Towards Software Analytics: Modeling Maintenance Activities

STANISLAV LEVIN and AMIRAM YEHUDAI, Tel Aviv University, Israel

Lehman's Laws teach us that a software system will become progressively less satisfying to its users over time, unless it is continually adapted to meet new needs. Understanding software maintenance can potentially relieve many of the pains currently experienced by practitioners in the industry and assist in reducing uncertainty, improving cost-effectiveness, reliability and more. The research community classifies software maintenance into 3 main activities: *Corrective*: fault fixing; *Perfective*: system improvements; *Adaptive*: new feature introduction.

In this work we seek to model software maintenance activities and design a commit classification method capable of yielding a high quality classification model. We performed a comparative analysis of our method and existing techniques based on 11 popular open source projects from which we had manually classified 1151 commits, over 100 commits from each of the studied projects. The model we devised was able to achieve an accuracy of 76% and Kappa of 63% (considered "Good" in this context) for the test dataset, an improvement of over 20 percentage points, and a relative improvement of ~40% in the context of cross-project classification.

We then leverage our commit classification method to demonstrate two applications: (1) a tool aimed at providing an intuitive visualization of software maintenance activities over time, and (2) an in-depth analysis of the relationship between maintenance activities and unit tests.

CCS Concepts: • Software and its engineering \rightarrow Software evolution; Maintaining software;

Additional Key Words and Phrases: Software Maintenance, Mining Software Repositories, Predictive Models, Human Factors

Authors' address: Stanislav Levin, stas.levin@cs.tau.ac.il; Amiram Yehudai, amiramy@tau.ac.il, The Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel.

1 SOFTWARE EVOLUTION & MAINTENANCE

The software evolution phenomenon was first identified in the late 60's. The term software evolution however, was coined by Lehman only years later (Lehman, 1969, 1978; Lehman and Ramil, 2003). Initial studies in this area took place during the 70's and concentrated primarily on measuring and interpreting the growth of software systems and evolutionary trends (Belady and Lehman, 1971; Lehman and Ramil, 2003). Belady and Lehman (1976) recognized that the process of large-scale program development and maintenance appeared to be unpredictable, its costs were high and its output was a fragile product. They advocated that one should try to reach beyond understanding and attempt to change the process for the better. Lehman et al. (2000) classify the field of software evolution research into two groups, the first considers the term evolution as a *verb* while the second as a *noun*.

- THE VERBAL VIEW research is concerned with the question of "how", and focuses on means, processes, activities, languages, methods and tools required to effectively and reliably evolve a software system.
- THE NOUNAL VIEW research is concerned with the question of "what" and investigates the nature of software evolution, as a phenomenon, and focuses on the nature of evolution, its causes, properties, characteristics, consequences, impact, management and control (Lehman and Ramil, 2003; Lehman et al., 2000).

Lehman and Ramil (2003); Lehman et al. (2000) suggest that both views are mutually supportive. Moreover, they suggest that the verbal view research will benefit from progress made in studying the nounal view, and both are required if the community is to advance in mastering software evolution.

Software maintenance activities are a key aspect of software evolution and have been a subject of research in numerous works (Levin and Yehudai, 2016; Lientz et al., 1978; Meyers, 1988; Mockus and Votta, 2000; Schach et al., 2003; Swanson, 1976). As a step towards enhanced Software Analytics (Buse and Zimmermann, 2010; Menzies and Zimmermann, 2013), we believe that a better understanding of software maintenance activities could help practitioners reduce uncertainty and improve cost-effectiveness (Swanson, 1976) by planning ahead and pre-allocating resources towards source code maintenance. To determine maintenance activity profiles, one must first classify the activities (i.e., developer commits to the version control system), into one of the 3 maintenance activities kinds: *Corrective*: fault fixing; *Perfective*: system improvements; *Adaptive*: new feature introduction.

A widely practiced method for commit classification has been inspecting the commit message¹ (Amor et al., 2006; Fischer et al., 2003; Mockus and Votta, 2000; Śliwerski et al., 2005). Works employing commit message based classification reported the accuracy to average below 60% when used in the scope of a single project, and below 53% when used in the scope of multiple projects, i.e., when a single model was used to classify commits from multiple projects (Amor et al., 2006; Hindle et al., 2009). Arguably, low accuracy may be a significant barrier preventing these classification methods from being used in professional tools.

¹Also known as "commit comment".

It would therefore be beneficial to devise maintenance classification methods with higher accuracy (and overall classification quality). Our work is also motivated by the following observations:

- Cross project classification quality leaves much to be desired.
 Existing results rarely consider cross-project classification, which threatens external validity. Hindle et al. (2009) explored cross-project classification and reported the accuracy to be ~52%, which is considerably lower than the ~60% range reported by studies dealing with a single project.
- (2) Cohen's Kappa is vital to determine imbalanced classification quality, but it is rarely reported. Existing classification results rarely report Cohen's kappa (hence forth Kappa) metric (see also Section 3.1), which accounts for cases where classification labels (a.k.a classes) are unevenly distributed. Such cases make the accuracy metric somewhat misleading. For example, if the corrective class accounted for 98% of the commits in a given dataset, and each of the remaining classes accounted for 1% of the commits, then a simple classification model which always classified commits as corrective would have an impressive accuracy of 98%. Its Kappa on the other hand, would be 0, making this model much less appealing.
- (3) High quality maintenance activity classification may benefit both previous and future work. Our previous work (Levin and Yehudai, 2016) shows that source code change types as defined by Fluri and Gall (2006) are statistically significant in the context of maintenance activities defined by Mockus and Votta (2000). We believe that increasing the accuracy and Kappa characteristics of commit classification into maintenance activities could improve the quality and accuracy of individual developer maintenance profiles as well as the ability to build predictive models thereof.

In contrast to standard version control systems (VCS) and traditional diff tools which model code changes on the text level, in this work we wish to study changes in object oriented entities such as classes, methods, and fields throughout the life span of a software repository. To this end we use Fluri's taxonomy of source code changes (Fluri and Gall, 2006) for object-oriented programming languages (OOPLs), which consists of 48 (47 + an "unknown type") different change types, all of which are project agnostic and describe a meaningful action performed by a developer in a commit (e.g., *statement_delete, statement_insert, removed_class, additional_class* etc). Our work explores the following research questions:

- **RQ 1.** Can fine-grained source code changes be utilized to improve the quality of commit classification into maintenance activities?
- **RQ 2.** How does the quality of models which utilize fine-grained source code changes compare to that of traditional models which rely on word frequency analysis only?
- RQ 3. How can our findings be useful for practitioners and researchers?

This paper is an extension of our previous work (Levin and Yehudai, 2017b), where we first suggested utilizing fine-grained source code changes to classify commits into maintenance activities. In this extended paper, we provide a detailed discussion of our commit classification and repository harvesting methods, as well as new perspectives on applications for the discussed methods and techniques. To that end, Section 4

provides detailed information about the methods we used to effectively process Big Code, and Section 8 showcases additional applications which focuses on two particular directions: (1) Software Maintenance Activity Explorer, a tool aimed at providing an intuitive visualization of software maintenance activities over time, and (2) an in-depth analysis of the relationship between maintenance activities and unit tests in software projects.

2 RELATED WORK

The research community classifies software maintenance into 3 main activities: CORRECTIVE, PERFECTIVE and ADAPTIVE. The interpretation of these categories, and namely, the criteria to be used to determine which commits fall under what activity type is yet to reach a consensus. Swanson (1976) and Ghezzi et al. (2002) suggested the following definitions:

- Corrective: rectify the bugs observed while the system is in use.
- Perfective: support new features or enhance performance according to user demand.
- Adaptive: run on new platforms, new operating systems or interface with new hardware or software.

Mockus and Votta (2000) used different definitions for the perfective and adaptive activities:

- Perfective: code (re-)structuring to accommodate future changes.
- Adaptive: new feature introduction.

In this study we adopt the definitions put forth by Mockus and Votta (2000) and use these definitions to devise a commit classification method that improves existing results. Having spent almost a decade and a half professionally developing commercial software for both start-ups and enterprises, the authors feel that the definitions suggested by Mockus et al. almost two decades ago, have stood the test of time and remain relevant and applicable to how modern software evolves. For example, relatively new techniques such as refactoring are now common for improving the quality of code. Despite the fact refactoring became common only years after the definition by Mockus et al. had been suggested, refactoring fits perfectly under their definition for perfective maintenance. The alternative maintenance definitions on the other hand, seem to struggle with accommodating refactoring in a sensible manner. Moreover, we favour the interpretation by Mockus and Votta of the "adaptive" maintenance as adding new features (rather than accommodating new operating systems and hardware) since it intuitively covers one of the most basic activities carried out by developers - extending existing software with new features. The alternative definition of the "adaptive" maintenance activity speaks of adapting software to new platforms, operating systems and hardware. We believe that the latter has become significantly less frequent in (modern) software evolution. Even when considering the appearance of smart-phones and other gadgets which required the adaptation of software to new platforms and hardware, the endless stream of new features developers are required to implement in today's software seems like a much more dominant factor.

Mockus and Votta suggested the hypothesis that a textual description of the source code change (a commit to the VCS) is essential to understanding why that change was performed. To test this hypotheses, an automatic classification algorithm for maintenance activities was designed based on the textual description of changes. The automatic classification was then verified by surveying 8 developers. The survey results were in line with the automatic classification results, paving the road to text based commit classification approaches. The reported accuracy was ~61%. Fischer et al. (2003); Hindle et al. (2009); Levin and Yehudai (2016); Mockus and Votta (2000); Śliwerski et al. (2005) employed similar, keywords based, techniques for classifying commits into maintenance activities.

Recent work explored using additional information such as commits' author and module, to classify commits both within a single software project, and cross-projects (Hindle et al., 2009). Within a single project, the reported accuracy ranged from \sim 35% to 70% (accuracy fluctuated considerably depending on the project). In a cross-project scope, Hindle et al. (2009) reported the classification accuracy to be \sim 52%. A slightly different technique was used by Amor et al. (2006), who explored classifying maintenance activities in the FreeBSD project by applying a Naive Bayes classifier on commits' comments without an apparent use of keywords. In FreeBSD, the reported accuracy of classifying a random sample (whose size was not specified) was \sim 70% (within the scope of the FreeBSD project).

A summary of the existing results for commit classification into maintenance activities can be found in Table 1. In this work we were able to improve upon previous results and achieve an accuracy of 76% and Cohen's kappa of 63% in the context of cross-project commit classification, an improvement of over 20 percentage points and a relative improvement of \sim 40% in accuracy compared to previous results.

Study	Scope	Accuracy	F1 Score	Public Dataset
Hindle et al (2009)	Single Project	70%	0.69	N/A
11111ule et al. (2007)	Cross Project	52%	0.51	N/A
Amor et al. (2006)	Single Project	70%	N/A	N/A
Mockus and Votta (2000)	Single Project	61%	N/A	N/A

Table 1. Classifying commits into maintenance activities, existing results (Amor et al., 2006; Hindle et al., 2009; Mockus and Votta, 2000)

In contrast to prior studies which typically used the commit message to devise commit classification models, in this work we leverage fine grained source code changes in combination with the commit message to achieve superior model quality. In addition, we design and evaluate our models in a cross project scope (see also Table 4), rather than a single project scope. That is, after performing the per-project stratified sampling to obtain the ground truth dataset (see also Section 4), our subsequent model training and evaluation do not limit the commits to a single project, and are performed on heterogeneous commits (see also Table 4).

We also extend our previous work (Levin and Yehudai, 2017a) which studied the co-evolution of test maintenance and code maintenance and showed that maintenance activities can be successfully used to model the number of test methods and test classes in software projects. In particular, we provide statistical

evidence showing that software maintenance activities play an important role in modeling test (method and class) counts.

3 RESEARCH METHOD

Our research method consists of the following stages:

- (1) Select candidate software repositories and harvest their commit data such as commit message and source code changes performed in the commits (see Section 4).
- (2) Create a labeled dataset by sampling commits and manually labeling them. Each label is a maintenance activity, i.e. one of the following: corrective, perfective, or adaptive (see Section 5).
 - (a) Inspect the agreement level on the manually classified commits by having both authors independently classify a 10% sample of commits (see Section 5).
- (3) Devise predictive models that utilize source code changes for the task of commit classification into maintenance activities (see Section 6).
- (4) Evaluate the devised models using two mutually exclusive datasets obtained by splitting the labeled dataset into (1) a *training* dataset, consisting of 85% of the labeled dataset, and (2) a *test* dataset, consisting of the remaining 15% of the labeled dataset. The test dataset was never used as part of the training process (see Section 7).

3.1 Statistical Methods

Picking the optimal classifier for a real-world classification problem is hardly a simple task (Fernández-Delgado et al., 2014), however, Random Forest (RF) (Breiman, 2001; Ho, 1998) and Gradient Boosting Machine (GBM) (Caruana and Niculescu-Mizil, 2006; Caruana et al., 2008; Friedman, 2001) based classifiers are generally considered well performing (Caruana and Niculescu-Mizil, 2006; Fernández-Delgado et al., 2014). In addition, we also use J48, a variation of the C4.5 (Quinlan, 2014) algorithm. The RF implementation (Andy Liaw, 2015; Liaw and Wiener, 2002) and the GBM's one (Ridgeway, 2007; Ridgeway and Others, 2015) are most likely to outperform the simpler J48 (Frank et al., 2005; Hornik et al., 2009; Witten and Frank, 2005), but the latter, in contrast to the formers, is capable of providing a human readable representation of its decision tree. We find this ability valuable since inspecting the decision tree may reveal further insights. An example of a decision tree produced by the J48 classifier can be found in Figure 2, which depicts our keyword based commit classification model described in Section 6.

To evaluate the different commit classification models we employ common statistical measures for classification performance. For a given class $L \in \{adaptive, corrective, perfective\}, TP_L$ is the number of commits correctly classified as class L; FP_L is the number of commits incorrectly classified as class L; FN_L is the number of commits of class L that were incorrectly classified.

• $Precision_L = \frac{1P_L}{TP_L + FP_L}$, the number of commits correctly classified as class *L*, divided by the total number of commits classified as class *L*.

Towards Software Analytics: Modeling Maintenance Activities

- $Recall_L = \frac{TP_L}{TP_L + FN_L}$, the number of commits correctly classified as class *L*, divided by the actual number of *L* class commits in the dataset.
- Accuracy = $\frac{\sum_{L \in \{a,c,p\}} TP_L}{\sum_{L \in \{a,c,p\}} (TP_L + FP_L)}$, the proportion of correctly classified commits out of all classified commits.
- *No Information Rate (NIR)*, measures the accuracy of a trivial classifier which classifies all commits with using a single class, the one that is most frequent, in our case corrective.
- Kappa = Accuracy-ExpectedAccuracy/1-ExpectedAccuracy, Cohen's kappa, often considered helpful as a measure that can handle both multi-class and imbalanced class problems (see Section 1). Cohen's kappa measures the agreement between the predictions and the actual labels based on both the actual and predicted distributions.
- *P-Value [Accuracy > NIR]*, the *p*-value for the null hypothesis that the "Accuracy ≤ NIR" (i.e., the accuracy of a given predictive model). A low *p*-value allows one to reject the null hypothesis in favor of the alternative hypothesis that the "Accuracy > NIR".

4 DATA COLLECTION

We use GitHub (GitHub Inc., 2010) as the data source for this work due to its popularity (GitHub Inc., 2018) and rich query options (GitHub Inc., 2013, 2015). Candidate repositories were selected according to the following criteria, aimed to capture data-rich repositories that:

- (1) Used the Java programming language (our tools were Java oriented)
- (2) Had more than 100 stars (i.e. more than 100 users have "liked" these repositories)
- (3) Had more than 60 forks (i.e., more than 60 users have "cloned" these repositories to their private/organization accounts)
- (4) Had their code updated since 2016-01-01 (i.e., these repositories are active)
- (5) Were created before 2015-01-01 (i.e., these repositories have existed for several years)
- (6) Had size over 2,000 KB (i.e. these repositories are of considerable size)

The criteria aimed at capturing data abundant projects, i.e., projects with plenty of revisions that were still being actively developed. We found that while popularity related metrics such as stars and forks were a good start, after sampling some of the candidates we identified a number of projects that had little data (revisions) and were therefore not an ideal choice for our study. A closer examinations of these projects revealed that more than a few of them turned out to be visually pleasing Android User Interface (UI) controls which had gone viral. To mitigate this, we set a threshold on the repository size in an attempt to filter out small (yet widely popular) projects with little data to analyze.

In light of limited resources we reduced the final candidate set to 11 well known projects from the open source arena, representing various software domains such as IDEs, programming languages (that were implemented in Java), distributed database and storage platforms, and integration frameworks. Following is the list of projected studies in this work (see also Table 2):

- (1) **RxJava** a library for composing asynchronous and event-based programs for the Java VM.
- (2) Intellij Community Edition a popular IDE for the Java programming language.
- (3) HBase a distributed, scalable, big data store.
- (4) **Drools** a business rules management system solution.
- (5) Kotlin a statically typed programming language for the JVM, Android and the browser by JetBrains.
- (6) **Hadoop** a framework that allows for the distributed processing of large data sets across clusters of computers.
- (7) Elasticsearch a distributed search and analytics engine.
- (8) Restlet a RESTful web API framework for Java.
- (9) **OrientDB** a distributed graph database with the flexibility of documents in one product.
- (10) Camel an open source integration framework based on known enterprise integration patterns.
- (11) **Spring Framework** an application framework and inversion of control container for the Java platform.

Project	Total Commits	Total Contributors
RxJava	5,413	211
Restlet	8,840	39
Drools	11,713	137
HBase	15,561	189
Spring Framework	16,927	291
OrientDb	17,035	120
Hadoop	19,541	137
Camel	32,967	410
Elasticsearch	39,958	1103
Kotlin	47,386	239
Intellij Community Edition	232,607	356

Table 2. Statistics for the 11 studied projects²

²Updated as of 2018, the original study was conducted in 2016.

Fine-grained source code changes are not directly available in traditional VCSs, Git included, and we therefore had to extract them based on the pre-change and post-change revisions of the changed Java files (which are available in the VCSs). The task of extracting fine-grained source code changes by comparing two source code files on the abstract syntax tree (AST) level was addressed by the ChangeDistiller (Fluri et al., 2007; S.E.A.L UZH, 2011) and GumTreeDiff (Falleri and Morandat, 2014; Falleri et al., 2014) projects. Both projects share a common trait, they were designed to operate on two ASTs at a time (typically two subsequent versions of a particular class), and do not support analyzing an entire source code repository's commit history. In order to distill (harvest) fine-grained source code changes from an entire repository's commit history, our solution design needed to address two main concerns:

- (1) **Multiple revisions**. In the context of modern VCS systems at any given time there is only one revision of each file available in the working tree of a given source code repository. Branches are either a different directory on the file-system³, or require switching to, in which case they swap the current revision for the new one in-place⁴. Since we are interested in analyzing a given file throughout all its revisions we need to work around this limitation so that for every revision *r* we have the file's revisions *r* and *r* + 1 available to the AST comparison tool.
- (2) Multiple files. A source code repository consists of numerous source code files, created and removed at different points in time throughout the repository's life-cycle. In order to analyze the entire repository an analysis needs to take place for all the source code files (and revisions).

The next stage was to build a mechanism that would replay all the changes made to a given repository according to its commit history so that the fine-grained source code changes could be recorded and repeat this process for every studied repository (see Listing 1). The Git VCS system (Torvalds, 2007), arguably the most popular VCS system in recent years (StackOverflow, 2017, 2018), and the one used by the prevalent repository hosting platform GitHub Inc. (2010), allows one to create a series of patch files, representing the repository's commit history (see also Listing 2). By applying these patches in a chronological order, one can essentially replay the changes made to a source code repository throughout its commit history (see Listing 3 to 5).

Given that we wish to analyze *n* repositories, after downloading (cloning) the repositories from GitHub, for each repository *r* where $1 \le r \le n$ we created a series of patch files $\{p_i^r\}_{i=1}^{N_r}$, where N_r is the latest revision number for repository *r*. We only considered the *master* branch⁵, which is the default branch name in Git. In exceptional cases where the master branch did not exist, we searched for the *trunk* branch, which is the default branch name in Subversion and can sometimes be found in Git repositories that follow Subversion's naming patterns⁶. Each patch file p_i^r is responsible for transforming repository *r* from revision r_{i-1} to revision r_i , where r_0 is the empty repository. By initially setting repository *r* to revision 1 (i.e. the

³As implemented in Subversion, see also http://svnbook.red-bean.com/en/1.7/svn.branchmerge.using.html.

⁴As implemented in Git, see also https://git-scm.com/book/en/v2/Getting-Started-Git-Basics.

⁵See also "Git Branching", https://git-scm.com/book/en/v1/Git-Branching-What-a-Branch-Is.

⁶See also "Recommended Repository Layout", http://svnbook.red-bean.com/en/1.7/svn.tour.importing.html.

1

2

3

4

5

6

7

initial revision) and then applying all patches $\{p_i^r\}_{i=2}^{N_r}$ in a sequential manner, the revision history for that repository is essentially replayed. Conceptually, this is equivalent to having all developers perform their commits sequentially one by one according to their chronological order.

Listing 1. Distilling fine-grained source changes from multiple repositories

```
distillRepos(repos) {
  for(repo in repos) {
    patches = prepareRepo(repo)
    changes = distillPatches(patches)
    write(changes) // persist the distilled fine-grained changes
  }
}
```

Listing 2. Preparing a source code repository for distilling fine-grained source code changes

```
prepareRepo(repo) {
    checkoutRevision(repo, LAST)
    patches = createPatches(repo) // leverage the git-format-patch command
    checkoutRevision(repo, FIRST)
    return orderByPatchId(patches, ASCENDING)
    }
```

Listing 3. Distilling fine-grained source code changes from a sequence of patches

```
distillPatches(patches) {
                                                                                            1
  for(patch in patches) {
                                                                                            2
    beforeAfterPairs = recordFileChanges(patch)
                                                                                            3
     for((revision<sub>i</sub>, revision<sub>i+1</sub>) in beforeAfterPairs) {
                                                                                             4
         currentChanges = distillChanges(revision<sub>i</sub>, revision<sub>i+1</sub>)
                                                                                            5
         changes.add(currentChanges)
                                                                                             6
    }
                                                                                            7
     return changes
                                                                                            8
  }
                                                                                            9
}
                                                                                            10
```

Towards Software Analytics: Modeling Maintenance Activities

```
recordFileChanges(patch) {
  javaFiles = onlyJavaFilesIn(patch)
  beforeContent = readContent(javaFiles)
  applyPatch(patch) // transform the repo to the next revision
  afterContent = recordContent(javaFiles)
  return zip(beforeContent, afterContent)
}
```

Listing 5. Distilling fine-grained source changes from two files, typically a before-and-after pair

```
distillChanges(left, right) {
  return distillerTool.distill(left, right)
}
```

We chose ChangeDistiller to perform the fine-grained source code change extraction (i.e., the distillerTool in Listing 5) due to its popularity in the research community (Falleri et al., 2014; Fluri and Gall, 2006; Fluri et al., 2007, 2008, 2009; Gall et al., 2009; Giger et al., 2011, 2012; Martinez et al., 2013) and its native Java support. ChangeDistiller required that both the before and after revisions of a source code file were present as physical files on the file system to perform the analysis (S.E.A.L UZH, 2014). This design choice presented some challenges in the face of analyzing multiple projects at scale. Fortunately, ChangeDistiller is an open source tool (S.E.A.L UZH, 2011) and we were able to easily obtain the source code and surgically resolve this and other issues we encountered. After the distilling stage was completed, the resulting datasets were manipulated using Apache Spark (Apache Spark, 2016), a state of the art framework for large data processing.

Harvesting a real-world software project may yield a great amount of fine-grained source code changes, easily adding up to millions and dozens of millions of records. Manipulating a dataset of this magnitude is no longer as trivial as inputting it into a spreadsheet or even massaging it in a native R environment (R Development Core Team, 2008). As data sizes have outpaced the capabilities of single machines both in terms of memory capacity and CPU speed, users need new frameworks to scale out their computations. As a result, there has been an explosion of new cluster programming models targeting diverse computing workloads (Zaharia et al., 2016) in the "Big Data" (Diebold, 2012) ecosystem.

Our framework of choice for this work was Apache Spark (Apache Spark, 2016) (henceforth Spark). Spark has one of the largest developer and user communities⁷ and we found its programming model quite intuitive.

1

2

3

4

5

6 7

1

2 3

⁷As indicated by a survey conducted by databricks in 2016, see https://goo.gl/w92BB5.

It also offers a native Scala language (Scala, 2015) application programming interface⁸ (API), which was a great fit in light of the authors' prior experience with Scala.

One of the fundamental abstractions in Spark is the resilient distributed datasets (RDD) (Zaharia et al., 2012). Spark exposes RDDs through a functional programming API where users can pass local functions to run on the cluster (local or distributed). Operations on an RDD are divided into transformations and actions. Transformations derive new RDDs from existing ones, while actions compute and return a concrete result to the program. Spark evaluates RDDs lazily, allowing it to find an efficient plan for the user's computation. In this regard, transformations return a new RDD objects representing the result of a computation but do not immediately compute it. The actual computation takes place when an RDD action is called.

We extensively used Spark to produce *data aggregations* to significantly reduce a dataset's size so it is sufficiently compact to lend itself to interactive exploration in the R environment. Most of our data aggregations begin with reading all the fine-grained source code changes we have already harvested on a per-project basis and stored them as files on disk, see Listing 6. Transformations are highlighted in blue, Scala type annotations are in violet. Type annotations for local variables can often be omitted in Scala, we explicitly provide them in some of the cases for the sake of clarity.

Listing 6. Loading all the fine-grained source code changes from the harvested projects

	val	<pre>sc = new SparkContext() // initiate a Spark context</pre>	1
			2
	val	<pre>perProjectData: Set[RDD[Array[String]]] =</pre>	3
		projects	4
0		.map(prj => sc.textFile(inputNameFor(prj))	5
2		<pre>.map(line => line.split("#")))</pre>	6
			7
3	val	<pre>fineGrainedChanges: RDD[Array[String]] = sc.union(perProjectData)</pre>	8

The variable projects is a collection of project names, over which we iterate and apply a map transformation (see bookmark) in Listing 6) that builds an RDD from each project's fine-grained source code changes stored as text files on disk. Each line in these files is a string concatenation of values separated by a "#" (pound) sign. We split the lines by the pound sign (see bookmark 2) in Listing 6) so that each element in the resulting RDD is of type Array[String]. Since we have multiple projects the perProjectData variable is of type Set[RDD[Array[String]]]. This set of RDDs is then unified into a single RDD for further manipulation using the union operation provided by Spark (see bookmark 3). Each element in this RDD is an array of strings representing parsed lines from the original files. Since RDDs are lazy data structures, no actual processing is done at this point, and it will only take place once an action (e.g., printing, counting, etc.) is invoked on the fineGrainedChanges RDD (as indicated in Listing 10)

⁸Spark provides APIs for a growing number of other programming languages, see https://spark.apache.org/docs/2.3.0/api.html.

Towards Software Analytics: Modeling Maintenance Activities

The aggregations we perform on the fineGrainedChanges RDD usually fall into one of the following categories:

- Per-commit, to explore commit level activity
- Per-developer, to explore developer level activity
- Per-project, to explore project level activity

0 0

• Global, to explore the entire dataset's properties

For example, to compute the frequencies of the different fine-grained source code changes per commit, i.e., how many times each fine-grained source code change appeared in the commits in our dataset we use the code in Listing 7.

Listing 7. Computing the fine-grained source code change frequencies per commit

```
val perCommitFrequencies: RDD[(String, Map[String, Int])] = 1
fineGrainedChanges 2
.groupBy(lineValues => lineValues(COMMIT_ID)) 3
.mapValues(countChanges) 4
```

This computation (Listing 7) uses the groupBy and mapValues transformations. The groupBy transformation takes an element from the collection it is applied on, i.e., fineGrainedChanges, and extracts a key that is used to group all elements with the same key into a single group. Since we would like to compute the frequencies of the different fine-grained source code changes per commit, we first group our records per commit. To accomplish this we specify the key to be the commit id⁹. This groupBy transformation (see bookmark ① in Listing 7) derives a new RDD where each element is a pair of type (String, Iterable[Array[String]]). The first tuple component (a.k.a "key") is the commit id, and the second (a.k.a "value") is a collection of all the elements that had this particular key. Next we apply a mapValues transformation (see bookmark ② in Listing 7) which iterates over these pairs and transforms their value while retaining the key. The transformation logic we provide to mapValues is one that calculates the frequencies of each fine-grained source code change, see Listing 8.

Listing 8. Computing the fine-grained source code change frequencies

<pre>def countChanges(lines: Iterable[Array[String]]): Map[String, Int] =</pre>	1
lines	2
.map(lineValues => lineValues(CHANGE_TYPE))	3
.groupBy(identity)	4
.mapValues(size)	5

The countFrequencies is a method which receives an iterable of lines and represents all changes performed in a given commit, it returns a mapping (Map[String, Int]) between the fine-grained source

⁹Also known as "commit hash" in git, see also https://git-scm.com/book/en/v2/Git-Basics-Viewing-the-Commit-History.

code change type (e.g., "ADDITIONAL_CLASS") and its frequency. Note that countFrequencies does not operate on RDDs but on Scala native collections. One of the benefits of using Spark's Scala API is that it is consistent with Scala's native collections. In particular, the name and semantics of the mapValues and groupBy transformations for Scala collections and Spark RDDs are the same.

First each line is mapped to its corresponding fine-grained source code change type (bookmark ① in Listing 8), then all values are grouped using the identity key extractor (bookmark ② in Listing 8), forming tuples where the key is the fine-grained source code type and the value is a collection of all the corresponding fine-grained source code change types equal to the key. Finally, we map the tuples' values (bookmark ③ in Listing 8) to the sizes of their value component. This results in tuples where the key is the fine-grained source code change type, and the value is the key's frequency. Since the keys in these tuples are the fine-grained source code change types, we end up with a mapping (Map[String, Int]) between the fine-grained source code change type (e.g., "ADDITIONAL_CLASS") and its frequency.

For example, if a project's raw data file contains the following pound separated values:

```
1a2b3c#PARAMETER_INSERT#file1.java
1a2b3c#ADDITIONAL_FUNCTIONALITY#file3.java
1a2b3c#DOC_DELETE#file2.java
1a2b3c#PARAMETER_INSERT#file1.java
1a2b3c#PARAMETER_INSERT#file1.java
1a2b3c#DOC_DELETE#file2.java
```

The map transformation (bookmark 1) in Listing 8) results in:

```
{PARAMETER_INSERT}
{ADDITIONAL_FUNCTIONALITY},
{DOC_DELETE}
{PARAMETER_INSERT}
{PARAMETER_INSERT}
{DOC_DELETE}
```

The groupBy transformation (bookmark 2) in Listing 8) results in:

```
(PARAMETER_INSERT -> {PARAMETER_INSERT, PARAMETER_INSERT, PARAMETER_INSERT})
(ADDITIONAL_FUNCTIONALITY -> {ADDITIONAL_FUNCTIONALITY})
(DOC_DELETE -> {DOC_DELETE, DOC_DELETE})
```

The mapValues transformation (bookmark 3 in Listing 8) results in:

```
(PARAMETER_INSERT -> 3)
(ADDITIONAL_FUNCTIONALITY -> 1)
(DOC_DELETE -> 2)
```

The perCommitFrequencies RDD (see Listing 7) will therefore contain the element:

(1a2b3c -> {PARAMETER_INSERT -> 3, ADDITIONAL_FUNCTIONALITY -> 1, DOC_DELETE -> 2})

Per-developer and per-project aggregations are performed similarly to what we have shown for the per-commit aggregation, the main change being the key passed to the groupBy transformation (see book-marks 1 and 2 in Listing 9). Global operations require no prior aggregations and can be performed directly on the fineGrainedChanges RDD, see Listing 10.

Listing 9. Per-developer and per-project aggregations

```
// aggregate per developer (email) and project
                                                                             1
   val perDeveloper: RDD[((String, String), Iterable[Array[String]])] =
                                                                             2
       fineGrainedChanges
                                                                             3
       .groupBy(entry => (entry(EMAIL), entry(PROJECT)))
0
                                                                             4
                                                                             5
   // aggregate per project
                                                                             6
   val perProject: RDD[(String, Iterable[Array[String]])] =
                                                                             7
       fineGrainedChanges
                                                                             8
       .groupBy(entry => entry(PROJECT))
                                                                             9
0
```

Listing 10. Global operations

```
// unlike previous examples, count() is an "action" which triggers 1
// an actual computation that returns a result to the program 2
// rather than deriving a new RDD 3
val total: Long = fineGrainedChanges.count() 4
```

5 CREATING A GROUND TRUTH DATASET

The first author *manually* classified a randomly sampled set of ~100 commits from each of the studied 11 repositories. To improve classification quality the projects' issue tracking systems, e.g. JIRA (Atlassian, 2014), was often used. The JIRA contained the tickets occasionally referenced in developers' commits messages (e.g., "[PRJ-NAME 1234] Fixed some bug"). Such tickets (a.k.a. issues) typically contain additional information about the feature or bug the referencing commit was trying to address. Moreover, tickets sometimes had their own classification labels such as "feature request", "bug", "improvement" etc., but unfortunately they were not very reliable as developers were not always consistent with their labeling (classification). For instance, in some cases bug fixes were labeled as "improvement", and while fixing a bug is indeed an improvement, according to the maintenance activities we use (Mockus and Votta, 2000), bug fixes should be classified corrective while improvements should be classified perfective. Some developers used the term "fix" even when they referenced feature requests, e.g. "fixed issue #N", where "issue #N" spoke of a new feature or an improvement that did not necessarily report a bug. These observations are consistent with Herzig et al. (2013) who reported that 33.8% of the bug reports they studied were misclassified.

In cases where the lack of supporting information (e.g., not enough information in the corresponding ticket and / or commit message) prevented us from classifying a certain commit with satisfactory confidence, that commit was discarded from the dataset and replaced by a new one, selected randomly from the same project repository (by re-sampling a commit). If we were unable to classify the replacement commit as well, we would repeat this routine until we found a commit that we were able to confidently classify. Further rules of thumb we used for classifying were as follows:

- *Javadoc and comment updates were considered perfective maintenance.* Rational: these changes improve the system.
- *Fixing a broken unit test or build was considered corrective maintenance.* Rational: we assume that tests break in the presence of bugs.
- Adding new unit test(s) was considered perfective maintenance.
 Rational: we assume that new tests improve coverage.
 We conjecture that more often than not, developers who add tests aim to improve system coverage.
- Performance improvements that resulted from an open ticket in the issue tracking system were considered corrective maintenance.
 - Rational: we assume that tickets that were reported on performance issues resulted from pains on the user side, and addressing these pains is more corrective in nature than perfective.
- Performance improvements that did NOT result from an open ticket in the issue tracking system were considered perfective maintenance.

Rational: we assume that developers may occasionally seize an opportunity to improve code performance, however, if there were no users suffering the problem being fixed, we consider the maintenance to be of a perfective nature, rather than corrective one.

We made efforts to avoid class starvation (i.e., not having enough instances of a certain class) by inspecting the proportion of each class within a given sample for a given project. An imbalanced training dataset could substantially degrade models' performance, and in case we detected a considerable imbalance in some project's classes, we added more commits of the starved class from the same project by means of repeatedly sampling and manually classifying commits until a commit of the starved class was found.

To alleviate the challenges involved in reproducing our study we have made our dataset publicly accessible online (Levin and Yehudai, 2017c). This dataset consists of 1151 manually classified commits, 100-115 commits from each of the 11 studied project. Among these commits 43.4% (500 instances) were corrective, 35% (404 instances) were perfective, and 21.4% (247 instances) were adaptive. The commits in this dataset sum up to 33,149 fine-grained source code changes.

In order to inspect manual classification agreement, we randomly selected 110 commits out of the 1151 commits, 10 random commits from each of the 11 projects, and had both authors classify it. At first the agreement stood at 79%. After discussing the conflicts and sharing the guidelines in more detail, the agreement level rose to 94.5%. According to the one sample proportion test (Altman, 1990), the error margin

for our observed agreement level was 4.2%, and the estimated asymptotic 95% confidence interval was [90.3%, 98.7%]. This indicates that both authors were in agreement about the labels for the vast majority of cases once they employed the same guidelines (see Section 5). Regarding some of the commits, no consensus was reached. Consider a commit with the following message: *"add hasSingleArrayBackingStorage allow foroptimization only when there really is a single array, and not when there is- a multi dimensional one"*. One of the annotators had labeled it "Corrective", assuming this commit fixed a bug, while the other had labeled it "Perfective" assuming this was an optimization which improved performance but did not necessarily fix a known bug. Since there was no JIRA ticket associated with this commit it was difficult to ascertain which label is more plausible. Similarly, consider a commit with the message: *"Timeouts for row lock and scan should- be separate*". Based on the message, this commit could be considered any of the maintenance activities, it could be fixing a bug, improving design (by separating concerns) or adding a new feature (e.g., allowing different timeouts for lock and scan). In this particular case, the referenced JIRA ticket indicated it was an "improvment" and thus "Perfective", but had it not been for the JIRA ticket it would have been quite challenging to determine the associated maintenance activity.

6 COMMIT CLASSIFICATION MODELS

We performed our statistical computations in the R statistical environment (R Development Core Team, 2008), where we extensively used the R caret package (Kuhn, 2017; Kuhn et al., 2017) for the purpose of model training and evaluation.

We split the labeled dataset into a training dataset and a test dataset, 85% and 15% respectively, in order to have the test dataset completely isolated from any training procedures. The split was performed by using R's *createDataPartition* function (Kuhn, 2018), with the percentage of data that goes to training set to 0.85. The createDataPartition function uses random sampling within the labels (Corrective, Perfective, Adaptive) in an attempt to balance the class distributions within the splits, see also Table 3 for a detailed description of the train and test splits.

Dataset	Corrective	Perfective	Adaptive
Train (979/1151 instances)	425	344	210
Test (172/1151 instances)	75	60	37

Table 3. Number of instances per class in the train and test datasets in our study¹⁰

Project	Commits in test dataset
RxJava	14
Restlet	17
Drools	16
HBase	17
Spring Framework	16
OrientDb	14
Hadoop	13
Camel	13
Elasticsearch	15
Kotlin	17
Intellij Community Edition	20

Table 4. Number of commits per studied project in the test dataset

The model training phase consists of using 5 time repeated 10-fold validation for each compound model on the training dataset (which boils down to performing a 10-fold cross validation process 5 different times and averaging the results). Then, the trained models were evaluated using the test dataset - the 15% split that did not take part in the model training process.

6.1 Utilizing word frequency analysis

First we classified the test dataset (the 15% of the entire labeled dataset) using a naive method to set an initial baseline. The naive method is based on a classification technique described in our previous work (Levin and Yehudai, 2016), and consists of searching for pre-defined words (see Table 5), and assigning the most frequent class (i.e., corrective) in case none of the keywords were present in the commit message, see Table 6 for more details. Assigning the most frequent class to an instance is far from ideal, however, when models find no features to rely on, using the overall distribution of the training dataset is a common technique (also called 'No Information Rate', see Section 3.1).

¹⁰The entire labeled dataset, consisting of 1151 labeled commits, is publicly available at https://doi.org/10.5281/zenodo.835534, see also Levin and Yehudai (2017c).

Table 5. Stemmed keywords used	by the "naive method"	' as described in (Levir	ı and Yehudai, 2016)
--------------------------------	-----------------------	--------------------------	----------------------

Corrective	fix, esolv, clos, handl, issue, defect, bug, problem, ticket
Perfective	refactor, re-factor, reimplement, re-implement, design, replac, modify, up- dat, upgrad, cleanup, clean-up
Adaptive	add, new, introduc, implement, extend, feature, support

true class classified as	Adaptive	Corrective	Perfective	
Adaptive	18	2	16	
Corrective	18	72	37	
Perfective	1	1	7	
Recall:	48%	96%	11%	
Precision:	50%	56%	77%	
Accuracy:		56%		
Карра:		29%		
F1 Score (micro-averaged):	0.56			
F1 Score (macro-averaged):	0.46			
No Information Rate (NIR):	43%			
P-Value [Accuracy > NIR]:	: 0.0005			

Table 6. Naive model's confusion matrix

The results showed that 34.8% of the commits in the test dataset (60 commits) did not have any of the keywords present in their commit message, and were therefore automatically classified corrective. In addition, the low recall of the perfective class was particularly notable, as opposed to the high recall of the corrective class (which accounts for most of the commits in the classified dataset). The noticeable difference

between the micro-averaged and macro-averaged F1 scores, 0.56 vs. 0.46 respectively, also indicates that the current model (based on the naive method) does not perform equally well for all classes.

The high percentage of commits without any keywords prompted us to try to fine-tune the keywords we were searching for. We performed an additional experiment using the same classification method, only this time the keywords were obtained by employing a word frequency analysis and normalization for the commit messages. This time 28% of the commits did not have any of the keywords present in their commit message. These findings led us to believe that the high number of commit messages containing none of the keywords could be playing a significant role in determining the overall classification quality.

6.2 Utilizing source code changes

Techniques for dealing with missing values in classification problems are broadly covered by Saar-Tsechansky and Provost (2007), who describe two common methods used to overcome such issues: (1) imputation, where the missing values are estimated from the data that are present, and (2) reduced-feature models, which employ only those features that will be known for a particular test case (i.e., only a subset of the features that are available for the entire training dataset), so that imputation is not necessary. Since our dataset consists of two different data types, keywords and source code changes, we use reduced-feature models, which are reported to outperform imputation and represent our use-case more naturally. In addition, since the missing feature patterns in our dataset are known in advance, i.e., given a commit only the keywords can be missing, its source code changes are always present, we can pre-compute and store two models; one to be used when all features are present (keywords + source code changes), and the other when only a subset is available (source code changes only). We define the notion of a **compound** model (similarly to the "classifier lattice" described by Saar-Tsechansky and Provost) which uses two separate models for classifying commits with, and without (pre-defined) keywords in their commit message. The classify routine of the compound model is pseudo-coded in Listing 11.

Listing 11. The compound model's classify routine

1
2
3
4
5
6
7

Given a commit *C*, the compound model first checks if *C*'s commit message has any keywords, if so, the model defined as *model*_{*KW*} is used to classify *C* (see bookmark) in Listing 11), otherwise (i.e., no keywords found in *C*'s commit message), the model defined as *model*_{*KW*} is used to classify *C* (see bookmark) in Listing 11).

Each of the models $model_{KW}$ and $model_{\overline{KW}}$ may or may not be a reduced-feature model, depending on whether it employs the full set of features (both keywords and source code changes), or only a subset of it (either keywords or source code changes).

We define $model_{\overline{KW}}$ and $model_{KW}$ to be one of the following model types:

- <u>Keywords</u> model, which relies solely on keywords to classify commits. The features used by this model are keywords obtained by performing the following transformations on the commit message field:
- (1) Stripped special characters
- (2) Made lower case (case-folding)
- (3) Stripped English stopwords
- (4) Stripped punctuation
- (5) Striped white-spaces
- (6) Performed stemming
- (7) Adjusted frequencies so that each comment can contribute a given word only once
- (8) Stripped custom words such as developer names, projects names, VCSs lingo (e.g., head, patch, svn¹¹, trunk, commit), domain specific terms (e.g., http, node, client): "patch", "hbase", "checksum", "code", "version", "byte", "data", "hfile", "region", "schedul", "singl", "can", "yarn", "contribut", "commit", "merg", "make", "trunk", "hadoop", "svn", "ignoreancestri", "node", "also", "client", "hdfs", "mapreduc", "lipcon", "idea", "common", "file", "ideadev", "plugin", "project", "modul", "find", "border", "addit", "changeutilencod", "clickabl", "color", "column", "cach", "jbrule", "drool", "coprocessor", "regionserv", "scan", "resourcemanag", "cherri", "gong", "ryza", "sandi", "xuan", "token", "contain", "shen", "todd", "zhiji", "tan", "wangda", "timelin", "app", "kasha", "kashacherri", "messag", "spr", "camel", "http", "now", "class", "default", "pick", "via".
- (9) We then selected the 10 most frequent words from each of the three maintenance activities in the test dataset:
 - Corrective: (1) fix (2) test (3) issu (4) use (5) fail (6) bug (7) report (8) set (9) error (10) npe
 - Perfective: (1) test (2) remov (3) use (4) fix (5) refactor (6) method (7) chang (8) add (9) improv (10) new
 - Adaptive: (1) support (2) add (3) implement (4) new (5) allow (6) use (7) method (8) test (9) set (10) chang

It can be seen that some of the words (as obtained by our commit message word frequency analysis) overlap between maintenance activities. The words "test" and "use" appear in all three maintenance activities; the word "fix" appears in both the corrective and perfective maintenance activity; the words "method", "chang", "add" and "new" appear both in the perfective and adaptive maintenance activities. These word overlaps may indicate that keywords alone are insufficient to accurately classify commits

¹¹Subversion is commonly abbreviated to SVN after its command name svn.

into maintenance activities, and need to be augmented with additional information in order to improve classification accuracy.

For the purpose of building the Keywords model type, we remove multiple occurrences of the same word (so that each word appears only once in the combined list) and remain with the following set of words: (1) add (2) allow (3) bug (4) chang (5) error (6) fail (7) fix (8) implement (9) improv (10) issu (11) method (12) new (13) npe (14) refactor (15) remov (16) report (17) set (18) support (19) test (20) use.

- (Source Code) Changes based model, which relies solely on source code changes to classify commits. The features used by this model are source code change types (Fluri and Gall, 2006) obtained by distilling commits, as described earlier in this section.
- <u>Combined</u> (Keyword + Source Code Change Types) model, which uses both keywords and source code change types to classify commits. The features used by this type of models consist of both keywords and source code change types.

A word-cloud visualization of the keyword distribution in each of the maintenance activities can be found in Figure 3, Figure 4, Figure 5. A summary of the model components can be found in Table 7.

Model Type	Model Features
Keywords	Words
Changes	Fine-grained Source Code Change Types
Combined	Words + Fine-grained Source Code Change Types

Table 7. Reduced-feature mo	odel components
-----------------------------	-----------------

For example, a commit where two methods were added (fine-grained source code change type "additional_functionality"), and one statement was updated (fine-grained source code change type "statement_updated") and has a commit message that says "Refactored blob logic into separate methods" will be treated differently by each of the model types indicated in Table 7.

The Keywords model extracts features represented by tuples of size 20, and given the commit above would 20extract the following features: $(0 \dots 1 \dots 1 \dots 0)$ with "1" in the coordinates that represent the words *"refac-tor"* and *"method"*. The count of each keyword is at most one, i.e., duplicate keywords are counted only once. Source code changes are ignored, since the Keywords model type does not consider source code changes. The Changes model extracts features represented by tuples of size 48 (since there are 48 different source code

change types), and given the commit above would extract the following features: $(0 \dots 2 \dots 1 \dots 0)$ with "2"

in the coordinate that represents the fine-grained source code change type "additional_functionality" and "1" in the coordinate that represents "statement_updated". In contrast to the case of the Keywords model, all occurrences of every fine-grained source code change type are counted in. Keywords in the commit message are ignored, since the Changes model type does not consider keywords.

The Combined model extracts features represented by tuples of size 68 (= 48 fine-grained source code change types + 20 keywords), and given the commit above would extract the following features: $\frac{68}{68}$

 $(\underbrace{0\ldots 1\ldots 1\ldots 0}_{20}\underbrace{\ldots 0\ldots 2}_{48})$, with "2" in the coordinate that represents the fine-grained source

code change type "additional_functionality", and "1" in the coordinates that represent the fine-grained source code change type "statement_updated", the keyword "refactor", and the keyword "method". The Combined model type captures both keywords and fine-grained source code change types - hence its name.

In the next sections we evaluate and compare different compound models by considering the different combinations of their $model_{KW}$ and $model_{\overline{KW}}$ model components. The evaluation process consists of the following steps:

- (1) Select the model component $model_{KW}$
- (2) Select the model component $model_{\overline{KW}}$
- (3) Select an underlying classification algorithm for the compound model, which determines the algorithm to be used by each of the model components $Model_{KW}$ and $Model_{\overline{KW}}$ (J48, GBM, or RF, see also Section 3.1).

7 EVALUATION

We describe an exhaustive set of combinations for selecting the pair of $(Model_{KW}, Model_{\overline{KW}})$ models in Table 8, where the pairs can be one of the three model types defined in Table 7. Each row in Table 8 represents a compound model, defined by the selection of $(Model_{KW}, Model_{\overline{KW}})$. The classification accuracy and Kappa achieved by a given compound model are reported in the corresponding Accuracy and Kappa columns. The best performing compound model for each classification algorithm is highlighted in lime-green, and the keywords based model (where both $Model_{KW}$ and , $Model_{\overline{KW}}$ are of the Keywords model type) is highlighted in orange so that it can be easily compared to compound models that utilize fine-grained source code changes.

Alg.	<i>Model_{KW}</i>	$Model_{\overline{KW}}$	Accuracy	Карра
	Com	bined	69.0%	51.7%
	Combined	Keywords	67.7%	50.2%
	Combined	Changes	69.2%	51.9%
	Keywords	Combined	69.8%	53%
J48	Keyv	words	68.5%	51.5%
	Keywords	Changes	69.9%	53.2%
	Changes	Combined	48.7%	20.1%
	Changes	Keywords	47.4%	17.2%
	Cha	nges	48.8%	18.6%
	Com	bined	72.0%	56.2%
GBM	Combined	Keywords	69.0%	51.8%
	Combined	Changes	72.0%	55.9%
	Keywords	Combined	71.6%	56.0%
	Vorm		68 5%	51 4%
GBM	Keyv	voras	00.370	J1.470
GBM	Keywords	Changes	71.5%	55.6
GBM	Keywords Changes	Changes Combined	71.5% 54.1%	55.6 26.9%
GBM	Keywords Changes Changes	Changes Combined Keywords	71.5% 54.1% 51.0%	55.6 26.9% 22.4%
GBM	Keywords Changes Changes Changes	Changes Combined Keywords nges	71.5% 54.1% 51.0% 54.3%	51.4% 55.6 26.9% 22.4% 26.9%
GBM	Keywords Changes Changes Changes Cha	Changes Combined Keywords nges bined	00.3% 71.5% 54.1% 51.0% 54.3% 73.1%	55.6 26.9% 22.4% 26.9% 57.8%
GBM	Keywords Changes Changes Cha Com Combined	Changes Combined Keywords nges bined Keywords	00.3% 71.5% 54.1% 51.0% 54.3% 73.1% 69.5%	55.6 26.9% 22.4% 26.9% 57.8% 52.6%
GBM	Keywords Changes Changes Cha Com Combined Combined	Changes Combined Keywords nges bined Keywords Changes	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7%
GBM	Keywords Changes Changes Cha Com Com Com Com Som Com Som Com Som Com Som Com Som Com Som Com Som Com Som Changes	Changes Combined Keywords nges bined Keywords Changes Changes	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9% 72.2%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7% 56.4%
GBM	Keywords Changes Changes Cha Com Combined Combined Keywords Keywords	Changes Combined Keywords nges bined Keywords Changes Changes Changes	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9% 72.2% 73.6%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7% 56.4% 58.9%
GBM	Keywords Changes Changes Changes Changes Combined Combined Keywords Keywords	Changes Combined Keywords nges bined Keywords Changes Changes Changes	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9% 72.2% 73.6% 69.8%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7% 56.4% 53.4%
GBM RF	Keywords Changes Changes Changes Changes Combined Combined Keywords Keywords Keywords	Changes Combined Keywords Iges Keywords Keywords Changes Changes Changes Combined	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9% 72.2% 73.6% 69.8% 54.5%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7% 56.4% 53.4% 26.6%
GBM	Keywords Changes Changes Changes Changes Combined Keywords Keywords Keywords Keywords	Changes Combined Keywords Iges bined Keywords Changes Changes Changes Combined Xords	71.5% 54.1% 51.0% 54.3% 73.1% 69.5% 71.9% 72.2% 73.6% 69.8% 54.5% 50.6%	51.4% 55.6 26.9% 22.4% 26.9% 57.8% 52.6% 55.7% 56.4% 53.4% 26.6% 21.1%

Table 8. Training dataset compound models performance

Alg.	Min.	1-st Q.	Median	Mean	3-rd Q.	Max.
J48	60.8%	66.4%	70.1%	69.9%	73.4%	80.6%
GBM	60.8%	69.2%	72.1%	72.0%	75.2%	80.8%
RF	65.6%	70.4%	73.4%	73.6%	76.6%	82.8%

Table 9. Training dataset accuracy, best model per algorithm

Table 10. Training dataset Kappa, best model per algorithm

Alg.	Min.	1-st Q.	Median	Mean	3-rd Q.	Max.
J48	38.4%	47.9%	53.4%	53.2%	58.8%	69.7%
GBM	38.3%	51.8%	56.9%	56.2%	60.6%	70.0%
RF	45.5%	54.1%	58.6%	58.9%	63.3%	73.5%

Following our main research questions (see Section 1), the accuracy and Kappa results for each compound model during the training (see Table 8) reveal that the compound models that use either $Model_{\overline{KW}} = Combined$ or $Model_{\overline{KW}} = Changes$ achieve higher accuracy and Kappa when compared to models with the same $Model_{\overline{KW}}$ component but with $Model_{\overline{KW}} = Keywords$, regardless of the underlying classification algorithm (J48, GBM or RF). This comes as no surprise, as one could expect keyword based models would have trouble accurately classifying commits that do not have any keywords in their commit message. Table 8 also reveals that models that rely solely on commit messages have higher accuracy and Kappa than models that rely solely on fine-grained source code changes (under all three algorithms).

Further accuracy and Kappa statistics pertaining to the training stage of the best performing model for each algorithm can be found in Table 9 and Table 10 respectively. From Table 9 and Table 10 we can learn that during the training stage, the RF model consistently outperforms the J48 and even the GBM model, in both accuracy and Kappa, across all of the cuts: minimum, 1-st quartile (25-th percentile), median, mean, 3-rd quartile (75-th percentile) and maximum. In particular, the minimum accuracy and Kappa of the RF are notably higher than its competitors.

A comparison between the best compound models from each of the underlying classification algorithm category can be found in Figure 1. The top performing models were then used to classify the test dataset, consisting of 15% of the entire labeled dataset, see Table 11. The ultimate winner was the RandomForest compound model with $Model_{KW} = Keywords$ and $Model_{\overline{KW}} = Combined$. A detailed confusion matrix for this champion model can be found in Table 12.



Fig. 1. Training dataset accuracy and Kappa, best model by algorithm

Table 11.	Test dataset classificat	ion performance
-----------	--------------------------	-----------------

Algorithm	$Model_{KW}$	$Model_{\overline{KW}}$	Accuracy	Kappa
J48	Keywords	Changes	70 %	53%
GBM	Com	bined	72 %	57 %
RF	Keywords	Combined	76 %	63 %

true class classified as	Adaptive	Corrective	Perfective		
Adaptive	28	5	5		
Corrective	6	63	14		
Perfective	3	7	41		
Recall:	75%	84%	68%		
Precision:	73%	75%	80%		
Accuracy:	76%				
Kappa:	a: 63%				
F1 Score (micro-averaged):	0.76				
F1 Score (macro-averaged):	0.76				
No Information Rate (NIR):		43%			
P-Value [Accuracy > NIR]:		$< 2e^{-16}$			

Table 12. RF based Keywords-Combined compound model's confusion matrix for the test dataset

The decision tree built by the J48 algorithm for our keyword based model (see Figure 2) provides some interesting insights regarding its classification process. The word "fix" is the single most indicative word of corrective commits, which aligns well with our intuition, according to which commits that fix faults are likely to include the "fix" noun or verb in the commit message. Given that "fix" did not appear, the words "support" and "allow" are most indicative of adaptive commits, presumably these words are used by developers to indicate the support of a new feature, or the fact that something new is now "allowed" in the system. The combination "implement chang" (stemmed), given that "fix", "support" and "allow" did not appear, is very indicative of either perfective or corrective commits, if however, "implement" is not accompanied by the word "chang" (stemmed), the commit is likely to be adaptive. The (stemmed) word "remov", given that the words "fix", "support", "allow" and "implement" did not appear, is very indicative of an evelopers often use it to describe a modification where they remove an obsolete mechanism in favor of a new one.



Fig. 2. A J48 Keywords model type ("a" stands for adaptive, "c" for corrective, and "p" for perfective)

We also visualized the keyword frequency in maintenance activities using a word-cloud (see Figure 3, Figure 4, Figure 5), which revealed that the word "test" is particularly common in perfective commits, but is generally common in all three maintenance activity types. The word "use" is also common in all three maintenance activity types, but is particularly frequent in the perfective maintenance activity. The words "fix", "remov" and "support" are quite distinctive of their corresponding maintenance activity types: corrective, perfective and adaptive (respectively). The word "add" is common in adaptive commits, as well as "allow".



Fig. 3. Word-cloud for the "Corrective" maintenance activity



Fig. 4. Word-cloud for the "Perfective" maintenance activity



Fig. 5. Word-cloud for the "Adaptive" maintenance activity

Similarly, we visualized the fine-grained source code changes frequencies using a source-code-change-type-cloud which revealed that statement related changes, e.g., "statement_insert", "statement_update" and "statement_delete" are the most common change types in all three maintenance activities (corrective,

perfective, adaptive). The fine-grained source code change type "additional_functionality" is common in both perfective and adaptive commits, but less so in corrective commits.

The term-cloud and J48 keyword based decision tree visualizations provide an intuition for why J48 is likely to outperform a simple word-frequency based classification. In contrast to the word-cloud, which provides "flat" frequencies, the J48 is capable of capturing information pertaining to the presence of multiple keywords in the same commit message, as indicated by the decision tree.

We depict the 20 most important predictors for our champion RF model in Table 13. The rank score is scaled from 0 to 100 and is based on the contribution each predictor makes towards the quality of the RF classification model. Not all predictors are equally important for all three maintenance activities. Some play a bigger role in classifying one maintenance activity over the others. It is worth noting that numerous fine-grained source code changes are ranked high in the list, which confirms their contribution to the model's quality.

Feature (keyword/fine grained source code change)	Addaptive	Corrective	Perfective
fix	100.00	100.00	90.42
ADDITIONAL_FUNCTIONALITY	75.72	72.07	75.72
STATEMENT_INSERT	62.17	40.19	62.17
support	54.92	54.92	53.20
ADDITIONAL_OBJECT_STATE	42.27	42.27	38.01
add	32.00	32.00	25.36
ALTERNATIVE_PART_INSERT	30.04	16.71	30.04
remov	27.47	23.46	27.47
DOC_UPDATE	22.51	26.77	26.77
test	26.60	13.62	26.60
STATEMENT_DELETE	25.98	25.98	17.37
REMOVED_FUNCTIONALITY	15.89	25.51	25.51
implement	22.97	22.97	18.60
COMMENT_INSERT	22.36	20.48	22.36
PARAMETER_INSERT	20.74	20.74	18.76
issu	15.04	18.05	18.05
REMOVED_OBJECT_STATE	12.01	17.98	17.98
allow	17.11	17.11	14.15
new	15.78	15.78	10.66
ADDITIONAL_CLASS	15.44	15.44	10.82

Table 13. The 20 most important features in the best RF compound model, the score is scaled from 0 to 100

8 APPLICATIONS

Lehman's Laws teach us that a software system will become progressively less satisfying to its users over time, unless it is continually adapted to meet new needs. The field of software evolution research can be classified into two groups, the first considers the term evolution as a verb while the second as a noun (Lehman et al., 2000). The *verbal* view is concerned with the question of "how", and focuses on means, processes, activities, languages, methods and tools required to effectively and reliably evolve and maintain a software system. The *nounal* view is concerned with the question of "what" and investigates the nature of software evolution, as a phenomenon, and focuses on the nature of evolution, its causes, properties, characteristics, consequences, impact, management and control. Both views are mutually supportive (Lehman and Ramil, 2003; Lehman et al., 2000). Moreover, they advocate that the verbal view research will benefit from progress made in studying the nounal view, and both are required if the community is to advance in mastering software evolution. We follow this thinking and put forth two applications.

8.1 Software Maintenance Activity Explorer

In the spirit of the verbal view (Lehman et al., 2000) which focuses on studying the means, methods and tools required to effectively evolve a software system, we implement a tool for exploring software maintenance activities aimed to assist practitioners. The Software Maintenance Activity Explorer tool (Levin, 2017) is aimed at providing an intuitive visualization of software maintenance activities over time. We believe this visualization may be useful to project and team managers who seek to recognize inefficiencies and monitor the health of a software project and its corresponding source code repository. The Software Maintenance Activity Explorer was built with Few's (2009) and Cleveland's (1985) principles in mind, which advocate for encoding data using visual cues such as variation in size, shape, color, etc'. We chose stacked bar diagrams to visualize data since they allow for an easy comparison both between maintenance activities within a given time frame (e.g., what maintenance activity dominated a given time frame), and between different time frames (e.g., which of the time frames had more maintenance of a given type). In addition, bar diagrams allow users to quickly detect anomalies such as peaks and deeps in one maintenance activity or another compared to past periods.

Project Activity Visualization. The project activity visualization (see Figure 6) allows users to examine the volumes of the different maintenance activities over time, and can be sliced and diced according to a specified date range and an activity period (e.g., from date x until date y, in time frames of 28 days). The stacked bar plot allows for an easy comparison between the maintenance activity types, as well as trend detection.

Developer Activity Visualization. The developer activity visualization (see Figure 7) is a segmentation of the data by a specific developer. Users can examine the data for a specific developer, adjusting the period of interest and date range. Developers identity can be determined by their name, email or both, a feature that

can be useful when developers perform commits using different emails, e.g., when working on an open source project from both their private account and their cooperate account.

Publicly Accessible Data. The Software Maintenance Activity Explorer's about page provides an option to explore the data in-line (see Figure 8), or download it in a CSV format for an offline analysis.

Publicly Accessible Code. The code for this tool is publicly available on GitHub (Levin, 2018).



Fig. 6. Software Maintenance Activity Explorer's project activity tab

Towards Software Analytics: Modeling Maintenance Activities



Fig. 7. Software Maintenance Activity Explorer's developer activity tab

Explore raw data	Lownload raw data				Search:	
CommitId	⇒	Project 🔶	Contributor	Date 🔶	MaintenanceType 🔶	Comment
All		•	All	-	All	All
32d5da5dc8947b31	c289588c85937d8f89612371	kotlin		2010-12-19	Corrective	comment/uncomment actions for the jet
346c0fd980d6723b	3e74fc9d5e80dbc660a68ab3	kotlin		2010-12-20	Corrective	Recovery fixed for properties and decomposers
c614d1dadc23d5cd	bd664c4316415cbd84c1de81	kotlin		2010-12-20	Adaptive	Intermediate implementation for constructors
6361d0eb0a793152	9452141a4a6a0d8cc4bb0b42	kotlin		2010-12-20	Corrective	Functions
ba1cc369f821fbab9	ceef073ad42f4b419cbd7e2	kotlin		2010-12-20	Corrective	Extensions
6f241b7f9f97a3b67f	31f9f115805b621518556b2	kotlin		2010-12-20	Perfective	Parsing property declarations by looking ahead
4d54f704b9962a84f	4274aee67125737d66f3b2b	kotlin		2010-12-21	Corrective	Soft keywords fixed
7547b36eb0aff03d0	1c76b09433bd4e737385e02	kotlin		2010-12-21	Perfective	Class objects
198c46528c1c6e79	669e8fbc9c65fc50f84002a	kotlin		2010-12-21	Corrective	More tests, minor fixes
b521a7a244dbf8bfc	1d4ece09d0b590ab7b4aeb4	kotlin		2010-12-21	Perfective	Tests for basic class member cases
Showing 21 to 30 of 1	12,492 entries		Prev	ious 1	3 3 4 5	11250 Next

Fig. 8. Software Maintenance Activity Explorer's data exploration tab

We conjecture that a balanced maintenance activity profile, i.e., a profile which includes all three maintenance activity kinds (corrective, perfective, adaptive) may help developers be more effective and engaged with the project they work on. It may also be the case that different project managers will choose different thresholds for what a balanced (or unbalanced) profile is, in the context of their project. Nonetheless, once these thresholds have been set our method provide means to identify opportunities for improvement. This may be of particular interest in open source projects, which tend to heavily rely on community efforts. To that end, well balanced maintenance activity profiles may be something the community needs to drive development forward and ensure that the project gets a fair share of new features, bug fixing, and design improvements - activities which tend to compete for resources in real-world scenarios.

We use our dataset and the software maintenance activity explorer to identify homogeneous activity profiles, i.e., profiles of developers who performed only *one* kind of maintenance activity, see Figure 9a and Figure 9b.



Fig. 9. Maintenance activity profiles for two developers from the Kotlin project

The visualization offered by our tool makes it easier to identify these homogeneous maintenance activity profiles and encourage developers to take on a more varied set of tasks. We performed the homogeneous maintenance activity profiles test for 10 projects (see also Table 2) in our study and report the results in Table 14. According to our data, the Camel project had an extremely low portion of homogeneous maintenance activity profiles. It may be the case that Camel's contributors were indeed developers who were inclined towards heterogeneous maintenance activities. Alternatively, one could suggest a number of possible scenarios. It is possible that the Camel project had a significant number of contributors whose contribution to the project did not include Java code, i.e., it revolved around documentation, configuration files, and so forth. This would mean that the percentage of homogeneous maintenance profiles is actually

higher and it might be best to compute it by considering only Java contributors. Another possibility is that the number of contributors to the Camel project significantly increased since we had originally processed its commit history. In which case it would be necessary to re-collect and re-process the project's data to produce a more accurate result.

Our analysis indicates that homogeneous maintenance activity profiles were not uncommon in the projects we inspected (see Table 14). We believe that unbalanced (i.e., where a significant disproportion between maintenance activities is present), and homogeneous maintenance activity profiles in particular, are an opportunity for managers to reach out to developers and suggest taking on tasks that will balance their maintenance activity profiles. A possible way to identify suitable tasks would be using projects' task management systems (e.g., a JIRA system) which provide contextual and detailed information about the available tasks. We also hope that this kind of tool will empower both managers and developers to monitor the ongoing maintenance activities and assist in keeping them varied and balanced. Moreover, such a tool may serve as an alerting mechanism in situations which call for special attention, e.g., when the proportion of unbalanced maintenance profiles exceeds a given threshold.

	Corrective only	Perfective only	Adaptive only	Homogeneous contributors (% of total, truncated)	Total Contributors ¹³
Restlet	11	3	2	41%	39
Drools	22	20	8	36%	137
OrientDb	16	14	4	28%	120
Spring Framework	25	18	30	25%	291
RxJava	14	24	8	19%	211
Hbase	8	16	8	16%	189
Elasticsearch	79	54	33	14%	1103
Kotlin	14	10	6	14%	356
Hadoop	4	12	1	12%	137
Camel	2	1	0	<1%	410

Table 14. Homogeneous maintenance activity profile statistics¹²

¹²Due to certain technical difficulties we had to exclude the IntelliJ Community Edition project from homogeneous maintenance activity analysis.

¹³The total number of contributors is updated as of 2018, maintenance activity profiles were computed as part of the original study conducted in 2016.

8.2 Utilizing Software Maintenance Activities to Model Test Counts

In the spirit of the nounal view (Lehman et al., 2000) which investigates the nature of software evolution as a phenomenon, we conduct a study which leverages our method to demonstrate the importance of maintenance activities for modeling the number of tests in a software project (see Section 8.2).

Automated testing, and automatic unit tests (Hamill, 2004) in particular, is a popular technique for improving software quality. As this technique is gaining popularity and becoming ubiquitous among practitioners it is beneficial to have a good understating of its nature, which as it turns out can be alluding. Beller et al. (2015) conducted a large-scale field study, where 416 software engineers were closely monitored over the course of five months. Their findings indicate that software developers spend a quarter of their work time engineering tests, whereas they think they test half of their time.

In our previous work (Levin and Yehudai, 2017a) we studied 61 open source projects (Levin and Yehudai, 2017d) and established a connection between maintenance activities and test (method and classes) counts in software projects. In this section we extend our previous results and focus on the viability of maintenance activities to modeling the number of test methods and test classes in a software project.

The generalized regression models (GLM, McCullagh and Nelder (1989); Venables and Ripley (2013)) we devised were of the following form:

$$Test^{M}(prj) = Constant^{M} + \sum_{i=1}^{|Predictors|} (coeff_{i}^{M} * predictor_{i}^{M}(prj))$$

where: $M \in \{Methods, Classes\}$ is the test metric we model; *Predictors* is the set of predictors; $coeff_i^M$ are the predictor coefficients; $predictor_i^M(prj)$ are predictor values; and $Constant^M$ is the model constant.

The corresponding models for *Test*^{*Methods*} and *Test*^{*Classes*} can be found in Table 15.

All predictors were log transformed to alleviate skewed data, a common practise when dealing with software metrics (Camargo Cruz and Ochimizu, 2009; Shihab, 2012). Statistically significant predictors of interest are highlighted in lime-green, and the standard error is reported in parenthesis below the estimated coefficients. In addition to the variables we are directly interested in, such as the log(corrective), log(perfective) and log(adaptive) we also use log(LOC), log(age) and log(developers) as control variables, in order to reduce the effect of lurking variables which correlate both with the predictors and the predicted (outcome) variable. Control variables are highlighted in light-bisque.

The ANOVA type-II analysis computes the changes in the model given any single predictor is dropped and it therefore does not depend on the order of the predictors in the model. Employing ANOVA type-II analysis helps in avoiding situations where regression models may lead to the conclusion that certain predictors possess greater explanatory powers than others only because they appear first (Hassan, 2017). The ANOVA type-II analysis for the predictive models *Test*^{Methods} and *Test*^{Classes} can be found in Table 16 Table 17 respectively. Each row indicates the change in the residual deviance and the "AIC" measure (Akaike information criterion, an estimator of the relative quality of statistical models) induced by removing a given predictor from the model. The statistical significance for each row is indicated in the rightmost column. By inspecting the "AIC" column in Table 16 Table 17 we learn which predictors can be *excluded* in order to achieve a lower (better) AIC. By inspecting the "Deviance" column we learn a given predictor's contribution to "explaining" the predicated variables. The "base" model's deviance and AIC are indicated in the "none" row.

It is statistically significant that removing the *log(corrective)* predictor will result in the model's deviance rising from 72 to 95 and its AIC rising from 993 to 1,015. Higher deviance indicates that the new model will have less explanatory power, and higher AIC indicates that it will be worse than the one it is compared to, i.e., the model where the *log(corrective)* was present. Similar arguments can be applied to the *log(perfective)* predictor. The ANOVA analysis confirms that both perfective and corrective maintenance activities are vital to the model, and an attempt to remove either will significantly and adversely affect the model's quality.

Also worth noting is the LOC predictor, its AIC and deviance indicate that it demonstrates statistically significant high explanatory power in both predictive models. This implies that the size of the project has a considerable effect on the number of test methods and test classes it contains.

	Predicted variable:		
Predictor	$Test^{Methods}$	$Test^{Classes}$	
log(corrective)	-1.696***	-1.351***	
	(0.314)	(0.285)	
log(perfective)	1.621***	1.583***	
	(0.397)	(0.358)	
log(adaptive)	-0.247	-0.173	
	(0.366)	(0.329)	
log(developers)	0.318*	0.105	
	(0.182)	(0.163)	
log(LOC)	1.189***	1.053^{***}	
	(0.171)	(0.154)	
log(age)	0.770***	0.686***	
	(0.205)	(0.185)	
Constant	-12.326^{***}	-13.289***	
	(1.873)	(1.702)	
Number of observations	61	61	

Table 15. Negative Binomial GLM for test method and test class counts (Levin and Yehudai, 2017a)

*p<0.1; **p<0.05; ***p<0.01

	Df.	Deviance	AIC	F value	$\Pr(>F)$
<none></none>		72.311	993.480		
log(corrective)	1	95.903	1,015.071	17.617	0.0001***
log(perfective)	1	86.599	1,005.767	10.669	0.0018**
log(adaptive)	1	72.762	991.930	0.336	0.564
log(developers)	1	75.007	994.175	2.013	0.162
log(loc)	1	106.675	1,025.843	25.661	5.077 <i>e</i> – 06***
log(age)	1	84.408	1,003.576	9.033	0.0040**

 Table 16. ANOVA for Test^{Methods}

Table 17. ANOVA for Test^{Classes}

	Df.	Deviance	AIC	F value	$\Pr(>F)$
<none></none>		71.873	812.533		
log(corrective)	1	89.012	827.673	12.877	0.0007***
log(perfective)	1	87.015	825.676	11.377	0.0013***
log(adaptive)	1	72.125	810.786	0.189	0.665
log(developers)	1	72.256	810.916	0.288	0.594
log(loc)	1	105.557	844.217	25.308	5.75 <i>e</i> – 06***
log(age)	1	84.017	822.678	9.124	0.0038**

Following the insights provided by these test regression models, we performed a deeper inspection of two outlier projects, "XPrivacy" and "Omni-Notes" (see Figure 10), that had extremely high values of corrective activity (per 1 LOC) combined with a low number of tests (per 1 LOC).



Fig. 10. Corrective activities and unit tests per 1 LOC for 61 projects (see also Levin and Yehudai (2017a))

Our analysis of XPrivacy did not reveal any unit tests in its codebase. Its README page on GitHub had a designated testing section which revealed that a separate application had been written for testing purposes. The test application's (GitHub) project was nowhere as popular as XPrivacy itself (more than 1.5K stars vs. less than 10 stars) implying it may not have been widely used by developers upon contributing code. It is possible that since the test application project was separate from the original application, it was not executed frequently (and automatically) enough, rendering it less effective in preventing defects. This may account for the high amount of corrective activity performed in this project. Omni-Notes, the second outlier project we inspected, had only 12 tests spread over 8 suites according to our analysis. Its README page on GitHub also had a designated section for testing which specified the build command developers should execute when contributing code. While the presence of a designated test section in its README page may indicate testing was quite important to the project's owner, the great amount of corrective activity performed in this project in this project may suggest it could have benefited from more unit tests. Gaining fine grained

visibility into anomalies (e.g., as indicated in Figure 10) will allow managers to identify potential issues by examining abnormal values even without knowing the root cause. Having identified potential issues, mangers can then shift focus towards investigation and resolution.

To conclude this section, while regression models do not provide means to ascertain causality, the negative correlation between corrective commits and tests (i.e., both methods and classes) is worth considering. Potentially, one could argue that projects with tests may only need little corrective activity due to the high quality of the codebase. The opposite direction, may imply that corrective activity may be required when the test count of a project is low, and the codebase's quality is poor. It is also possible, that test counts and corrective commits do not have a cause and effect relationship at all, in which case they just tend to happen together and are connected via a lurking variable. Either of these narratives requires further evidence before it can be reliably established, but to the very least, the empirically evident negative correlation between corrective activity and tests is yet another reminder of the relationship between automated testing and the nature and volume of the maintenance activities a project is likely to require in the future.

8.3 Future Applications

Identifying Anomalies In Development Processes. The manager of a large software project should aim to control and manage its maintenance activity profiles, i.e., the volume of commits made in each maintenance activity. Monitoring for unexpected spikes in maintenance activity profiles and investigating the reasons (root cause) behind them could assist managers and other stakeholders to plan ahead and identify areas that require additional resource allocation. For example, lower corrective profiles could imply that developers are neglecting bug fixing. Higher corrective profiles could imply an excessive bug count. Finding the root cause in cases of significant deviations from predicted values may reveal essential issues the removal of which can improve projects' health. Similarly, exceptionally well performing projects can be a good subject for case studies, so as to identify positive patterns.

Improving development team's composition. Building a successful software team is hardly a trivial task as it involves a delicate balance between technological and human aspects (Gorla and Lam, 2004; Guinan et al., 1998). We believe that by using commit classification it would be possible to build reliable developer maintenance activity profiles which could assist in composing balanced teams. We conjecture that composing a team that heavily favors a particular maintenance activity (e.g. adaptive) over the others could lead to an unbalanced development process and adversely affect the team's ability to meet typical requirements such as developing a sustainable number of product features, adhering to quality standards, and minimizing technical debt so as to facilitate future changes.

9 THREATS TO VALIDITY

Threats to Statistical Conclusion Validity are the degree to which conclusions about the relationship among variables based on the data are reasonable.

- <u>Classification Models</u>. Our commit classification results were based on manually classifying 1151 commits, over 100 commits from each of the studied 11 projects. The projects originated from various professional domains such as IDEs, programming languages, distributed database and storage platforms, and integration frameworks. Each compound model was trained using 5-time repeated 10-fold cross validation. In addition, our commit classifications evaluations demonstrated *p*-value below 0.01, supporting the statistical validity of the hypothesis accuracy > NIR with high confidence.
- Regression Models. Our dataset for the regression analysis consisted of 61 projects and over 240,000 commits. Both the model coefficients and the predictions were annotated with statistical significance levels to indicate the strength of the signal. Most of the coefficients were statistically significant (p-value < 0.01). To compare distributions we used the Wilcoxon-Mann-Whitney test and reported its high significance level (p-value < 0.01).

We assume commits are independent, however, it may be the case that commits performed by the same developer share common properties.

Threats to Construct Validity consider the relationship between theory and observation, in case the measured variables do not measure the actual factors.

- <u>Manual Commit Classification</u>. We took the following measures to mitigate manual classification related errors:
- Projects' issue tracking systems were used, and often provided additional information pertaining to commits.
- (2) Commits that did not lend themselves to classification due to lack of supporting information were removed from the dataset and replaced by other commits from the same repository (see Section 5).
- (3) A sample of 10% out of all manually labeled commits was independently classified by both authors. The observed agreement level was 94.5%, and the asymptotic 95% confidence interval for the agreement level was [90.3%, 98.7%] indicating that both authors agreed about the labels for the vast majority of cases.
- Fined-grained Source Code Change Extraction. ChangeDistiller and the VCS mining platform we have built on top of it are both software programs, and as such, are not immune to bugs which could result in inaccurate or incomplete data.
- <u>Test Maintenance Classification</u>. We used a widely practiced conventions and heuristics (Maven Surefire Plugin, 2017; Zaidman et al., 2011) for detecting JUnit test methods and test classes. However, the use of heuristics may lead to undetected test maintenance.

• Data Cleaning. Prior to devising regression models, we removed extreme data points using a technique suggested in (Hubert and Vandervieren, 2008). Despite the fact we removed only ~10% of the data, this process could have introduced bias into the dataset we operated on.

Threats to External Validity consider the generalization of our findings.

- Programming Language Bias. All analyzed commits were in the Java programming language since the tool we used to distill fine grained source code changes (ChangeDistiller) was Java oriented. It is possible that developers who use other programming languages, have different maintenance activity patterns which have not been explored in the scope of this work.
- Open Source Bias / GitHub. The repositories studied in this paper were all popular open source projects from GitHub, selected according to the criteria described in Section 4. It may be the case that developers' maintenance activity profiles are different in an open source environment when compared to other environments.
- <u>Popularity Bias</u>. We intentionally selected the popular, data rich repositories. This could limit our results to developers and repositories of high popularity, and potentially skew the perspective on characteristics found only in less popular repositories and their developers.
- <u>Limited Information Bias</u>. The entire dataset, both the training and the test datasets, contained only those commits that we were able to manually classify. At the stage of VCS inspection it can be essentially impossible to actually ascertain the maintenance activities of commits that do not provide enough information traces (comment, ticket id, etc.). The true maintenance activity for such commits may only be known to the developers who made them, and even they may no longer recall it soon after they have moved on to their next task.
- <u>Mixed Commits</u>. Recent studies (Kirinuki et al., 2014; Nguyen et al., 2013) report that commits may involve more than one type of maintenance activity, e.g. a commit that both fixes a bug, and adds a new feature. Our classification method does not currently account for such cases, but this is definitely an interesting direction to be considered for future work (see Section 10).
- <u>Activity Boundary</u>. In this work we assume a commit serves as a logical boundary of an activity. It may be the case, that developers perform test maintenance as part of activities that span multiple commits. Such work patterns were not considered in the scope of this work, but are definitely an interesting direction for future work in this area.

10 SUMMARY

We suggested a novel method for classifying commits into maintenance activities and used it to devise and evaluate a number of models that utilize fine-grained source code changes and the commit message for the purpose of cross-project commit classification into maintenance activities. These models were then evaluated and compared using the accuracy and Kappa metrics with different underlying classification algorithms. Our champion model showed a promising accuracy of 76% and Kappa of 63% when applied

48

on the test dataset which consisted of 172 commits originating from various projects. These results show an improvement of over 20 percentage points, and a relative improvement of over 40% when compared to previous results (Table 1). A comparison between the widely used classifier and our champion classifier can be found in Table 6 and Table 12, respectively. Our evaluation was based on studying 11 popular open source projects from various professional domains, from which we manually classified 1151 commits, ~100 from each of the studied projects. The suggested models were trained using repeated cross validation on 85% of the dataset, and the remaining 15% of the dataset were used as a test set.

We conclude that the answer to RQ 1. is that fine-grained source code changes can indeed be successfully used to devise high quality models for commit classification into maintenance activities.

The answer to RQ 2. is that models that utilize source code changes are capable of outperforming the reported accuracy of word frequency based models (Amor et al., 2006; Hindle et al., 2009) from ~60% to ~75%, even when classifying cross-project commits. In addition, we make the following observations based on our study:

- Using text cleaning and normalization, our word frequency based models were able to achieve an accuracy of 68-69% with Kappa of 51-53% for cross-project commits classification (see Table 8).
- Compound models employing both (commit message) word frequency analysis and source code change types for the task of cross-project commit classification were able to achieve up to 73% accuracy with Kappa 59% during the training stage, and up to 76% accuracy with Kappa of 63%, considered "Good" (Altman, 1990), for the test dataset.
- The RF algorithm outperformed the GBM and J48 in classifying cross-project commits (see Table 11 and Table 12).

To explore RQ. 3 we demonstrated two applications for our classification and repository harvesting methods, one in the spirit of the verbal view, and the other in the spirit of the nounal view.

- The Software Maintenance Activity Explorer, a tool that is aimed at providing an intuitive visualization of code maintenance activities over time. It provides users with both project wide, and developer centring views of maintenance activities over various periods of time. We then showed how the software maintenance activity explorer and our dataset can be used to identify homogeneous maintenance activity profiles, which we believe managers should be made aware of and act upon.
- Detecting software projects which may be lacking in tests and potentially require extensive corrective maintenance. The suggested application employs insights obtained from modeling the relationship between commit classification (into maintenance activities) and the number of test methods in a software project.

11 FUTURE WORK

We believe that our methods and results can be leveraged to further explore numerous directions in the field of software evolution and software analytics in particular. For example, it would be interesting to learn

whether our software maintenance activity explorer could appeal to practitioners working on open source and/or commercial projects. It would also be beneficial to learn what real-life tasks they believe this tool can help with, and/or what changes they would like to suggest to make it useful for their needs. In addition, it may be of particular interest to get feedback from developers who took part in the projects we analyzed as part of our publicly available version of the software maintenance activity explorer¹⁴.

Some commits may involve more than one type of maintenance activity, and some activities may span more than one commit. It would therefore be beneficial to explore whether extended activities and mixed commits lend themselves to automatic and accurate classification.

The availability of an accurate classification model may make it possible to automatically classify an unprecedentedly large number of projects and commit activities. This, in turn, could shed new light on the distribution of maintenance activities in software projects (Lientz et al., 1978; Schach et al., 2003), a subject the research community is yet to agree upon.

¹⁴Available at https://soft-evo.shinyapps.io/maintenance-activities.

REFERENCES

D. G. Altman. Practical statistics for medical research. CRC press, 1990.

- J. J. Amor, G. Robles, J. M. Gonzalez-Barahona, and A. Navarro. Discriminating development activities in versioning systems: A case study. In *Proceedings PROMISE*. Citeseer, 2006.
- L. B. A. C. Andy Liaw, Matthew Wiener. randomforest: Breiman and cutler's random forests for classification and regression. https://CRAN.R-project.org/package=randomForest, 2015. [Online; accessed Nov-2016].

Apache Spark, 2016. Lightning-fast cluster computing. http://spark.apache.org/, 2014. [Online; accessed 11-April-2016].

- Atlassian. The #1 software development tool used by agile teams. https://www.atlassian.com/software/jira, 2014. [Online; accessed 20-Mar-2017].
- L. Belady and M. Lehman. Programming System Dynamics Or Thmeta-dynamics of Systems in Maintenance and Growth. IBM Thomas J. Watson Research Center, 1971.
- L. A. Belady and M. M. Lehman. A model of large program development. IBM Systems journal, 15(3):225-252, 1976.
- M. Beller, G. Gousios, A. Panichella, and A. Zaidman. When, how, and why developers (do not) test in their ides. In Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, pages 179–190. ACM, 2015.
- L. Breiman. Random forests. Machine learning, 45(1):5-32, 2001.
- R. P. Buse and T. Zimmermann. Analytics for software development. In Proceedings of the FSE/SDP workshop on Future of software engineering research, pages 77–80. ACM, 2010.
- A. E. Camargo Cruz and K. Ochimizu. Towards logistic regression models for predicting fault-prone code across software projects. In Proceedings of the 2009 3rd International Symposium on Empirical Software Engineering and Measurement, pages 460–463. IEEE Computer Society, 2009.
- R. Caruana and A. Niculescu-Mizil. An empirical comparison of supervised learning algorithms. In Proceedings of the 23rd international conference on Machine learning, pages 161–168. ACM, 2006.
- R. Caruana, N. Karampatziakis, and A. Yessenalina. An empirical evaluation of supervised learning in high dimensions. In Proceedings of the 25th international conference on Machine learning, pages 96–103. ACM, 2008.
- W. S. Cleveland, R. McGill, et al. Graphical perception and graphical methods for analyzing scientific data. Science, 229(4716):828–833, 1985.
- F. X. Diebold. On the origin (s) and development of the term Big Data. PIER Working Paper, 2012.
- J. Falleri and F. Morandat. Gumtree a neat code differencing tool. https://github.com/GumTreeDiff/gumtree, 2014. [Online; accessed 11-March-2017].
- J. Falleri, F. Morandat, X. Blanc, M. Martinez, and M. Monperrus. Fine-grained and accurate source code differencing. In ACM/IEEE International Conference on Automated Software Engineering, ASE '14, Vasteras, Sweden - September 15 - 19, 2014, pages 313–324, 2014. doi: 10.1145/2642937.2642982. URL http://doi.acm.org/10.1145/2642937.2642982.
- M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim. Do we need hundreds of classifiers to solve real world classification problems. J. Mach. Learn. Res, 15(1):3133–3181, 2014.
- S. Few. Now you see it: simple visualization techniques for quantitative analysis. Analytics Press, 2009.
- M. Fischer, M. Pinzger, and H. Gall. Populating a release history database from version control and bug tracking systems. In Software Maintenance, 2003. ICSM 2003. Proceedings. International Conference on, pages 23–32. IEEE, 2003.
- B. Fluri and H. C. Gall. Classifying change types for qualifying change couplings. In Program Comprehension, 2006. ICPC 2006. 14th IEEE International Conference on, pages 35–45. IEEE, 2006.
- B. Fluri, M. Wursch, M. PInzger, and H. C. Gall. Change distilling: Tree differencing for fine-grained source code change extraction. Software Engineering, IEEE Transactions on, 33(11):725–743, 2007.
- B. Fluri, E. Giger, and H. C. Gall. Discovering patterns of change types. In Automated Software Engineering, 2008. ASE 2008. 23rd IEEE/ACM International Conference on, pages 463–466. IEEE, 2008.

- B. Fluri, M. Würsch, E. Giger, and H. C. Gall. Analyzing the co-evolution of comments and source code. Software Quality Journal, 17 (4):367–394, 2009.
- E. Frank, M. Hall, G. Holmes, R. Kirkby, B. Pfahringer, I. H. Witten, and L. Trigg. Weka. In Data Mining and Knowledge Discovery Handbook, pages 1305–1314. Springer, 2005.
- J. H. Friedman. Greedy function approximation: a gradient boosting machine. Annals of statistics, pages 1189-1232, 2001.
- H. C. Gall, B. Fluri, and M. Pinzger. Change analysis with evolizer and changedistiller. *IEEE Software*, 26(1):26, 2009.
- C. Ghezzi, M. Jazayeri, and D. Mandrioli. Fundamentals of software engineering. Prentice Hall PTR, 2002.
- E. Giger, M. Pinzger, and H. C. Gall. Comparing fine-grained source code changes and code churn for bug prediction. In Proceedings of the 8th Working Conference on Mining Software Repositories, pages 83–92. ACM, 2011.
- E. Giger, M. Pinzger, and H. C. Gall. Can we predict types of code changes? an empirical analysis. In *Mining Software Repositories* (*MSR*), 2012 9th IEEE Working Conference on, pages 217–226. IEEE, 2012.
- GitHub Inc. New year, new company. https://blog.github.com/2010-01-22-new-year-new-company/, 2010. [Online; accessed 18-April-2016].
- GitHub Inc. A whole new code search. https://blog.github.com/2013-01-23-a-whole-new-code-search/, 2013. [Online; accessed 11-April-2016].
- GitHub Inc. About the search api. https://developer.github.com/v3/search/, 2015. [Online; accessed 11-April-2016].
- GitHub Inc. Github the largest open source community in the world. https://github.com/about, 2018. [Online; accessed 18-October-2018].
- N. Gorla and Y. W. Lam. Who should work with whom?: building effective software project teams. *Communications of the ACM*, 47(6): 79–82, 2004.
- P. J. Guinan, J. G. Cooprider, and S. Faraj. Enabling software development team performance during requirements definition: A behavioral versus technical approach. *Information Systems Research*, 9(2):101–125, 1998.
- P. Hamill. Unit Test Frameworks: Tools for High-Quality Software Development. " O'Reilly Media, Inc.", 2004.
- A. E. Hassan. Empirical evaluations in software engineering research: A personal perspective. https://www.slideshare.net/SAIL_QU/ empirical-evaluations-in-software-engineering-research-a-personal-perspective, 2017. [Online; accessed 11-February-2018].
- K. Herzig, S. Just, and A. Zeller. It's not a bug, it's a feature: how misclassification impacts bug prediction. In Proceedings of the 2013 International Conference on Software Engineering, pages 392–401. IEEE Press, 2013.
- A. Hindle, D. M. German, M. W. Godfrey, and R. C. Holt. Automatic classication of large changes into maintenance categories. In Program Comprehension, 2009. ICPC'09. IEEE 17th International Conference on, pages 30–39. IEEE, 2009.
- T. K. Ho. The random subspace method for constructing decision forests. *IEEE transactions on pattern analysis and machine intelligence*, 20(8):832–844, 1998.
- K. Hornik, C. Buchta, and A. Zeileis. Open-source machine learning: R meets Weka. Computational Statistics, 24(2):225–232, 2009. doi: 10.1007/s00180-008-0119-7.
- M. Hubert and E. Vandervieren. An adjusted boxplot for skewed distributions. *Computational statistics & data analysis*, 52(12): 5186–5201, 2008.
- H. Kirinuki, Y. Higo, K. Hotta, and S. Kusumoto. Hey! are you committing tangled changes? In *Proceedings of the 22nd International Conference on Program Comprehension*, pages 262–265. ACM, 2014.
- M. Kuhn. The caret package. http://topepo.github.io/caret/index.html, 2017. [Online; accessed Nov-2016].
- M. Kuhn. caret v6.0-80, createdatapartition. https://www.rdocumentation.org/packages/caret/versions/6.0-80/topics/ createDataPartition, 2018. [Online; accessed 29-Jul-2018].
- M. Kuhn, A. W. C. K. A. E. T. C. Z. M. B. K. t. R. C. T. M. B. R. L. A. Z. L. S. Y. T. C. C. Jed Wing, Steve Weston, and T. Hunt. caret: Classification and regression training. https://CRAN.R-project.org/package=caret, 2017. [Online; accessed Nov-2016].
- M. M. Lehman. The programming process. internal IBM report, 1969.
- M. M. Lehman. Programs, cities, studentsj-limits to growth? In Programming Methodology, pages 42-69. Springer, 1978.

Towards Software Analytics: Modeling Maintenance Activities

- M. M. Lehman and J. F. Ramil. Software evolution-background, theory, practice. Information Processing Letters, 88(1):33-44, 2003.
- M. M. Lehman, J. F. Ramil, and G. Kahen. Evolution as a noun and evolution as a verb. In SOCE 2000 Workshop on Software and Organisation Co-evolution, volume 9, page 31, 2000.
- S. Levin. Software maintenance activities explorer. https://soft-evo.shinyapps.io/maintenance-activities, 2017. [Online; accessed 11-February-2018].
- S. Levin. Software maintenance explorer. https://github.com/staslev/software-maintenance-explorer, 2018. [Online; accessed 11-November-2018].
- S. Levin and A. Yehudai. Using temporal and semantic developer-level information to predict maintenance activity profiles. In *Proc. ICSME*, pages 463–468. IEEE, 2016.
- S. Levin and A. Yehudai. The co-evolution of test maintenance and code maintenance through the lens of fine-grained semantic changes. In 2017 IEEE International Conference on Software Maintenance and Evolution, ICSME 2017, Shanghai, China, September 20-22, 2017, pages 35–46, 2017a. doi: 10.1109/ICSME.2017.9. URL https://doi.org/10.1109/ICSME.2017.9.
- S. Levin and A. Yehudai. Boosting automatic commit classification into maintenance activities by utilizing source code changes. In Proceedings of the 13th International Conference on Predictive Models and Data Analytics in Software Engineering, PROMISE, pages 97–106, New York, NY, USA, 2017b. ACM. ISBN 978-1-4503-5305-2. doi: 10.1145/3127005.3127016. URL http://doi.acm.org/ 10.1145/3127005.3127016.
- S. Levin and A. Yehudai. 1151 commits with software maintenance activity labels (corrective,perfective,adaptive), July 2017c. URL https://doi.org/10.5281/zenodo.835534.
- S. Levin and A. Yehudai. Statistics for the studied 61 open source projects. https://github.com/staslev/paper-resources/blob/master/ The-Co-Evolution-of-Test-Maintenance-and-Code-Maintenance-through-the-lens-of-Fine-Grained-Semantic-Changes/studiedrepos.md, 2017d. [Online; accessed 11-February-2018].
- A. Liaw and M. Wiener. Classification and regression by randomforest. *R News*, 2(3):18-22, 2002. URL http://CRAN.R-project.org/doc/ Rnews/.
- B. P. Lientz, E. B. Swanson, and G. E. Tompkins. Characteristics of application software maintenance. *Communications of the ACM*, 21 (6):466–471, 1978.
- M. Martinez, L. Duchien, and M. Monperrus. Automatically extracting instances of code change patterns with ast analysis. arXiv preprint arXiv:1309.3730, 2013.
- Maven Surefire Plugin. Inclusions and exclusions of tests. http://maven.apache.org/surefire/maven-surefire-plugin/examples/ inclusion-exclusion.html, 2017. [Online; accessed Jan-2017].
- P. McCullagh and J. A. Nelder. Generalized linear models, volume 37. CRC press, 1989.
- T. Menzies and T. Zimmermann. Software analytics: so what? IEEE Software, (4):31-37, 2013.
- W. Meyers. Interview with wilma osborne. IEEE Software, 5(3):104-105, 1988.
- A. Mockus and L. G. Votta. Identifying reasons for software changes using historic databases. In Software Maintenance, 2000. Proceedings. International Conference on, pages 120–130. IEEE, 2000.
- H. A. Nguyen, A. T. Nguyen, and T. N. Nguyen. Filtering noise in mixed-purpose fixing commits to improve defect prediction and localization. In Software Reliability Engineering (ISSRE), 2013 IEEE 24th International Symposium on, pages 138–147. IEEE, 2013.
- J. R. Quinlan. C4. 5: programs for machine learning. Elsevier, 2014.
- R Development Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, 2008. URL http://www.R-project.org. ISBN 3-900051-07-0.
- G. Ridgeway. Generalized boosted models: A guide to the gbm package. Update, 1(1):2007, 2007.
- G. Ridgeway and Others. R gbm package. https://CRAN.R-project.org/package=gbm, 2015. [Online; accessed Nov-2016].
- M. Saar-Tsechansky and F. Provost. Handling missing values when applying classification models. *Journal of machine learning research*, 8(Jul):1623–1657, 2007.
- Scala, 2015. The Scala programming language. https://www.scala-lang.org/, 2015. [Online; accessed 11-February-2018].

- S. R. Schach, B. Jin, L. Yu, G. Z. Heller, and J. Offutt. Determining the distribution of maintenance categories: Survey versus measurement. *Empirical Software Engineering*, 8(4):351–365, 2003.
- S.E.A.L UZH. The changedistiller repository. https://bitbucket.org/sealuzh/tools-changedistiller/wiki/Home, 2011. [Online; accessed 26-March-2017].
- S.E.A.L UZH. The changedistiller api. https://bitbucket.org/sealuzh/tools-changedistiller/src/ feee5be3724a3eabfb7c415554cb26f2258a65f4/src/main/java/ch/uzh/ifi/seal/changedistiller/distilling/
- FileDistiller.java?at=master&fileviewer=file-view-default#FileDistiller.java-75, 2014. [Online; accessed 26-March-2017].
- E. Shihab. An exploration of challenges limiting pragmatic software defect prediction. PhD thesis, Citeseer, 2012.
- J. Śliwerski, T. Zimmermann, and A. Zeller. When do changes induce fixes? In ACM sigsoft software engineering notes, volume 30, pages 1–5. ACM, 2005.

StackOverflow. Developer survey results 2017. https://insights.stackoverflow.com/survey/2017, 2017. [Online; accessed 1-Nov-2017].
StackOverflow. Developer survey results 2018. https://insights.stackoverflow.com/survey/2018, 2018. [Online; accessed 26-March-2018].

- E. B. Swanson. The dimensions of maintenance. In *Proceedings of the 2nd international conference on Software engineering*, pages 492–497. IEEE Computer Society Press, 1976.
- L. Torvalds. Tech talk: Linus torvalds on git. https://www.youtube.com/watch?v=4XpnKHJAok8&, 2007. [Online; accessed 11-Mar-2018].
- W. N. Venables and B. D. Ripley. Modern applied statistics with S-PLUS. Springer Science & Business Media, 2013.
- I. H. Witten and E. Frank. Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann, San Francisco, 2nd edition, 2005.
- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association, 2012.
- M. Zaharia, R. S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M. J. Franklin, et al. Apache spark: a unified engine for big data processing. *Communications of the ACM*, 59(11):56–65, 2016.
- A. Zaidman, B. Van Rompaey, A. van Deursen, and S. Demeyer. Studying the co-evolution of production and test code in open source and industrial developer test processes through repository mining. *Empirical Software Engineering*, 16(3):325–364, 2011.