accompanying renaming.

Figure .5 Developers' opinion on cost of renaming.



Figure .6 How do developers rename?

on the system was the show-stopper and finally, 25% recall deciding not to rename because of the high effort required.

We also asked participants whether a set of predefined factors would impact the decision to undertake a renaming. Results are not surprising (Fig. .8). The majority of participants consider important all those factors. However, the factor that is worth highlighting here is the impact on other projects—69% of participants say that this would definitely impact their decision.

Fig. .9 shows when developers feel the need to rename. As expected, the majority (66%) of participants clearly state that they will definitely rename an entity if its name is not consistent with its functionality. They made less strong statements about naming conventions, spelling errors, and hard to understand words, but still the majority of participants report that they will probably rename in such cases. Surprisingly, only 13% of participants will probably rename if an entity contains an abbreviation—the majority of participants (56%) will not rename. Finally, when the name of an entity contains a negation, *e.g.*, `notOpen`, 30% of the participants will rename, while 46% will not.

We asked participants whether they consider useful automatically documenting renaming; results are shown in Fig. .10. The majority (52%) are positive. 33% of participants are negative about automatic documentation. The remaining participants did not provide their opinion.

Fig. .11 reports participants' opinion on renamings that are useful to document. More specifically, developers see the usefulness of automatically documenting renamings of public APIs, *i.e.*, classes and methods, and renamings concerning the meaning of the name—renaming towards opposite meaning, towards unrelated words, towards more general/specific names, adding/removing meaning. Surprisingly, the percentage of participants that see a benefit from automatic renaming towards synonyms is lower—36%. Similarly, but not surprisingly, a small number of participants see a benefit from automatically documenting renaming of entities with local scope, renamings towards abbreviations/expansions, and spelling errors.

Figure .7 Reasons for which developers already postponed or canceled a renaming.

Legend: ■ Definitely YES ■ Probably YES ▦ Undecided ■ Probably NO ▢ Definitely NO

	Definitely YES	Probably YES	Undecided	Probably NO	Definitely NO
Code ownership	38%	41%	10%	6%	3%
Many uses	31%	32%	4%	24%	7%
Uses in other projects	69%	27%			1%
Close to deadline	44%	32%	10%	8%	3%
Insufficient (lack of) domain knowledge	58%	27%	7%	6%	1%

Figure .8 Factors impacting developers decision to undertake a renaming.

Legend: ■ Definitely YES ■ Probably YES ▦ Undecided ■ Probably NO ▢ Definitely NO

	Definitely YES	Probably YES	Undecided	Probably NO	Definitely NO
The name and functionality are not consistent	66%	32%			1%
The name does not follow the language naming conventions	34%	46%	10%	7%	3%
The name does not follow the team naming conventions	34%	38%	17%	8%	
The name contains an abbreviation/acronym		13%	30%	49%	7%
The name contains a spelling error	44%	41%	7%	8%	
The name contains misleading/hard to understand words	31%	48%	13%	7%	
The name contains a negation	7%	23%	21%	38%	8%

Figure .9 When will developers rename ?

Legend: ■ Definitely YES ■ Probably YES ■ Probably NO ▢ Definitely NO

	Definitely YES	Probably YES	Probably NO	Definitely NO
Is automatically documenting identifier renamings useful?	18%	34%	25%	8%

Figure .10 Developers' opinion on the usefulness of documenting renamings.
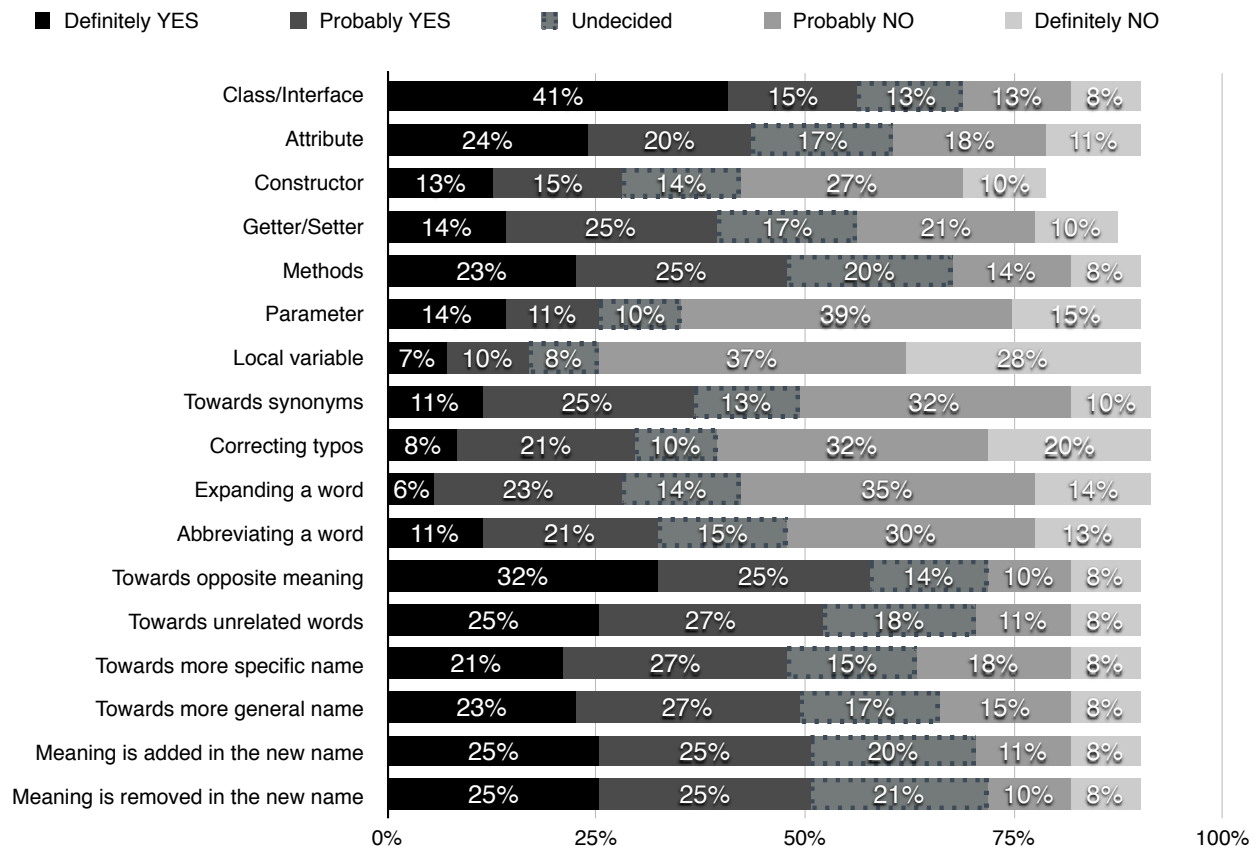
Figure .11 Developers' opinion on renamings that are useful to document.

The majority (68%) of participants see a benefit of automatic recommendations for renaming (Fig. .12) provided that such recommendations are non-intrusive and offer reliable suggestions. Fig. .12 reports developers' opinion on the usefulness of recommending renamings. Finally, participants see the benefit of recommendations regarding the majority of renamings (Fig. .13).

Figure .12 Developers' opinion on the usefulness of recommending renamings.

Figure .13 Developers' opinion on renamings that are useful to recommend.

Thresholds for Detection and Classification

**Thresholds for Detecting Renamings**   REPENT uses three sets of thresholds, where each threshold varies in the range $[0, 1]$. Specifically, the REPENT thresholds are described as follows :

– **Statement Similarity Threshold (SST) :** used to match def-uses of the mapped entities.
– **Number of Matched Statements Threshold (NST) :** used to decide if the mapped entities have a sufficiently high number of matched def-uses statements, to support the evidence of a real renaming.
– **Declaration Similarity Threshold (DST) :** used to match the declaration of two mapped entities, whenever one or both entities do not have uses.

While DST has been set to a constant value not depending on the kind of entity, *SST* and *NST* require different calibrations to be able to work effectively on different kinds of entities, in order to balance false positives as well as false negatives. The rationale is that for different entities (class names, method names, etc.) the syntax and frequency of statements where def-uses occur may be quite different.

We calibrate the thresholds by varying NST between zero and one, with step 0.1 ; for each fixed value of *NST* (*e.g.*, 0.4), we made *SST* vary (between zero and one with step 0.1).

We used Tomcat as a calibration data set, *i.e.*, we used the oracle of 724 renamings detected in our previous study Eshkevari *et al.* (2011) and manually classified. In addition, to better distinguish how REPENT performs with different thresholds, we computed the set of renamings that change when thresholds vary—we computed the union and intersection of the detected renamings for all combinations of the different values of *NST* and *SST*. We then randomly sampled renamings from the complement of the intersection until reaching an oracle that sufficiently discriminates the different thresholds—we stopped when the oracle reached 2,265 renamings. Next, we computed the True Positive Rate (TPR) and the False Positive Rate (FPR) for each entity kind on the calibration set. TPR and FPR were used to generate a family of Receiver Operating Characteristic (ROC) curves. ROC curves plot the TPR over the FPR at various threshold settings. Notice that TPR is equivalent to recall. The relation between FPR (TPR) and precision is more complex as precision is the ratio of true positives over the sum of true positives and false positives ; however, reducing FPR increases the number of correctly classified items, and thus it improves precision.

Figure .1 REPENT class FielD (FD) ROC curves as functions of SST and NST.

Table .1 Thresholds chosen for the study as well as corresponding TPR and FPR on the calibration set.

Kind of entity	SST	NST	DST	TPR		FPR		Precision
Type	0.8	0.3	0.7	89%	(17/19)	0%	(0/1)	100%
Constructor	0.6	0.8	0.7	100%	(7/7)	0%	(0/4)	100%
Field	0.8	0.4	0.7	86%	(262/303)	12%	(20/163)	93%
Method	0.8	0.3	0.7	95%	(349/368)	18%	(11/61)	97%
Getter/Setter	0.8	0.2	0.7	93%	(250/270)	18%	(22/122)	92%
Parameter	0.7	0.5	0.7	64%	(214/334)	36%	(48/132)	82%
Local variable	0.8	0.3	0.7	77%	(172/224)	4%	(11/257)	94%
Overall				83%	(1271/1525)	15%	(112/740)	92%

Fig. .1 shows a subset of ROC family computed for class FielD (FD); top right curve (symbol +) was obtained setting *NST*= 0.1, while the bottom left curve (symbol ▲) corresponds to NST=0.9. REPENT works on the ROC curve with *NST*= 0.4 (symbol ⊡). The *SST* value was fixed at 0.8 which gives (on the calibration set) TPR of 86% and FPR of 12%. As shown in Fig. .1, these two values are the threshold values giving (on our FD calibration set) the highest TPR with a reasonably low FPR for FD; no other ROC curve has a lower FPR with a higher (or equal) TPR. For example, the topmost-right ROC curve (denoted with +) indicates a higher TPR at a price of a higher FPR. In essence, depending on the goal to be achieved, REPENT can be calibrated to favor low FPR, high TPR or a compromise between the two. This latter choice was used to select, via the same analysis process, for each entity kind *SST* and *NST* values reported in Table .1.

Thresholds for Classifying Renamings To classify renamings as spelling errors RE-PENT uses a threshold for the Levenshtein distance between the old and new name of an entity. Lower values for the Levenshtein distance threshold ensure a high precision in the classification, but also a higher number of false negatives. Table .2 shows the accuracy of the classification of renamings as spelling errors on Tomcat when the Levenshtein distance threshold varies from 1 to 5. A threshold of 1 ensures 0% FPR while failing to classify as spelling errors renamings such as `refeelReadBuffer` → `refillReadBuffer`. Increasing the threshold to 2 solves this issue while keeping a low FPR (4% on Tomcat). An example of misclassified renaming when the Levenshtein distance threshold is set to 2 is `class` → `clazz`.

Table .2 Accuracy of the classification of renamings as spelling errors using different Levenshtein distance thresholds on Tomcat.

Threshold value	TPR	FPR
1	86%	0%
2	100%	4%
3	100%	25%
4	100%	50%
5	100%	100%

Sampling for Evaluating Detection and Classification

Sampling for Evaluating the Detection   Table .1 reports the sample size and the number of detected renamings, overall and for each kind of entity (we highlight in bold face the size of the significant sample). For example, for ArgoUML (first system in Table .1) we sampled 352 renamings for validation out of the 3,994 detected renamings for that program.

Sampling for Evaluating the Classification   Tables .2 to .4 report the sample size and the number of correctly classified renamings for each dimension of taxonomy. We highlight in bold face the significant samples and the corresponding precision. For example, regarding *forms of renaming* (Table .2) we sampled 96 *simple* renamings and 93 of them are correctly classified thus resulting in a precision of 97%.

We do not evaluate the accuracy of the classification for the entity kind dimension as it is correct by construction, *i.e.*, it is extracted when parsing the source code, and its only imprecision could also be due to mistakes in the Eclipse JDT parser we used.

Table .1 Sample size to estimate the precision of REPENT.

	ArgoUML		dnsjava		Eclipse-JDT		JBoss		Tomcat		Overall	
Package	1 /	1	0 /	0	1 /	4	1 /	7	0 /	0	3 /	12
Type	2 /	18	17 /	67	5 /	180	18 /	656	9 /	70	51 /	991
Constructor	1 /	16	40 /	159	4 /	139	13 /	475	7 /	49	65 /	838
Field	190 /	2,156	15 /	59	58 /	1,945	49 /	1,808	67 /	498	379 /	6,466
other-MD	17 /	200	44 /	177	58 /	1,966	66 /	2,397	59 /	441	244 /	5,181
Getter/Setter	17 /	192	11 /	45	37 /	1,244	44 /	1,595	57 /	423	166 /	3,499
Parameter	61 /	696	126 /	506	96 /	3,226	94 /	3,419	75 /	561	452 /	8,408
Local variable	63 /	715	37 /	149	114 /	3,853	90 /	3,261	59 /	439	363 /	8,417
OVERALL	**352** /	3,994	**290** /	1162	**373** /	12,557	**375** /	13,618	**333** /	2,481	1,723 /	33,812

Table .2 Evaluation of classification for "Forms of renaming" - Java programs.

	ArgoUML		dnsjava		Eclipse-JDT		JBoss		Tomcat		Overall	
Simple	100%	(4/4)	100%	(1/1)	94%	(34/36)	100%	(43/43)	92%	(11/12)	**97%**	(93/**96**)
Complex	100%	(2/2)	100%	(2/2)	96%	(47/49)	94%	(32/34)	100%	(8/8)	**96%**	(91/**95**)
Formatting only	100%	(58/58)	100%	(6/6)	100%	(3/3)	100%	(24/24)	100%	(2/2)	**100%**	(93/**93**)
Term reordering	-		-		100%	(23/23)	100%	(18/18)	100%	(5/5)	**100%**	(46/**46**)
Overall	100%	(64/64)	100%	(9/9)	96%	(107/111)	98%	(117/119)	96%	(26/27)	98%	(323/330)

Table .3 Evaluation of classification for "Semantic changes"- Java programs.

		ArgoUML		dnsjava		Eclipse-JDT		JBoss		Tomcat		Overall	
Preserve	Synonym	50%	(1/2)	-		94%	(33/35)	97%	(36/37)	88%	(7/8)	**94%**	(77/**82**)
	Synonym phrase	-		-		0%	(0/1)	100%	(2/2)	-		**67%**	(2/**3**)
	Spelling error	-		100%	(1/1)	100%	(32/32)	94%	(32/34)	100%	(9/9)	**97%**	(74/**76**)
	Expansion	100%	(17/17)	100%	(1/1)	81%	(17/21)	90%	(19/21)	100%	(6/6)	**91%**	(60/**66**)
	Abbreviation	67%	(2/3)	-		93%	(13/14)	90%	(36/40)	100%	(9/9)	**91%**	(60/**66**)
	Overall	91%	(20/22)	100%	(2/2)	92%	(95/103)	93%	(125/134)	97%	(31/32)	93%	(273/293)
Change	Opposite	-		-		100%	(25/25)	88%	(14/16)	100%	(5/5)	**96%**	(44/**46**)
	Opposite phrase	0%	(0/1)	-		29%	(6/21)	44%	(8/18)	0%	(0/3)	**33%**	(14/**43**)
	Whole-part	-		-		-		-		100%	(2/2)	**100%**	(2/**2**)
	Whole-part phrase	-		-		-		-		-		-	
	Unrelated	67%	(4/6)	-		78%	(31/40)	79%	(33/42)	71%	(5/7)	**77%**	(73/**95**)
	Overall	57%	(4/7)	-		72%	(62/86)	72%	(55/76)	71%	(12/17)	72%	(133/186)
Narrow	Specialization	0%	(0/3)	100%	(1/1)	78%	(31/40)	35%	(11/31)	50%	(2/4)	**57%**	(45/**79**)
	Specialization phrase	100%	(3/3)	50%	(1/2)	70%	(28/40)	63%	(24/38)	56%	(5/9)	**66%**	(61/**92**)
	Overall	50%	(3/6)	67%	(2/3)	74%	(59/80)	51%	(35/69)	54%	(7/13)	62%	(106/171)
Broaden	Generalization	67%	(2/3)	-		85%	(40/47)	75%	(15/20)	63%	(5/8)	**79%**	(62/**78**)
	Generalization phrase	67%	(2/3)	-		65%	(33/51)	58%	(15/26)	36%	(4/11)	**59%**	(54/**91**)
	Overall	67%	(4/6)	-		74%	(73/98)	65%	(30/46)	47%	(9/19)	69%	(116/169)
Add		100%	(1/1)	-		84%	(47/56)	87%	(27/31)	43%	(3/7)	**82%**	(78/**95**)
Remove		100%	(3/3)	100%	(3/3)	94%	(46/49)	88%	(30/34)	67%	(4/6)	**91%**	(86/**95**)
None		100%	(40/40)	100%	(5/5)	100%	(16/16)	100%	(27/27)	100%	(5/5)	**100%**	(93/**93**)
Overall		88%	(75/85)	92%	(12/13)	82%	(398/488)	79%	(329/417)	72%	(71/99)	80%	(885/1102)

Table .4 Evaluation of classification for "Grammar changes"- Java programs.

			ArgoUML		dnsjava		Eclipse-JDT		JBoss		Tomcat		Overall	
Grammar change	Part of speech change	Singular/Plural	100%	(1/1)	-		100%	(57/57)	100%	(28/28)	100%	(2/2)	**100%**	(88/**88**)
		Verb conjugation	50%	(1/2)	-		79%	(23/29)	83%	(29/35)	70%	(7/10)	**79%**	(60/**76**)
		Other	40%	(2/5)	-		25%	(11/44)	9%	(3/33)	23%	(3/13)	**20%**	(19/**95**)
	None		100%	(26/26)	100%	(4/4)	100%	(24/24)	100%	(34/34)	100%	(8/8)	**100%**	(96/**96**)
Overall			88%	(30/34)	100%	(4/4)	75%	(115/154)	72%	(94/130)	61%	(20/33)	74%	(263/355)