

An automatically created novel bug dataset and its validation in bug prediction

Rudolf Ferenc^{a,*}, Péter Gyimesi^a, Gábor Gyimesi^a, Zoltán Tóth^a, Tibor Gyimóthy^{a,b}

^a Department of Software Engineering, University of Szeged, Hungary

^b MTA-SZTE Research Group on Artificial Intelligence, Szeged, Hungary

ARTICLE INFO

Article history:

Received 26 July 2019

Received in revised form 30 April 2020

Accepted 4 June 2020

Available online 7 June 2020

Keywords:

Bug dataset

Bug prediction

Static code analysis

Code metrics

Machine learning

GitHub

ABSTRACT

Bugs are inescapable during software development due to frequent code changes, tight deadlines, etc.; therefore, it is important to have tools to find these errors. One way of performing bug identification is to analyze the characteristics of buggy source code elements from the past and predict the present ones based on the same characteristics, using e.g. machine learning models. To support model building tasks, code elements and their characteristics are collected in so-called bug datasets which serve as the input for learning.

We present the *BugHunter Dataset*: a novel kind of automatically constructed and freely available bug dataset containing code elements (files, classes, methods) with a wide set of code metrics and bug information. Other available bug datasets follow the traditional approach of gathering the characteristics of all source code elements (buggy and non-buggy) at only one or more pre-selected release versions of the code. Our approach, on the other hand, captures the buggy and the fixed states of the same source code elements from the narrowest timeframe we can identify for a bug's presence, regardless of release versions. To show the usefulness of the new dataset, we built and evaluated bug prediction models and achieved F-measure values over 0.74.

© 2020 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The characterization of buggy source code elements is a popular research area these days. Programmers tend to make mistakes despite the assistance provided by different integrated development environments, and errors may also occur due to frequent changes in the code and inappropriate specifications; therefore, it is important to get more and/or better tools to help the automatic detection of errors (Johnson et al., 2013). For automatic recognition of unknown faulty code elements, it is a prerequisite to characterize the already known ones.

During software development, programmers use a wide variety of tools, including bug tracking, task management, and version control systems. There are numerous commercial and open-source software systems available for these purposes. Furthermore, different web services are built to meet these needs. The most popular ones like SourceForge, Bitbucket, and GitHub fulfill the above mentioned functionalities. They usually provide several services, such as source code hosting and user management. Their APIs make it possible to retrieve various kinds of data,

e.g., they provide support for the examination of the behavior or the co-operation of users or even for analyzing the source code itself. Since most of these services include bug tracking, it raises the idea to use this information in the characterization of buggy source code parts (Zhou et al., 2012). To do so, the bug reports managed by these source code hosting providers must be connected to the appropriate source code parts (Wu et al., 2011). A common practice in version control systems is to describe the changes in a comment belonging to a commit (log message) and often provide the identifier of the associated bug report which the commit is supposed to fix (Kalliamvakou et al., 2014). This can be used to identify the faulty versions of the source code (Dallmeier and Zimmermann, 2007a,b). Processing diff files can help us obtain the code sections affected by the bug (Toth et al., 2013). We can use source code metrics (Couto et al., 2012), for which we only need a tool that is able to produce them.

To build a dataset containing useful buggy code element characterization information, we chose GitHub, since it hosts several regularly maintained projects and also a well defined API that makes it possible to implement an automatic data retrieval tool. We selected 15 Java projects as our subject systems, which differ in many ways from each other (size, domain, number of bugs reported) to cover a wide and general set of systems.

* Correspondence to: H-6720 Szeged, Dugonics tér 13, Hungary.

E-mail addresses: ferenc@inf.u-szeged.hu (R. Ferenc), pgyimesi@inf.u-szeged.hu (P. Gyimesi), ggimesi@inf.u-szeged.hu (G. Gyimesi), zizo@inf.u-szeged.hu (Z. Tóth), gyimothy@inf.u-szeged.hu (T. Gyimóthy).

Previously published datasets follow a traditional concept to create a dataset that serves as a benchmark for testing bug prediction techniques (D'Ambrosi et al., 2010; Jureczko and Madeyski, 2010). These datasets include all code elements – both buggy and non-buggy – from one or more versions of the analyzed system. In this work, we created a new approach that collects before-fix and after-fix snapshots of source code elements (and their characteristics) that were affected by bugs and does not consider those code elements that were not touched by bugs. This kind of dataset helps to capture the changes in software product metrics when a bug is being fixed. Thus, we can learn from the differences in source code metrics between faulty and non-faulty code elements. As far as we know, there exists no other bug dataset yet that tries to capture this before-fix and after-fix state.

Our new dataset is called *BugHunter Dataset* and it is freely available (see Section 5.8). It can serve as a new kind of benchmark for testing different bug prediction methods since we included a wide range of source code metrics to describe the previously detected bugs in the chosen systems. We took all reported bugs stored in the bug tracking system into consideration. We used the usual methodology of connecting commits to bugs by analyzing the log messages and by looking for clues that would unambiguously identify the bug that was the intended target of the corresponding fixing commit(s). Commit diffs helped us detect which source code elements were modified by a given change set, thus the code elements which had to be modified in order to fix the bug.

The first version of this work was published in our earlier conference paper (Gyimesi et al., 2015). Since then, we have extended the list of subject projects (one project was replaced because it contained template files with .java extension which undermined the source code analysis) and we have also expanded the time interval of the analysis of the projects' history to cover three additional years (from 2014 to 2017). Furthermore, we have refined our technique and we have incorporated method level bug information into our dataset as well. The dataset we present in this paper has not been published before.

Here we also performed experiments to check whether our novel dataset is suitable for bug prediction purposes. During this investigation, we collected bug characterization metrics at three source code levels (file, class, method). After the dataset was constructed, we used different machine learning algorithms to analyze the usefulness of the dataset.

We also performed a novel kind of experiment in which we assessed whether the method level metrics are better predictors when projected to class level than the class level metrics themselves.

An important aspect to investigate is how the bug prediction models built from the novel dataset compare to the ones which used the traditional datasets as corpus. However, this comparison is hard in its nature due to the variability in multiple factors. One major problem is the difference in the corpus itself. The list of the included projects vary from dataset to dataset. In our previous work, we constructed a traditional dataset, the *GitHub Bug Dataset* (Tóth et al., 2016), which consists of the same 15 projects we also included in our novel bug dataset. This gives an opportunity to assess if there is any difference in the bug prediction capabilities of these two datasets.

To emphasize the research artifact contribution and the research questions, we highlighted them in the following box.

Research artifact: A freely available novel dataset containing source code metrics of buggy Java source code elements (file, class, method) before and after bug fixes were applied to them.

RQ1: Is the constructed dataset usable for bug prediction purposes?

RQ2: Are the method level metrics projected to class level better predictors than the class level metrics themselves?

RQ3: Is the *BugHunter Dataset* more powerful and expressive than the *GitHub Bug Dataset*, which is constructed with the traditional approach?

The rest of the paper is structured as follows. In Section 2 we discuss related work. We present some statistics about GitHub and the projects that we have chosen for this work in Section 3. Next, we introduce the metrics used for bug characterization in Section 4. We describe our approach for generating and validating the dataset in detail in Section 5. Then, we evaluate it by applying different machine learning algorithms in Section 6, where we also address our research questions. In Section 7, we present the threats to validity. Finally, we conclude and describe some future work directions in Section 8.

2. Related work

2.1. Bug localization and source code management

Many approaches have been presented dealing with bug characterization and localization (Saha et al., 2013; Wang and Lo, 2014; Davies et al., 2012). Zhou et al. published a study describing BugLocator (Zhou et al., 2012), a tool that detects the relevant source code files that need to be changed in order to fix a bug. BugLocator uses textual similarities (between initial bug report and the source code) to rank potentially fault-prone files. Prior information about former bug reports is stored in a bug database. Ranking is based on the assumption that descriptions with high similarities imply that the related files are highly similar too. A similar ranking is done by Rebug-Detector (Wang et al., 2010), a tool made by Wang et al. for detecting related bugs from source code using bug information. The tool focuses on overridden and overloaded method similarities. In our study, we constructed a dataset that stores information about formerly buggy code elements that are now fixed, thus the same method could be applied by using source code metrics for detecting similar source code elements that are possibly fault-prone.

ReLink (Wu et al., 2011) was developed to explore missing links between code changes committed in version control systems and the fixed bugs. This tool could be helpful for software engineering research based on linkage data, such as software defect prediction. ReLink mines and analyzes information like bug reporter, description, comments, and date from a bug tracking database and then tries to pair the bug with the appropriate source code files based on the set of source code information extracted from a version control system. Most of the studies dealing with this kind of linkage data use the SZZ algorithm, which has been improved over the years (Kim et al., 2006; Williams and Spacco, 2008). This approach uses file level textual features to extract extra information between bugs and the source code itself. We characterized the set of bugs in a given system at file, class, and method levels by assigning different source code metrics to the source code elements (methods, classes, and files) affected by bugs. Other approaches try to link information retrieved from only version control and bug tracking systems (Fischer et al., 2003; Mockus and Votta, 2000).

Kalliamvakou et al. mined GitHub repositories to investigate their characteristics and their qualities (Kalliamvakou et al., 2014). They presented a detailed study discussing different project characteristics, such as (in)activity, while also involving

further research questions – e.g., whether a project is standalone or part of a more massive system. Results showed that the extracted data can serve as a good input for various investigations, however one must use them with caution and always verify the usefulness and reliability of the data. It is a good practice to choose projects with many developers and commits, but it should always be kept in mind that the most important point is to choose projects that fit well for one's own purpose. In our case, we have tried to create a dataset that is large, reliable (through some manual validation) and general enough for testing different bug prediction techniques (Catal and Diri, 2009; Porter et al., 1990; Ostrand et al., 2005; Ma et al., 2006; Zhou and Leung, 2006), while still being created in an automatic way.

Mining software repositories can be a harsh task when an automatic mechanism is used to construct a large set of data based on the information gathered from a distributed software repository. As we used GitHub to address our research questions, we paid extra attention to prevent and avoid pitfalls. Bird et al. (2009) presented a study on distributed version control systems – focusing mainly on Git – that examined their usage and the available set of data (such as whether the commits are removable, modifiable, movable). The main purpose of the paper was to draw attention to pitfalls and help researchers to avoid them during the processing and analysis of a mined information set.

Many research papers showed that using a bug tracking system improves the quality of the software system under development. Bangcharoensap et al. introduced a method to locate the buggy files in a software system very quickly using the bug reports managed by the bug tracking system (Bangcharoensap et al., 2012). The presented method contains three different approaches to rank the fault-prone files, namely: (a) Text mining, which ranks files based on the textual similarity between a bug report and the source code itself, (b) Code mining, which ranks files based on prediction of the potential buggy module using source code product metrics, and (c) Change history, which ranks files based on prediction of the fault-prone module using change process metrics. They used the gathered project data collected on the Eclipse platform to investigate the efficiency of the proposed approaches and showed that they are indeed suitable to locate buggy files. Furthermore, bug reports with a short description and many specific words greatly increase the effectiveness of finding the “weak points” of the system.

Similarly to our study, Ostrand et al. investigated fault prediction by using source code metrics. However, only file level was considered as the finest granularity unit (Ostrand et al., 2007), while we have built a toolchain to also support class and method levels.

In addition to the above presented methods, a significant change in source code metrics can also indicate that the relevant source code files contain potential bugs (Gyimesi et al., 2005). Couto et al. presented a paper that shows the possible relationship between changed source code metrics (used as predictors) and bugs as well (Couto et al., 2012). They described an experiment to discover more robust evidences towards causality between software metrics and the occurrence of bugs. Although our method does not include this specific information, we still aim to show that considering methods as basic elements and including them in a dataset is also a way for building a working corpus for bug prediction techniques.

2.2. Public datasets

The previously mentioned approaches use self-made datasets for their own purposes, as illustrated in the work of Kalliamvakou et al. too (Kalliamvakou et al., 2014). Bug prediction techniques and approaches can be presented and compared in different

ways; however, there are some basic points that can serve as common components (Li et al., 2018). One common element can be a dataset used for the evaluation of the various approaches. PROMISE (Shirabad and Menzies, 2005) is a repository of datasets out of which several ones contain bugs gathered from open-source and also from closed-source industrial software systems. Amongst others it includes the NASA MDP dataset, which was used in many research studies and also criticized for containing erroneous data (Shepperd et al., 2013; Petrić et al., 2016). The PROMISE repository also contains an extensively referenced dataset created by Jureczko and Madeyski (2010), which provides object-oriented metrics as well as bug information for the source code elements (classes). This latter one includes open-source projects such as Apache Ant, Apache Camel, JEdit, Apache Lucene, forming a dataset containing 48 releases of 15 projects. The main purpose of these datasets is to support prediction methods and summarize bugs and their characterizations extracted from various projects. Many research papers used datasets from the PROMISE repository as an input for their investigations.

A similar dataset for bug prediction came to be commonly known as the *Bug prediction dataset*¹ (D'Ambros et al., 2010). The reason for creating this dataset was mainly inspired by the idea of measuring the performance of the different prediction models and also comparing them to each other. This dataset handles the bugs and the relevant source code parts at class level, i.e., the bugs are assigned to classes. As we have already mentioned, we do not only focus on file and class levels, but on method-level elements as well.

Zimmermann et al. (2007) used Eclipse as the input for a study dealing with defect prediction. They investigated whether the complexity metrics have the power to detect fault prone points in the system at package and file level. During the study, they constructed a public dataset² that is still available. It contains different source code metrics and a subset of the files/packages is marked as “buggy” if it contained any bugs in the interval between two releases.

A recent study showed that code smells also play a significant role in bug prediction (Hall et al., 2014) but the constructed dataset is not public. In our dataset, we also include code smell metrics to enhance its usefulness.

iBUGS (Dallmeier and Zimmermann, 2007b) provides a complex environment for testing different automatic defect localization methods. Information describing the bugs comes from both version control systems and from bug tracking systems. iBUGS used the following three open-source projects to extract the bugs from (the numbers of extracted bugs are in parentheses): AspectJ – an extension for the Java programming language to support aspect oriented programming (223); Rhino – a JavaScript interpreter written in Java (32); and Joda-Time – a quality replacement (extension) for the Java date and time classes (8). The authors attempted to generate the iBUGS dataset in an automatic way and they compared the generated set to the manually validated set of bugs (Dallmeier and Zimmermann, 2007a). iBUGS is a framework aimed more towards bug localization and not a standalone dataset containing source code elements and their characterizations (i.e., metrics).

The *Bugcatchers* (Hall et al., 2014) dataset is created by Hall et al. which is not only a bug dataset, but also contains bad smells detected in the subject systems. The selected three systems are Eclipse JDT Core, ArgoUML, and Apache Commons. The dataset is built and evaluated at file level.

¹ <http://bug.inf.usi.ch/>.

² <https://www.st.cs.uni-saarland.de/softevo/bug-data/eclipse/>.

Table 1
Comparison of the two types of datasets.

Feature	Traditional	Novel
Included time interval	Usually 6 months	Entire project history
Included source code elements	All the elements from a single version	Only the modified elements right before and after bug-fixes
Assumptions	Source code elements that are not included in any bug-fix are non-faulty	No assumptions needed
Uncertainty	The source code elements are faulty in the latest release version before the bug-fix and non-faulty after the fix	The source code elements are faulty right before the bug-fix and fixed afterwards

The *ELFF* dataset (Shippey et al., 2016) is a recent dataset proposed by Shippey et al. They experienced that only a few method level datasets exist, thus they created a dataset whose entries are methods. Additionally, they also made class level datasets publicly available. They used Boa (Dyer et al., 2013) to mine SourceForge repositories and collect as many candidates as they can, selecting 23 projects out of 50,000 that fulfilled their criteria (number of closed bugs, bugs are referenced from commits, etc.). They only kept projects with SVN version control systems which narrows down their candidate set. They used the classic and well-defined SZZ algorithm (Śliwerski et al., 2005) to find linkage between bugs and the corresponding source code elements.

The *Had-oops!* dataset (Harman et al., 2014) is constructed by a new approach presented by Harman et al. They analyzed 8 consecutive Hadoop versions and investigated the impact of chronology on fault prediction performance. They used Support Vector Machines (SVMs) with the Genetic Algorithm (for configuration) to build prediction models at class level. For a given version, they constructed a prediction model from all the previous versions and a model from only the current version and compared which one performed better. Results are not straightforward since they found early versions preferable in many cases as opposed to models built on recent versions. Moreover, using all versions is not always better than using only the current version to build a model from.

The *Mutation-aware fault prediction dataset* is a result of an experiment carried out by Bowes et al. on using mutation metrics as independent variables for fault prediction (Bowes et al., 2016). They used 3 software systems from which 2 projects (Eclipse and Apache) were open-source and one was closed. They used the popular PITest (or simply, PIT (Coles, 2018)) to obtain the set of mutation metrics that were included in the final dataset. Besides the mutation metrics, some static source code metrics (calculated by JHawk (2018)) were also included in the dataset for comparison purposes. This dataset is also built at class level.

Table 2
Comparison of the datasets.

Project	Level of bugs	Bug characteristics	# of projects
<i>NASA MDP Dataset</i>	Class	Static source code metrics	11
<i>Jureczko Dataset</i>	Class	Static source code metrics	15
<i>Bug prediction dataset</i>	Class	Static source code metrics, process metrics	5
<i>Eclipse dataset</i>	File, package	Complexity metrics	1
<i>iBUGS</i>	N/A	Bug-fix size properties, AST fingerprints	3
<i>Bugcatchers</i>	File	Code smells	3
<i>ELFF</i>	Class, method	Static source code metrics	23
<i>Had-oops!</i>	Class	Static source code metrics	1
<i>Mutation-aware fault prediction dataset</i>	Class	Static source code metrics, mutation metrics	3
<i>GitHub Bug Dataset</i>	File, class	Static source code metrics, code duplication metrics, code smell metrics	15
<i>Novel dataset</i>	File, class, method	Static source code metrics, code duplication metrics, code smell metrics	15

The *GitHub Bug Dataset* is a recent dataset that includes class and file level static source code metrics (Tóth et al., 2016) for 15 Java systems gathered from GitHub. Besides size, documentation, object-oriented, and complexity metrics, the dataset also contains code duplication and coding rule violation metrics. This dataset is our previous work that was still constructed in the “traditional” way. In Table 2, we compare the main characteristics of the mentioned datasets.

Our goal was to pick the strong aspects of all the previous datasets and put them together, as its positive effects are described by Li et al. (2019). Although the discussed works successfully made use of their datasets, an extended dataset can serve as a good basis for further investigations. Our dataset includes various projects from GitHub and includes numerous static source code metrics and stores a large number of entries in fine granularity (file, class, and method level as well). Furthermore, we also experimented with chronology, although in a different way compared to Harman et al. (2014). The differences between the traditional datasets and the proposed novel dataset are summarized in Table 1. See Section 5.2 for details about the process of selecting the bug related data for the novel dataset. The detailed comparison can be found in Section 6.3.

3. Data source

To address our research objectives, this section briefly introduces the version control system used (Git), its corresponding source code hosting service (GitHub), and their main functionalities that are closely related to the creation of linkage data between bugs and source code elements. Afterwards, we enumerate the chosen projects and give some further insight on the reasons why we chose them as our subject systems.

3.1. GitHub

GitHub is one of today’s most popular source code hosting services. It is used by several major open-source teams for managing their projects like Node.js, Ruby on Rails, Spring Framework, Zend Framework, and Jenkins, among others. GitHub offers public and private Git repositories for its users, with some collaborative services, e.g., built-in bug and issue tracking systems.

Bug reporting is supported by the fact that any GitHub user can add an issue, and collaborators can even label these issues for further categorization. The system provides some basic labels, such as “bug”, “duplicate”, and “enhancement”, but these tags can be customized if required. In an optimal case, the collaborators review these reports and label them with the proper labels, for instance, the bug reports with the “bug” label. The most important feature of bug tracking is that we can refer to an issue from the log message of a commit by using the unique identifier of the issue, thereby identifying a connection between the source code

Table 3

The number of repositories created between 01-01-2013 and 07-09-2017 for the top 10 languages.

Language	Number of repositories
JavaScript	2,019,215
Java	1,465,168
Ruby	1,379,225
Python	1,014,760
PHP	983,479
C	737,314
C++	619,914
CSS	568,493
C#	282,092
Shell	263,350

and the reported bug. GitHub has an API³ that can be used for managing repositories from other systems, or query information about them. This information includes events, feeds, notifications, gists, issues, commits, statistics, and user data.

With the GitHub Archive⁴ project that also uses this API, we can get up-to-date statistics about the public repositories. For instance, Table 3 presents the number of repositories created between 1 January 2013 and 7 September 2017, grouped by the main programming languages they use (only the top 10 languages are shown). Our approach uses Java repositories (the second most used platform on GitHub) to gather a proper information base for constructing a bug dataset.

Although extracting basic information from GitHub is easy, some version control features are hard to deal with, especially during the linking process when we try to match source code elements to bugs. For example, Git provides a powerful branching mechanism by supporting the creation, deletion, and selection of branches. In our case, we have to handle different branches because a fixing commit most often occurs on other – so called “topic” – branches and not on the master branch. Fortunately, the projects we analyzed often solved this problem by merging. During the merge, isomorphic commits are generated and placed on the master branch, thus all the desired analysis can be done by taking only the master branch with a given version as input. Another example is forking a repository, which is used worldwide. In our experiment, we do not handle forks, since it would have encumbered the above mentioned linking process and we would not gain significant additional information since bugs are often corrected in the original repository. These details can be viewed as our underlying assumptions regarding the usage of GitHub.

3.2. The chosen projects

We considered several criteria when searching for appropriate projects on GitHub. First of all, we searched for projects written in Java, especially larger ones, because those are more suitable for this kind of analysis. It was also important to have an adequate number of issues labeled as bugs, and the number of references from the log messages to certain commits is also a crucial factor (this is how we can link source code elements to bugs). Additionally, we preferred projects that are still actively maintained. Logged-in users can give a star for any repository and bookmark selected ones to follow. The number of stars and watches applied to repositories forms a ranking between them, which we will refer to as “popularity” in the following. We performed our search for candidate projects mainly based on popularity and activity. We also found many projects during the search that would have

fulfilled most aspects, had the developers not used an external bug tracker system – something we could not support yet.

In the end, we selected the 15 projects listed in Table 4 based on the previously mentioned criteria. As the descriptions show, these projects cover different domains; a good practice when the goal is creating a general dataset. The table contains the following additional data about the projects:

Stars the number of stars a project received on GitHub

Forks the number of forks of a project on GitHub

KLOC the thousand lines of code a project had at September, 2017

Recently, the repository of the Ceylon project was moved to a new location and the old repository is not available anymore. Due to this reason we could not obtain the total number of stars and the total number of forks of this repository, resulting the low values in the table.

Besides knowing each project’s domain, further descriptors can help us get a more precise understanding. Table 5 provides a more accurate picture of the projects by showing different characteristics (related to the repositories) for each project. This table sums up the occurrences of various bug reports and commits of the projects present at September, 2017. Considering the total number of commits (TNC) is a good starting point to show the scale and activity of the projects. The number of commits referencing a (closed) bug (NCRB) shows how many commits out of TNC referenced a bug by using the pattern ‘#x’ in their commit log messages, where x is a number that uniquely identifies the proper issue that is labeled as a bug (Mockus and Votta, 2000). NCBR (Number of Closed Bug Reports) is also important, since we only consider closed bug reports and the corresponding commits in this context. The abbreviations we used stand for the following:

TNC Total Number of Commits

NCRB Number of Commits Referencing a Bug

NBR Number of Bug Reports

NOBR Number of Open Bug Reports

NCBR Number of Closed Bug Reports

ANCBR Average Number of Commits per closed Bug Reports (NCRB/NCBR)

It is apparent that the projects are quite different according to the number of bug reports and the lines of code they have. NCRB is always lower than NCBR except in three cases (ANTLR v4, Oryx, Broadleaf Commerce) which means that not all bug reports have at least one referencing commit to fix the bug. This is possible since closing a bug is viable not only from a commit but directly from GitHub’s Web user interface without committing anything.

Fig. 1 depicts the number of commits for each closed bug report. One insight here is that the rate of closed bug reports is high where not even a single commit is present to fix the bug. There are several possible causes for this, for example, the bug report is not referred from the commit’s log message, or the error has already been fixed.

Fig. 2 shows the ratio of the number of commits per projects, illustrating the activity and the size of the projects. Neo4j is dominant if we consider only the number of commits, however bug report related activities are slight. The presented figures show the variability of the selected software systems, which ensures the construction of a heterogeneous dataset.

³ <https://developer.github.com/v3/>.

⁴ <https://www.gharchive.org/>.

Table 4

The selected projects and their descriptions.

Project name			Description
Stars	Forks	kLOC	
Android Universal Image Loader ^a 16,521	6357	13	An Android library that assists the loading of images.
ANTLR v4 ^b 6030	1559	68	A popular software in the field of language processing. It is a powerful parser generator for reading, processing, executing, or translating structured text or binary files.
Elasticsearch ^c 42,685	14,303	995	A popular RESTful search engine.
jUnit ^d 7536	2826	43	A Java framework for writing unit tests.
MapDB ^e 3700	745	68	A versatile, fast and easy to use database engine in Java.
mcMMO ^f 511	448	42	An RPG game based on Minecraft.
Mission Control Technologies ^g 818	280	204	Originally developed by NASA for the space flight operations. It is a real-time monitoring and visualization platform that can be used for monitoring any other data as well.
Neo4j ^h 6643	1636	1027	The world's leading graph database with high performance.
Netty ⁱ 20,006	9128	380	An asynchronous event-driven networking framework.
OrientDB ^j 3919	792	621	A popular document-based NoSQL graph database. Mainly famous for its speed and scalability.
Oryx 2 ^k 1633	388	34	An open-source software with machine learning algorithms that allows the processing of huge datasets.
Titan ^l 4931	1015	108	A high-performance, highly scalable graph database.
Eclipse plugin for Ceylon ^m 56	30	181	An Eclipse plugin which provides a Ceylon IDE.
Hazelcast ⁿ 3211	1169	949	A platform for distributed data processing.
Broadleaf Commerce ^o 1266	1020	322	A framework for building e-commerce websites.

^a<https://github.com/nostra13/Android-Universal-Image-Loader>.^b<https://github.com/antlr/antlr4>.^c<https://github.com/elastic/elasticsearch>.^d<https://github.com/junit-team/junit4>.^e<https://github.com/jankotek/MapDB>.^f<https://github.com/mcMMO-Dev/mcMMO>.^g<https://github.com/nasa/mct>.^h<https://github.com/neo4j/neo4j>.ⁱ<https://github.com/netty/netty>.^j<https://github.com/orientechnologies/orientdb>.^k<https://github.com/OryxProject/oryx>.^l<https://github.com/thinkaurelius/titan>.^m<https://github.com/eclipse/ceylon-ide-eclipse>.ⁿ<https://github.com/hazelcast/hazelcast>.^o<https://github.com/BroadleafCommerce/BroadleafCommerce>.

4. Metrics

A software metric is a quantified measure of a property of a software project. By using a set of different metrics, we can measure the properties of a project objectively from various points of view. Metrics can be obtained from the source code, from the project management system, or even from the execution traces of the source code. We can deduce higher-level software characteristics from lower level ones (Bakota et al., 2011), such as the maintainability of the source code or the distribution of defects, but they can be also used to build a cost estimation model, apply performance optimization, or to improve activities supporting software quality (Boehm et al., 2000; Bán et al., 2018;

Bán and Ferenc, 2014). In this work, we used static source code metrics (also known as software product metrics).

The area of object-oriented source code metrics has been researched for many years (Chidamber and Kemerer, 1994; Basili et al., 1996; Bruntink and Van Deursen, 2004), thus no wonder that several tools exist for measuring them. These tools are suitable for detailed examination of systems written in various programming languages. The source code metrics provide information about the size, inheritance, coupling, cohesion, or complexity of the code. We used the *OpenStaticAnalyzer*⁵ tool to obtain various software product metrics for the selected systems.

⁵ <https://github.com/sed-inf-u-szeged/OpenStaticAnalyzer>.

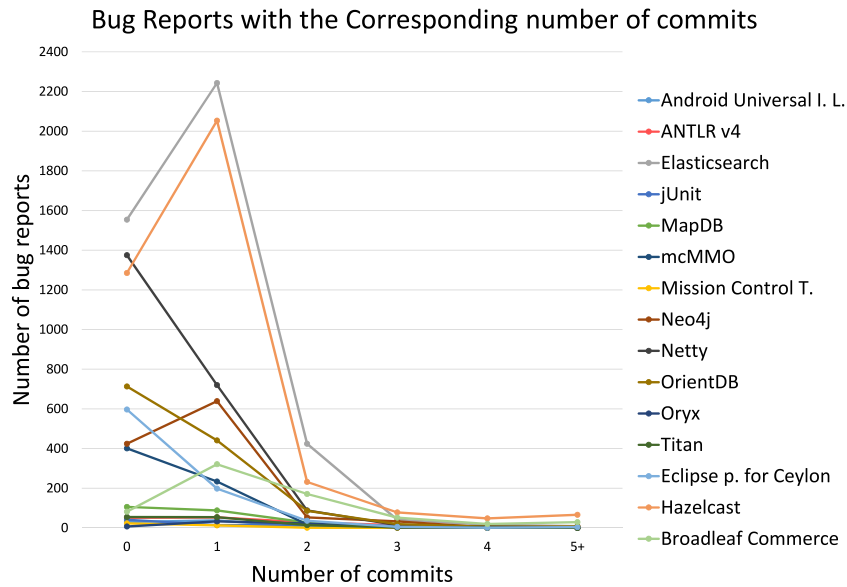


Fig. 1. The number of bug reports with the corresponding number of commits.

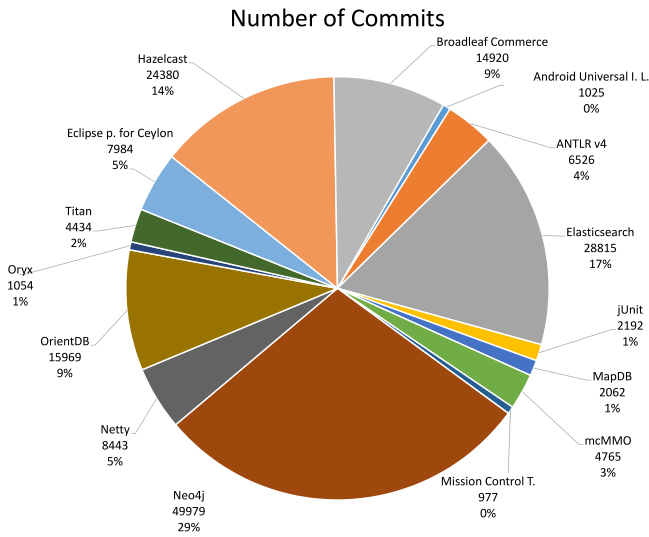


Fig. 2. The number of commits per projects.

Table 5
Statistics about the selected projects.

	TNC	NCRB	NBR	NOBR	NCBR	ANCBR
Android Universal I. L.	1,025	52	90	15	75	0.69
ANTLR v4	6,526	162	179	23	156	1.04
Elasticsearch	28,815	2,807	4,494	207	4,287	0.65
jUnit	2,192	72	90	6	84	0.86
MapDB	2,062	167	244	16	228	0.73
mcMMO	4,765	268	664	8	656	0.41
Mission Control T.	977	15	46	9	37	0.40
Neo4j	49,979	781	1,268	116	1,152	0.68
Netty	8,443	956	2,240	33	2,207	0.43
OrientDB	15,969	722	1,522	250	1,272	0.57
Oryx	1,054	69	67	2	65	1.06
Titan	4,434	93	135	8	127	0.73
Eclipse p. for Ceylon	7,984	316	923	82	841	0.38
Hazelcast	24,380	3,030	3,882	120	3,762	0.81
Broadleaf Commerce	14,920	1,051	703	28	675	1.56

The full list of the object-oriented metrics we used is shown in Table 6. The last three columns of the table indicate the kind

of elements the given metric is calculated for, namely method, class, and file. The presence of 'X' indicates that the metric is calculated for the given source code level. Most of the blanks in the table come from the fact that the metric is defined only for a given level. For instance, *Weighted Methods per Class* cannot be interpreted for methods and files. Other blanks come from the limitations of the used static source code analyzer (i.e. OpenStaticAnalyzer).

One special metric category is provided by source code duplication detection (Roy et al., 2009). OpenStaticAnalyzer is able to detect Type-1 (exact copy of code, not considering white spaces and comments) and Type-2 clones (syntactically identical copy of code where variable, function or type identifiers can be different; also not considering white spaces and comments) in software systems (Bellon et al., 2007) and also supports clone management tasks, such as:

- Clone tracking: clones are tracked during the source code analysis of consecutive revisions of the analyzed software system.
- Calculating clone metrics: a wide set of clone related metrics is calculated to describe the properties of a clone in the system (for example, risk of a clone or the effort needed to eliminate the clone from the system).

Basic clone related metrics that are calculated for methods and classes are presented in Table 7.

OpenStaticAnalyzer also provides a coding rule violation detection module. The presence of rule violations in a source code element can cause errors (Boogerd and Moonen, 2008) in a later phase (can easily be a ticking bomb); thus the number of different rule violations located in the source code element can serve as good predictors and the dataset encapsulates this information too.

5. Dataset creation

In this section, we introduce the methodology we used to create the dataset. We carried out the data processing in multiple steps using the toolchain shown in Fig. 3. Each of these steps – and their corresponding components – are detailed in their dedicated sections below.

Table 6

Source code metrics used for characterization.

Abbreviation	Full name	Method	Class	File
CLOC	Comment Lines of Code	X	X	X
LOC	Lines of Code	X	X	X
LLOC	Logical Lines of Code	X	X	X
NL	Nesting Level	X	X	
NLE	Nesting Level Else-If	X	X	
NII	Number of Incoming Invocations	X	X	
NOI	Number of Outgoing Invocations	X	X	
CD	Comment Density	X	X	
DLOC	Documentation Lines of Code	X	X	
TCD	Total Comment Density	X	X	
TCLOC	Total Comment Lines of Code	X	X	
NOS	Number of Statements	X	X	
TLOC	Total Lines of Code	X	X	
TLLOC	Total Logical Lines of Code	X	X	
TNOS	Total Number of Statements	X	X	
McCC	McCabe's Cyclomatic Complexity	X		X
PDA	Public Documented API		X	X
PUA	Public Undocumented API		X	X
HCPL	Halstead Calculated Program Length	X		
HDIF	Halstead Difficulty	X		
HEFF	Halstead Effort	X		
HNDB	Halstead Number of Delivered Bugs	X		
HPL	Halstead Program Length	X		
HPV	Halstead Program Vocabulary	X		
HTRP	Halstead Time Required to Program	X		
HVOL	Halstead Volume	X		
MIMS	Maintainability Index (Microsoft version)	X		
MI	Maintainability Index (Original version)	X		
MISEI	Maintainability Index (SEI version)	X		
MISM	Maintainability Index (SourceMeter version)	X		
NUMPAR	Number of Parameters	X		
LCOM5	Lack of Cohesion in Methods 5		X	
WMC	Weighted Methods per Class		X	
CBO	Coupling Between Object classes		X	
CBOI	Coupling Between Object classes Inverse		X	
RFC	Response set For Class		X	
AD	API Documentation		X	
DIT	Depth of Inheritance Tree		X	
NOA	Number of Ancestors		X	
NOC	Number of Children		X	
NOD	Number of Descendants		X	
NOP	Number of Parents		X	
NA	Number of Attributes		X	
NG	Number of Getters		X	
NLA	Number of Local Attributes		X	
NLG	Number of Local Getters		X	
NLM	Number of Local Methods		X	
NLPA	Number of Local Public Attributes		X	
NLPM	Number of Local Public Methods		X	
NLS	Number of Local Setters		X	
NM	Number of Methods		X	
NPA	Number of Public Attributes		X	
NPM	Number of Public Methods		X	
NS	Number of Setters		X	
TNA	Total Number of Attributes		X	
TNG	Total Number of Getters		X	
TNLA	Total Number of Local Attributes		X	
TNLG	Total Number of Local Getters		X	
TNLM	Total Number of Local Methods		X	
TNLPA	Total Number of Local Public Attributes		X	
TNLPM	Total Number of Local Public Methods		X	
TNLS	Total Number of Local Setters		X	
TNM	Total Number of Methods		X	
TNPA	Total Number of Public Attributes		X	
TNPM	Total Number of Public Methods		X	
TNS	Total Number of Setters		X	

Table 7

Clone metrics used for characterization.

Abbreviation	Full name
CC	Clone Coverage
CCL	Clone Classes
CCO	Clone Complexity
CI	Clone Instances
CLC	Clone Line Coverage
CLLC	Clone Lines of Code
LDC	Lines of Duplicated Code
LLDC	Logical Lines of Duplicated Code

changing on GitHub due to the activities in the projects, we need a consistent data source for the analysis. The data we save includes the list of users assigned to the repository (Contributors), the open and closed bug reports (Issues), and all of the commits. For open issues, we stored only the date of their creation. For closed issues, we stored the creation date, closing date, and the hash of the fixing commits with their commit dates. Additionally, we focused exclusively on bug related issues, so closed bugs that were not referenced from any commit were not stored. This filtering is based on the issue labels provided by GitHub and the set of labels we manually selected for each project. The data we stored about the commits includes the identifier of the contributor, the parent(s) of the commit, and the affected files with their corresponding changes. All this raw information is stored in an XML format, ready for further processing.

5.2. Processing raw data

While the data saved from GitHub includes all commits, we only need the ones that relate to the bug reports. These commits are then divided into different subsets, as depicted in Fig. 4. Green nodes are directly referencing the bug report (fixing intention). Gray nodes are commits applied between the first fix and the last fix but not referencing the bug id in their commit log messages. One extra commit taken into consideration is the one right before the first fix (colored with orange). This commit holds the state when the source code is buggy (not fixed yet), thus a snapshot (source code analysis) will be performed at that point too. Although the orange node represents the latest state where the bug is not fixed yet, the blue nodes also contain the bug so we mark the source code elements as buggy in these versions too. These blue markings are important for distinguishing commits that are involved in multiple bugs at the same time.

We have to perform code analysis on the orange and green commits to construct dataset entries. Two entries are created for every source code element they contain: one with the state (metrics) right before the fix was applied, and one with the state when the bug was fixed. At green commits except the last one, we do not need to perform a full code analysis, since at those points we are only interested in extracting the affected source code elements. Amongst the selected commits, some further ones can occur that need to be removed because they are no longer available through Git (deleted, merged). Moreover, we do not only search for links from the direction of commits but also from the direction of issues (bug reports). When considering a bug report, we can find a commit id showing that the bug was closed in that specific commit. At this point, the full list is constructed as a text file, which has all the commit ids (hash) for a selected project to undergo static analysis.

5.3. Source code analysis

After gathering the appropriate versions of the source code for a given project, feature extraction can begin. This component wraps the results of the OpenStaticAnalyzer tool that computes the source code metrics and determines the positions of the

5.1. Collecting data

First, we save data about the selected projects via the GitHub API. This is necessary, because while the data is continuously

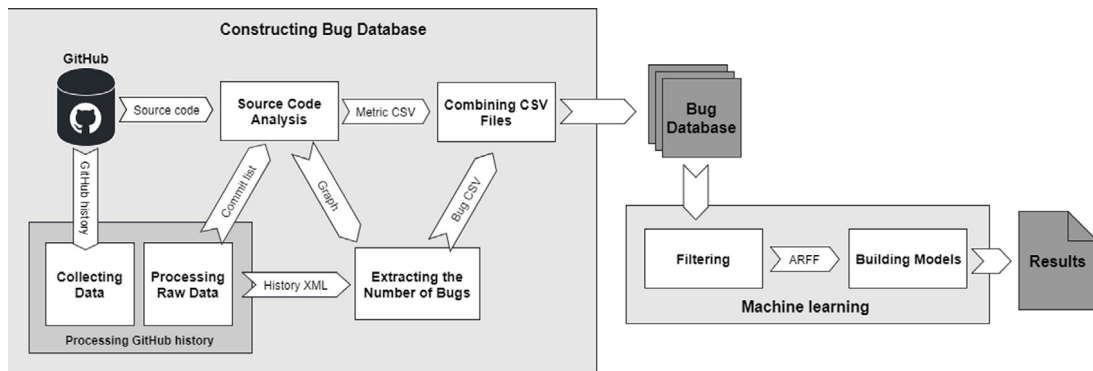


Fig. 3. The components of the process.

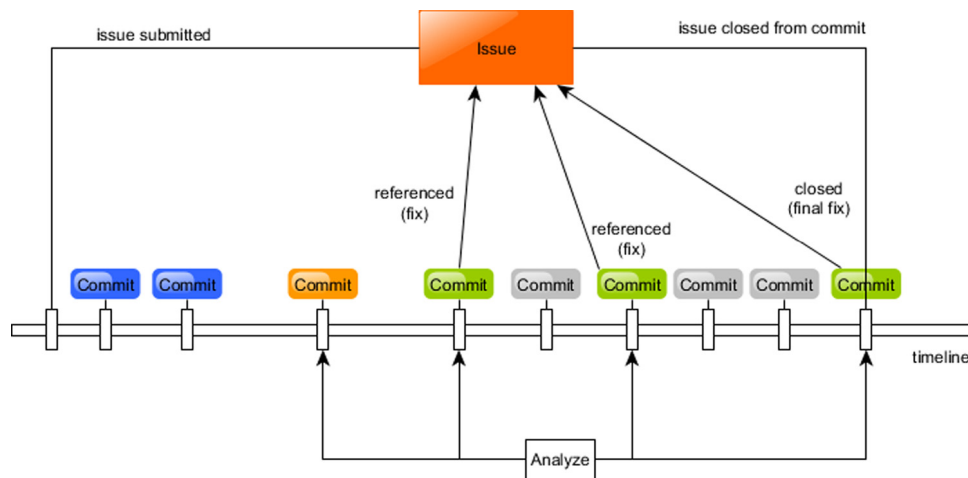


Fig. 4. The relationship between the bug reports and commits.

source code elements. Results are generated in a graph format, which contains the files, classes, and methods with the computed data that includes different software product metrics (described in Section 4). At this point we have all the raw data desired, including the source code elements located in the project and all the bug related information.

5.4. Extracting the number of bugs

The next step is to link the two datasets – the results of the code analysis and the data gathered from GitHub – and extract the characteristics of the bugs. Here, we determine the source code elements affected by the commits and the number of bugs in each commit for file, class, and method levels.

To determine the affected source code parts, an approach similar to the SZZ algorithm (Williams and Spacco, 2008) is used. However, we do not want to detect the fix inducing commits, only the mapping between the fixing code snippets and source code elements. For this purpose, we used the diff files – from the GitHub data we saved – that contain the differences between two source code versions in a unified diff format. An example unified diff file snippet is shown below.

```
--- /path/to/original''timestamp''
+++ /path/to/new''timestamp''
@@ -1,4 +1,4
@@+Added line
-Deleted line
This part of the
document has stayed the
```

same

Each diff contains a header information specifying the starting line number and the number of affected lines. Using this information, we can get the range of the modification (for a given file pair: original and new). To obtain a more accurate result, we subtracted the unmodified code lines from this range. Although the diff files generated by GitHub contain additional information about which method is affected, it does not carry enough information because the difference can affect multiple source code elements (overlapping cases that are not handled by GitHub). Thus, there is no further task but to examine the source code elements in every modified file and identify which ones of them are affected by the changes. The method uses the source code element positions, i.e., source line mappings from the output of the OpenStaticAnalyzer tool. We identified the source code elements by their fully qualified names that involve the name of the package, the class, the method, the type of the parameters, and the type of the return value.

Next, we take the commits that were selected by the “Processing Raw Data” step and mark the code sections affected by the bug in these commits. We do this by accumulating the modifications on the issue level and collecting the fully qualified names of the elements. Then, the algorithm marks the source code elements in the appropriate versions that will be entries in the dataset (touched in order to fix a bug). If a source code element in a specific version is marked by multiple issues, then it contains multiple bugs in that version. The dataset for files, classes, and methods are exported into three different files in a simple CSV format. The first row of these files contains the header

information, namely the commit id, the qualified name and the bug cardinality. Further lines store the data of the source code elements according to the header.

5.5. Combining CSV files

Now, the CSV outputs of OpenStaticAnalyzer and the previously described CSV output can be merged. In this phase, we attach the source code elements that are entries in the dataset to the calculated metrics. The output of this step is also a CSV file for each type of source code element, containing the hash code of the version, unique identifiers of the source code elements, identifiers of metrics, rule violation groups, and bug cardinality (the number of bugs located in the source code elements). One entry is equivalent to one source code element at a given time (the same source code element can occur more than once with a different commit id – hash).

5.6. Filtering

This dataset we compiled so far can contain various entries that complicate further investigations. As the dataset should be suitable for studying the connection between different metrics and bug occurrences, it should serve as a practical input for different machine learning algorithms. It is possible, however, to have entries in the dataset that have the same metric values with different number of bugs assigned to them. For example, let us consider a buggy method f with metric values M_{f_1} . After the bugfix, the metric values of f is changed to M_{f_2} . Similarly, let us consider another buggy method g with metric values M_{g_1} and M_{g_2} , respectively. These two methods could contain two different bugs that are present in a system for distinct periods of time. In this case, the dataset would contain 4 entries: M_{f_1} , M_{f_2} , M_{g_1} , M_{g_2} , where M_{f_1} and M_{g_1} are buggy and M_{f_2} and M_{g_2} are non-buggy entries. If any of these metric values are equal (e.g. $M_{f_1} = M_{g_2}$ or $M_{g_1} = M_{f_2}$), then redundancy occurs that can influence the accuracy of machine learning for bug prediction (overfitting, contradicting records).

To solve this issue, we used different approaches to filter the raw dataset and eliminate the redundant entries. We tried various methods to reduce the noise in the learning set, whose entries are classified into either buggy or not buggy.

- **Removal:** keep the entries located in the class with the larger cardinality (e.g., for a 10:20 distribution, the result is 0:20)
- **Subtract:** reduce the number of entries in the class with the larger cardinality by removing as many entries as the cardinality of the smaller class (e.g., for a 10:20 distribution, the result is 0:10)
- **Single:** remove the entries of the class with the smaller cardinality and hold only one entry from the larger one (e.g., for a 10:20 distribution, the result is 0:1)
- **GCF:** divide the number of entries of both classes by their greatest common factor (or greatest common divisor) and retain only the resulting amounts of entries from the classes (e.g., for a 10:20 distribution, the result is 1:2)

Each selected approach can seriously modify the result set, thus we investigated all four options and additionally the basic case when no filtering was applied. Tables 8, 10, and 12 present average F-measure values calculated for all of the machine learning algorithms we used for all of the projects. From these tables we can see that the Single and GCF methods performed quite similarly but were less effective than Subtract or Removal.

We employed a statistical significance test, namely the Friedman test (Friedman, 1940) with a threshold of $\alpha = 0.05$ to assess

Table 8

Filtering results at method level.

Method	Precision	Recall	F-measure
No filter	0.5553	0.5501	0.5317
Removal	0.6070	0.5963	0.5773
Subtract	0.5974	0.5893	0.5717
Single	0.5495	0.5448	0.5250
GCF	0.5445	0.5408	0.5218

Table 9

Significance test results for method level filtering.

	No filter	Removal	Subtract	Single
Removal	0.001 (16.88)			
Subtract	0.001 (13.87)	0.210 (3.00)		
Single	0.736 (1.67)	0.001 (18.54)	0.001 (15.54)	
GCF	0.020 (4.29)	0.001 (21.17)	0.001 (18.16)	0.343 (2.62)

the significance of the differences between the averages, as it was done similarly in previous bug prediction studies (Herbold et al., 2018; Ghotra et al., 2015). Our data does not follow normal distribution, it consists of dependent samples and we have five paired groups, thus the Friedman test is the appropriate choice. The null hypothesis is that the multiple paired samples have the same distribution. The tests resulted in very low p values ($p_{method} = 5.32e-80$, $p_{class} = 2.03e-77$, $p_{file} = 1.83e-40$); therefore, we reject the null hypothesis which means the distributions are not equal. Then, we applied the Nemenyi post-hoc test (Nemenyi, 1963) ($\alpha = 0.05$) that is usually used after a null hypothesis is rejected to gain more insight on the significance of the differences. The critical value for 5 groups and 176 samples (11 machine learning algorithms \times 16 databases) based on the studentized range table is $q_{crit} = 3.9$. Tables 9, 11, and 13 list the resulted p values with the corresponding rank difference in parentheses.

Let us consider the method level F-measure values in Table 8 where Removal has the highest average F-measure (0.5773) and Subtract is a close second (0.5717). In Table 9, the results of the significance tests for method level show that the p value of the test between Subtract and No filter is below the threshold ($p = 0.001 < \alpha = 0.05$); therefore, the difference is significant and with Subtract having a higher average F-measure (0.5717) than No filter (0.5317), we can state that it is significantly better. We can conclude the same when comparing Subtract with Single ($p = 0.001 < \alpha = 0.05$) or with GCF ($p = 0.001 < \alpha = 0.05$). The p value between Subtract and Removal is $p = 0.210 > \alpha = 0.05$ which is not significant.

Similar results can be concluded for class level and for file level as well. We can state that the Removal and Subtract methods performed significantly better than the other methods in all three cases. The difference between the Removal and Subtract methods is not significant.

We speculate that a disadvantage to Single is that it drops the multiplicity of the records (i.e., the weight information). The problem with GCF, on the other hand, is that it will only perform filtering when the greatest common factor is not one, and that it does not eliminate the noise completely (i.e., it will keep at least one entry from both classes). Removal removes the noise entirely, but it suffers from the fact that it ignores the minority.

The Subtract method, however, neutralizes the positive and negative entries with identical feature vectors. This means that it removes the noise while also keeping the weight of the records, so this filtering method seems to be the best choice. Presenting all the five different sets would be lengthy, thus we will only present the results achieved by the Subtract method.

5.7. Classification and resampling

As trying to predict the exact number of bugs in a given source code element would be much more difficult – and would

Table 10
Filtering results at class level.

Method	Precision	Recall	F-measure
No filter	0.5265	0.5235	0.5128
Removal	0.5567	0.5528	0.5419
Subtract	0.5541	0.5499	0.5393
Single	0.5236	0.5206	0.5090
GCF	0.5221	0.5201	0.5077

Table 11
Significance test results for class level filtering.

	No filter	Removal	Subtract	Single
Removal	0.0010 (15.92)			
Subtract	0.0010 (15.40)	0.9000 (0.52)		
Single	0.9000 (0.76)	0.0010 (16.69)	0.0010 (16.16)	
GCF	0.0496 (3.86)	0.0010 (19.78)	0.0010 (19.26)	0.1828 (3.10)

Table 12
Filtering results at file level.

Method	Precision	Recall	F-measure
No filter	0.5160	0.5117	0.4883
Removal	0.5451	0.5414	0.5194
Subtract	0.5407	0.5371	0.5147
Single	0.5187	0.5148	0.4910
GCF	0.5172	0.5129	0.4889

Table 13
Significance test results for file level filtering.

	No filter	Removal	Subtract	Single
Removal	0.0010 (14.21)			
Subtract	0.0010 (13.21)	0.9000 (1.00)		
Single	0.0682 (3.69)	0.0010 (10.51)	0.0010 (9.51)	
GCF	0.9000 (0.36)	0.0010 (13.85)	0.0010 (12.85)	0.1262 (3.34)

presumably require much larger datasets – we chose to restrict our study to predicting a boolean “flag” for the presence of any bugs. Thus, we applied only classification algorithms, and to do so, classes need to be formed from the bug numbers. For the binary classification, we selected instances with zero bugs into one class (non-buggy), and the remaining ones – with one or more bugs – into the second class (buggy).

Another problem we faced is that imbalanced learning sets could be formed from the dataset, where the positive or negative entries are in a majority which could also be misleading for model training. For example, the ratio of buggy and non-buggy source code elements or files can be totally different.

To handle this issue at the machine learning level, we used random under sampling (He et al., 2009; Wang and Yao, 2013) to obtain an equivalent number of elements in the two categories. For instance, if we have a final set as a corpus at method level that contains 10 buggy methods and 50 non-buggy methods, we use random under sampling for the non-buggy set to decrease the number of samples and balance the ratio to 10-10. The training process – using this random under sampling – is repeated multiple times and finally, an average is calculated. Without random under sampling, the machine learning algorithms achieved very high precision, recall, and F-measure values (e.g., by classifying all elements as non-buggy) because a significantly large difference was usually present in the number of entries for the two classes (non-buggy vs. buggy). For this reason, we present only the results achieved by using random under sampling.

5.8. The BugHunter dataset

As a result and a main contribution, we constructed a novel kind of bug dataset that contains before/after fix states of source code elements at file, class, and method levels. We produced

Table 14
BugHunter dataset metadata.

Type	Bug prediction dataset
Granularity	File, class, method
Number of projects	15
Number of metrics	Static source code metrics (66), Code duplication metrics (8), Code smell metrics (35)
Number of entries	File: 70,088 Class: 84,562 Method: 159,078
Ratio of the faulty entries	File: 1.03 (35,507/34,581) Class: 0.95 (41,098/43,475) Method: 0.59 (58,810/100,268)

a dataset for every project (see Table 4) and also a combined one with all the projects included. In Table 14, we collected the general metadata about the dataset which is scattered throughout in the paper. The *Ratio of the faulty entries* varies between the different granularity levels due to the nature of the bug-fixing changes. For example, on method level it is more common to split a method into multiple parts or to introduce new methods to the source code during a fix, which results in more non-buggy entries than buggy, hence the low ratio. On class level, however, the ratio is closer to 1.0 because it is less likely to create a new class in order to fix a bug. Furthermore, on class and file levels, a new factor also contributes to the ratio. Because classes and files represent a larger part of the source code than methods, it happens that in a fixed state of a bug the containing class or file contains other bugs as well. In these cases, the entries related to the fixed state are marked as faulty in the dataset, which results in higher ratio values, above 1.0 in the case of files.

The resulting BugHunter Dataset 1.0 is available as an online appendix at: <http://www.inf.u-szeged.hu/~ferenc/papers/BugHunterDataSet/> or <http://dx.doi.org/10.17632/8tx7kjbkg4.2>

The BugHunterDataset-1.0.zip file contains the dataset in CSV format as described above. The directory named `full` contains the unfiltered database. The remaining four directories, namely `gcf`, `remove`, `single`, and `subtract` contain the results of the different filtering methods. Each of these directories contain 15 subdirectories – one for each subject system – and an additional directory named `all` which contains the summarized dataset. Three CSV files are placed in these directories for file, class, and method levels, respectively. There is also a fourth CSV file in each directory, called `method-p.csv` which contains the method level dataset extended with an additional column, the name of the parent class (see Section 6.2). Additionally, the `appendix.zip` file contains the analysis results presented in Section 6 in spreadsheet files.

5.9. Validation

When constructing a dataset in an automatic way, one always has to validate the constructed set. As seen previously, this kind of generated data should always be handled with mistrust (Kalliamvakou et al., 2014). We chose the mid-sized project JUnit for such manual validation, which contains 90 bug reports (6 open and 84 closed) and a total of 72 referencing commits to the bugs, thus this project seems to be suitable for manual validation investing a reasonable amount of effort. We validated the 84 closed bug reports manually to verify whether the bugs are valid, and whether the matching algorithm works well. Table 15 summarizes our findings.

From the total number of closed bugs, only 37 are present in the dataset because many fixing commits are not related to the source code (e.g., documentation) or Java code (e.g., bug in build

Table 15
Validation results.

Closed bugs	Bugs in dataset	Commits	Java code	Commit mismatch
84	37	72	61	5

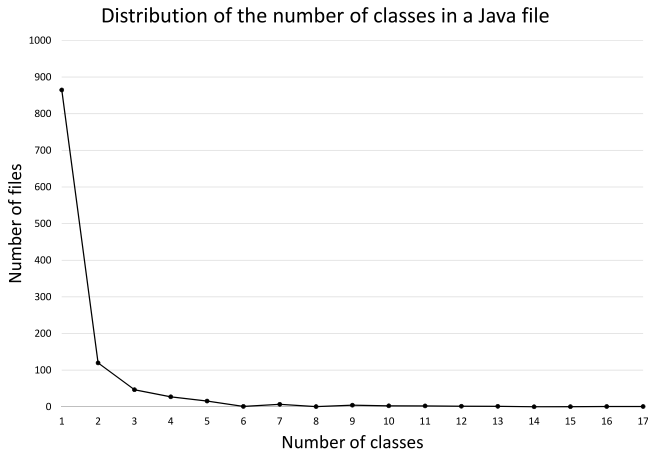


Fig. 5. Distribution of the number of classes in a Java file.

XML). This is shown in the Java code column of the table that summarizes the number of commits that contain Java language code (61). These commits that include code are referencing 37 bugs, thus at least one bug exists that is referenced from multiple commits according to the pigeonhole principle.

We found 5 “commit mismatch” cases in total, where only comments were modified in the source code – this means 7 entries in the dataset. The dataset created for JUnit has 734 entries in total (92 files, 216 classes and 426 methods), thus a very small (0.95%) number of entries was incorrectly included. Out of the 734 entries, 286 are not related to test code (43 files, 77 classes, and 166 methods). Based on this, we can presume that our validated dataset can be an appropriate corpus for further investigations and that our bug extraction mechanism is working quite reliably.

5.10. Relationship of files and classes

In case of Java, there is usually one class per *.java* file. We examined how true this is for our subject projects. We randomly chose 100 commits that we analyzed from each project and we counted the number of classes in each file, not including test files. After we calculated the frequency of these values for each commit, we calculated the average frequency (see Fig. 5). The diagram shows that most of the files contain only a single class (865), but there is a significant number of files with more than one class (120 files with 2 classes, 46 files with 3 classes, etc.).

Although we have a larger set of metrics on class level than on file level, we cannot associate a file to a single class due to the one-to-many relationship between them.

5.11. Computational cost of extending the dataset

Extending the dataset comes with computational cost that depends on multiple factors. Adding a new project to the dataset requires the project to have bug reports with bug-fixing commits. Finding such projects is time-consuming, because it mostly requires manual work to select good candidates (see Section 3.2).

The most critical step is to collect appropriate bugs. For this initial dataset, we collected projects from GitHub, since its API makes it easy to gather the required information about bugs automatically. The actual run-time of this step depends on the

size of the project, e.g. number of commits and number of bug reports, but for the selected 15 projects, it took just a few hours to save the required data. Selecting bugs manually would take considerably more time. GitHub has a limit on the number of API requests per hour which increased the total run-time. It is possible to collect data from other sources, although it may require a different amount of work.

The most time and resource consuming task is the source code analysis. We used the OpenStaticAnalyzer tool which performs deep static analysis; therefore, it requires more resources than a simple parser tool. During this step, we extract the static source code metrics and the source code positions of the classes and methods, as described in Sections 5.2 and 5.3, respectively. The computational cost of this step highly depends on the number of bug-fixing commits, the size of the source code, and the analyzer tool. It took days to analyze each of the nearly 10,000 bug-fixing commits. There are other tools that could be used to extract the source code positions and other tools to compute metrics with potentially less run-time, but the tool we used produces a wide range of metrics and rule violations accurately in a well processable format.

The next step, determining the buggy source code elements, is a simple algorithm that does not require much resources. The run-time here mostly depends on the number of bug-fixing commits. It took only a few hours for the 15 projects.

For example, to process a smaller project such as JUnit, it took around 2 h of machine time: 10 min to download the data from GitHub, 110 min to analyze 107 versions of the project (on average 1 min per version) and around 2 min to produce the bug dataset entries. Regarding a larger project, Elasticsearch, it took around 6 h to download the data from GitHub, around 1600 h to analyze 4881 versions of the project (on average 20 min per version) and it took around 90 min to produce the bug dataset.

At this point, the data is ready to be added to the dataset. The last step is to match the format of the dataset (see Section 5.8). Since the dataset consist of CSV files, it is very easy to extend it with new projects or with additional bugs for the projects that are already present.

6. Evaluation

To evaluate the usefulness of our new bug dataset, we created bug prediction models by using machine learning algorithms. During the training, we used 10-fold cross validation to measure the accuracy of the models. To compare the models, we used precision, recall, and F-measure metrics that are defined as follows:

$$\text{precision} = \frac{TP}{TP + FP} \quad \text{recall} = \frac{TP}{TP + FN}$$

$$F\text{-measure} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}},$$

where *TP* (True Positive) is the number of methods/classes/files that were predicted as faulty and observed as faulty, *FP* (False Positive) is the number of methods/classes/files that were predicted as faulty but observed as not faulty, *FN* (False Negative) is the number of methods/classes/files that were predicted as non-faulty but observed as faulty. We used the following algorithms from the Weka library to find out how they perform on our dataset:

- NaiveBayes
- NaiveBayesMultinomial
- Logistic
- SGD
- SimpleLogistic

Table 16

TOP 5 machine learning algorithms for method level based on F-measure.

Algorithm	Precision	Recall	F-measure
trees.RandomForest	0.6335	0.6324	0.6319
trees.J48	0.6147	0.6134	0.6119
trees.RandomTree	0.6115	0.6113	0.6110
functions.SimpleLogistic	0.6062	0.6043	0.6031
rules.DecisionTable	0.6138	0.6073	0.5983

- VotedPerceptron
- DecisionTable
- OneR
- J48 (C4.5)
- RandomForest
- RandomTree

We have 3 source code levels (method, class, file), 15 chosen projects (plus the summarized dataset), and 11 machine learning algorithms, implying that full tables with the obtained results would be too large to present in the paper, thus we introduce only the best algorithms here for the overall dataset. Please note that the online appendix (see Section 5.8 for the Web link) contains all the analysis results in spreadsheet files.

6.1. First research question

The first research question we will answer is the following:

RQ1: *Is the constructed dataset usable for bug prediction purposes?*

To answer RQ1, we present the best results obtained by different machine learning algorithms at method, class, and file level. Similar to Section 5.6, we used the Friedman test and the Nemenyi post-hoc test to check whether the distributions of the samples are equal or not. We observed the same, very low p values ($p_{method} = 1.21e-14$, $p_{class} = 1.54e-07$, $p_{projected} = 3.77e-14$, $p_{file} = 1.06e-05$), thus the distributions are not equal. With $\alpha = 0.05$, the critical value in this case for 11 groups (machine learning algorithms) and 16 samples (databases) based on the studentized range table is $q_{crit} = 5.256$. Due to the size of the tables, we do not include them here, instead, the complete tables are available in the [Appendix A](#) (see [Tables 29, 30, 31](#), and [32](#)).

6.1.1. Method level

We trained models to use method level metrics to predict future failures at method level. The results are shown in [Table 16](#) containing the best five algorithms selected by F-measure values.

The fifth best algorithm with 0.5983 F-measure value is DecisionTable ([Kohavi, 1995](#)). The SimpleLogistic algorithm resulted a slightly higher F-measure (0.6031). SimpleLogistic algorithm builds linear logistic regression models that uses automatic attribute selection ([Landwehr et al., 2005](#)). The first three algorithms are all from the tree family. J48 ([Quinlan, 1993](#)) that uses pruned or unpruned C4.5 decision tree to build a model is the second best algorithm (0.6119). The third and the first algorithms also use trees to produce prediction models. RandomForest (0.6319) builds a forest from RandomTrees to get a slightly better result than RandomTree (0.6110). The results of the statistical tests in [Table 29](#) show that the differences between the top five algorithms are not statistically significant ($p > \alpha = 0.05$) but the difference between the worst (NaiveBayes, NaiveBayesMultinomial and VotedPerceptron) and the best performing algorithms are significant.

At method level, trees were performing the best and can result up to 0.6319 when considering F-measure values. We also investigated the results by projects and found that specific projects

Table 17

The best F-measure values by projects at method level.

Project	F-measure	Algorithm
antlr4	0.7573	trees.RandomForest
BroadleafCommerce	0.7366	trees.RandomForest
hazelcast	0.7170	trees.RandomForest
mct	0.6876	trees.RandomForest
oryx	0.6678	trees.RandomForest
junit	0.6638	rules.DecisionTable
all	0.6622	trees.RandomForest
netty	0.6412	trees.RandomForest
elasticsearch	0.6411	trees.RandomForest
orientdb	0.6236	trees.RandomForest
titan	0.6216	functions.SGD
neo4j	0.6086	functions.Logistic
mcMMO	0.5815	trees.RandomForest
MapDB	0.5610	functions.SimpleLogistic
Android-Universal-Image-Loader	0.5569	functions.SGD
ceylon-ide-eclipse	0.5395	trees.RandomTree

Table 18

TOP 5 machine learning algorithms for class level based on F-measure.

Algorithm	Precision	Recall	F-measure
functions.SimpleLogistic	0.5760	0.5763	0.5685
rules.DecisionTable	0.5703	0.5705	0.5637
functions.SGD	0.5718	0.5676	0.5626
functions.Logistic	0.5561	0.5552	0.5537
trees.J48	0.5531	0.5530	0.5520

performed worse than others. Android Universal Image Loader and Ceylon were the worst, both when considering precision, recall, or F-measure. Achieved F-measure values depend highly upon the project itself. A possible factor that plays role in this is the size of the built corpus. These projects have a smaller training corpus and more inconsistency in the feature vectors, consequently it is harder to build a well-performing prediction model. This phenomenon does not appear only at method level but at class and file level too, since at these levels even less entries are created in the dataset. The best F-measure values (over 0.75 in one case) achieved on different projects are demonstrated in [Table 17](#).

6.1.2. Class level

When considering class level, we have quite a different set of algorithms in the top five than in the case of methods. Furthermore, the precision, recall, and F-measure values differ significantly from those we obtained at method level. We suspect that the main reason behind this is the different set of metrics used to build models, thus to predict the possibility of occurring bugs in a class. At class level, simple logistic, decision table, and SGD were the best. Function and rule based groups of machine learning algorithms can be emphasized as the best when considering class level. The best machine learning algorithms at class level are shown in [Table 18](#) with F-measure values around 0.56. In [Table 30](#), the results of the significance tests show that the best algorithm, SimpleLogistic with 0.5685 F-measure, achieved significantly better results than the worst two algorithms that are not in the top five (NaiveBayes $p = 0.019$ and VotedPerceptron $p = 0.001$). Between the top five algorithms, the differences are not significant ($p > \alpha = 0.05$).

The low F-measures values suggest that one cannot build efficient prediction models at class level.

However, we present the F-measure values of individual projects in [Table 19](#). Considering these F-measure values, we can see the same phenomenon as in the case of methods. McMMO, Android Universal Image Loader, and Neo4J are in the worst 5, which supports the previous experience according to which different projects provide different amount of “munition” for predicting faults. The best case, however, provides an F-measure of 0.74.

Table 19

The best F-measure values by projects at class level.

Project	F-measure	Algorithm
BroadleafCommerce	0.7400	trees.RandomForest
oryx	0.7095	functions.SGD
junit	0.6639	rules.DecisionTable
hazelcast	0.6175	trees.RandomTree
MapDB	0.6138	functions.SimpleLogistic
orientdb	0.6132	functions.SimpleLogistic
mct	0.5825	functions.SGD
elasticsearch	0.5817	rules.DecisionTable
all	0.5803	rules.DecisionTable
ceylon-ide-eclipse	0.5789	rules.DecisionTable
antlr4	0.5685	rules.OneR
mcMMO	0.5670	functions.SGD
titan	0.5614	functions.SimpleLogistic
netty	0.5537	functions.Logistic
neo4j	0.5413	functions.SimpleLogistic
Android-Universal-Image-Loader	0.4713	bayes.NaiveBayes

Table 20

TOP 5 machine learning algorithms for file level based on F-measure.

Algorithm	Precision	Recall	F-measure
trees.RandomTree	0.5484	0.5484	0.5476
trees.RandomForest	0.5458	0.5461	0.5455
functions.Logistic	0.5528	0.5474	0.5367
rules.OneR	0.5358	0.5359	0.5348
functions.SimpleLogistic	0.5491	0.5474	0.5321

6.1.3. File level

In Java context, a public class is almost equivalent to a file with a '.java' extension. However, despite the fact that we compute a different set of metrics for class and file level, the results are quite similar. Since we operate on a different set of metrics at class and file level, this explains that different machine learning algorithms performed the best. The best algorithms for this level use tree-based approaches to predict bugs as it is shown in Table 20. Similar to class level, the differences between the top five algorithms are not considered significant ($p > \alpha = 0.05$), as can be seen in Table 32. The best performing algorithm, achieving an F-measure value of 0.5476, is RandomTree. The top five algorithms achieved significantly better results compared to the worst algorithm (VotedPerceptron). Between the top five algorithms, the differences are not considered significant ($p > \alpha = 0.05$).

We also present the F-measure values obtained on projects in Table 21. The takeaway remains the same, Android Universal Image Loader and Neo4j are located in the worst five projects again. On the other hand, Broadleaf Commerce, Oryx, and Hazelcast seem to be appropriate to use in model building. The best F-measure value is over 0.77.

Answering RQ1: Considering the results we obtained, we can state that creating bug prediction models at method level is more successful than at file and class levels if we consider the full dataset. We also showed the diversion in F-measure values by projects, which strengthens our assumption that not all projects are capable of providing an appropriate training set. We can obtain F-measure values on separate projects up to 0.7573, 0.7400, and 0.7741 at method, class and file level, respectively, which is promising. In our next research question, we accomplish an experiment and its results are even better. However, even without knowing that there is a better solution, we can answer this research question in a positive manner and say that the constructed dataset is usable for bug prediction.

Table 21

The best F-measure values by projects at file level.

Project	F-measure	Algorithm
BroadleafCommerce	0.7741	trees.RandomForest
oryx	0.6458	bayes.NaiveBayesMultinomial
hazelcast	0.6417	trees.RandomTree
all	0.6234	trees.RandomTree
orientdb	0.6200	rules.DecisionTable
elasticsearch	0.6073	trees.RandomTree
ceylon-ide-eclipse	0.5857	trees.J48
titan	0.5793	functions.SimpleLogistic
mcMMO	0.5702	trees.RandomForest
MapDB	0.5525	rules.OneR
junit	0.5484	rules.OneR
netty	0.5344	trees.RandomTree
antlr4	0.5212	trees.RandomForest
neo4j	0.5138	rules.DecisionTable
Android-Universal-Image-Loader	0.4781	rules.DecisionTable
mct	0.4576	functions.Logistic

6.2. Second research question

The dataset contains the bug information on both method and class levels, and we also know the containing relationships between classes and methods. However, since classes have a different source code metrics set than methods, a question arose: can we (and more importantly, should we) use method level metrics to predict faulty classes? The second research question we will answer is the following:

RQ2: Are the method level metrics projected to class level better predictors than the class level metrics themselves?

We carried out an experiment where we projected the results of the method-level learning to the class level. During the cross-validation of the method level learning, we used the containing classes of the methods to calculate the confusion matrix from the number of classes classified as buggy and non-buggy. Classes containing at least one buggy method were considered as buggy.

We compared this result with the result of the class-level prediction. The results in Table 22 show that the projection method performs much better than the prediction with class level metrics.

We applied the Wilcoxon-signed-rank test (Wilcoxon, 1945) (a non-parametric paired test for dependent samples), with a threshold of $Z_{crit} = 1.96$ (for a two-tailed test with $\alpha = 0.05$) to check whether the difference is significant. We also calculated the effect size of these tests with the Pearson correlation coefficient (Pearson's r) from the formula $r = \frac{Z}{\sqrt{N}}$, where N is the total number of samples and Z is the z-score of the test (Rosenthal et al., 1994). According to Cohen (1992), the effect size is considered small if $r \approx 0.1$, medium if $r \approx 0.3$ or large if $r \approx 0.5$.

After the test, we can confirm that the difference between the projection method and the prediction with class level metrics is significant ($Z = 10.9 > Z_{crit} = 1.96$) and the effect size is considered large ($r = 0.58$).

We suspect that this is due to the generality of class-level metrics, which are therefore not powerful enough to effectively distinguish source code bugs. Although the bug information for methods does not include all bugs that affect the containing class (e.g. change of fields, interfaces or superclasses), method level metrics are more useful for bug prediction.

The results of the significance tests between the different machine learning algorithms is displayed in Table 31. The best performing algorithm is RandomForest with 0.7405 F-measure and it is significantly better than the worst three algorithms

Table 22

The results of projected learning.

Algorithm	Precision		Recall		F-measure	
	Projected	Class	Projected	Class	Projected	Class
trees.RandomForest	0.7471	0.5336	0.7370	0.5336	0.7405	0.5334
trees.RandomTree	0.7421	0.5381	0.7273	0.5380	0.7330	0.5376
functions.SGD	0.7441	0.5718	0.7288	0.5676	0.7322	0.5626
rules.DecisionTable	0.7425	0.5703	0.7404	0.5705	0.7309	0.5637
trees.J48	0.7390	0.5531	0.7250	0.5530	0.7290	0.5520

that are not displayed in Table 22 (NaiveBayes $p = 0.001$; NaiveBayesMultinomial $p = 0.001$; VotedPerceptron $p = 0.003$). The difference between the top five algorithms is not considered significant ($p > \alpha = 0.05$).

When using the projection approach to predict bugs in classes, the F-measure values reach 0.74. As an extension of the answer to RQ1, we can provide the above described mechanism to locate class level bugs with a higher accuracy in a software system.

Answering RQ2: Using method level metrics for class level bug prediction performed the best in our study. This fact also contributes to the answer given for RQ1. Furthermore, method level metrics are better predictors when projected to class level than class level metrics by themselves.

6.3. Third research question

The third research question we will answer is the following:

RQ3: Is the BugHunter Dataset more powerful and expressive than the GitHub Bug Dataset, which is constructed with the traditional approach?

Comparing the expressive power of different datasets is a harsh task since the various datasets were created with different purposes, they often have only few independent variables in common. The projects included in these datasets are different as well. Therefore, we provide an objective comparison between our previously published traditional bug dataset, the GitHub Bug Dataset (Tóth et al., 2016) and the BugHunter Dataset in the following. These two datasets include exactly the same 15 projects and the set of independent variables are common and also calculated in the same way with the same tool. We used the same machine learning algorithms to build prediction models. This

way, it is quite straightforward to compare the expressiveness and compactness of these datasets.

Firstly, we compare the size of the datasets expressed with the number of entries located in the datasets. Table 23 shows the number of entries at method, class, and file level. The number of entries contained in the traditional dataset are listed in the “Trad” column. The “BH” column represents the number of entries in the BugHunter dataset, while “Rate” is calculated as follows:

$$\text{Rate} = \frac{\# \text{ of entries in the traditional dataset}}{\# \text{ of entries in the BugHunter dataset}}$$

The obtained rate is higher than 1.0 for most of the projects in case of the method level, which shows that the new approach contains less entries at this level. A rate of 1.54 is achieved at method level, 0.41 at class level, and 0.33 at file level. It is important to note that the traditional dataset encompasses data for only a six-month long interval which has the most bugs in it. On the other hand, the BugHunter dataset contains information from the beginning of the project up to September, 2017. One would expect that the new approach will contain less entries than the traditional one since the BugHunter dataset contains only the entries which were affected by a closed bug. However, the traditional dataset only depends on the size (number of files, classes, and methods) of the projects included. In contrast, the BugHunter dataset highly depends on the number of closed bugs in the system (large projects can have small amount of reported bugs). Even if no feature development was performed on a project (the size of the project remains quasi-same: in general no new files and classes are added, only modified) the number of closed bugs imply more entries in the BugHunter dataset, while size is not affected in any way in the traditional dataset.

To sum up, we cannot clearly decide whether the new dataset is more compact; however it is clearly visible that BugHunter could compress the bug related information at method level. We achieved an F-measure value of 0.6319 at method level (see Table 24) and the composed dataset contains 58,464 less entries than the traditional one. In both datasets, the number of entries are sufficient to build a predictive model from; however we should investigate the predictive capabilities first to conclude our findings related to expressive power and compactness.

In the following, we present tables that capture the differences of the prediction capabilities between the two datasets (using F-measures, as before). Tables 24, 25, 26, and 27 present machine learning results for method, class, file levels and also the F-measure values for the projected method level predictors, respectively. The complete tables are not presented here due to lack of space, however average, standard deviation, min, and max

Table 23

Comparison of the size of the datasets.

Project	Method			Class			File		
	Trad	BH	Rate	Trad	BH	Rate	Trad	BH	Rate
Android-U-I-L	432	325	1.33	73	156	0.47	63	145	0.43
antlr4	3,640	840	4.33	479	314	1.53	411	347	1.18
BroadleafCommerce	14,651	4,709	3.11	1,593	2,957	0.54	1,719	2,969	0.58
ceylon-ide-eclipse	8,787	2,087	4.21	1,611	1,275	1.26	700	946	0.74
elasticsearch	34,324	35,862	0.96	5,908	24,994	0.24	3,035	17,724	0.17
hazelcast	21,642	32,973	0.66	3,412	19,845	0.17	2,228	14,913	0.15
junit	2,441	462	5.28	731	316	2.31	309	177	1.75
MapDB	2,913	1,456	2.00	331	899	0.37	138	482	0.29
mcMMO	2,531	1,184	2.14	301	732	0.41	267	678	0.39
mct	9,836	105	93.68	1,887	66	28.59	413	52	7.94
neo4j	30,256	7,030	4.30	5,899	3,701	1.59	3,278	2,934	1.12
netty	8,312	11,171	0.74	1,143	5,677	0.20	914	4,023	0.23
orientdb	17,013	9,445	1.80	1,847	4,134	0.45	1,503	3,564	0.42
oryx	2,506	810	3.09	533	598	0.89	281	536	0.52
titan	8,424	785	10.73	1,468	428	3.43	976	378	2.58
Total	167,708	109,244	1.54	27,216	66,092	0.41	16,235	49,868	0.33

Table 24

Predictive capabilities – Method level.

Dataset	Avg.	Std.dev.	Min	Max
BugHunter	0.6319	0.0836	0.3376	0.7573
Traditional	0.7348	0.0789	0.4019	0.8339

Table 25

Predictive capabilities – Class level.

Dataset	Avg.	Std.dev.	Min	Max
BugHunter	0.5685	0.0704	0.3572	0.7400
Traditional	0.7710	0.0869	0.3446	0.8331

Table 26

Predictive capabilities – File level.

Dataset	Avg.	Std.dev.	Min	Max
BugHunter	0.5147	0.0749	0.3328	0.7741
Traditional	0.6058	0.1076	0.2882	0.8247

Table 27

Predictive capabilities – Projected.

Dataset	Avg.	Std.dev.	Min	Max
BugHunter	0.7405	0.0914	0.3178	0.8386
Traditional	0.7831	0.0716	0.4399	0.8825

Table 28

Uncertainty in the traditional dataset.

Project	Average days	Average commits before reported	Average commits before fixed
Android-U.-I.-L.	78.78	179.04	22.82
antlr4	83.73	94.83	66.21
BroadleafCommerce	96.40	524.88	116.74
ceylon-ide-eclipse	136.05	442.00	20.22
elasticsearch	93.85	1004.60	382.79
hazelcast	84.61	1905.88	143.54
junit	91.94	76.71	171.09
MapDB	102.09	150.47	25.06
mcMMO	108.71	289.83	41.72
mct	64.00	203.00	55.93
neo4j	39.53	535.77	189.30
netty	83.65	411.60	48.96
orientdb	99.21	568.76	179.30
oryx	63.00	104.42	3.40
titan	51.35	65.91	59.85

values are calculated and included in the tables which provide a general picture for the comparison. The appendix.zip file supplied as an online appendix (see Section 5.8) contains the complete tables with all F-measure values. For the sake of clarity, we describe how we obtained the averages presented here in the paper in detail. First, since the traditional dataset consists of multiple versions with bugs from six-months long intervals, for each project we selected the version from the traditional dataset that has the most number of bugs assigned to them. After collecting the machine learning results of the selected versions, we calculated average F-measure values for each algorithm we used. Then we ranked the algorithms based on these averages and we selected the one with the highest average value. We used this average value for the traditional dataset in the comparison. From the BugHunter dataset, we used the average F-measure value of the previously selected algorithm calculated on the results obtained after applying the Subtract filtering. We performed this process for method, class, file, and projected levels separately.

On the traditional dataset, the machine learning algorithms performed better, achieving higher F-measure values in every case. The two kinds of datasets differ fundamentally, because they are constructed with two different methods. For the traditional dataset, we divided the history of the projects into six-month long intervals by selecting release versions from the GitHub repository. We collected the reported bugs from these intervals and we assigned the buggy source code elements to these release versions based on the bugfixes. Then we used the state of the source code elements from these selected versions to assemble the bug dataset. This method was used in several previous studies (D'Ambros et al., 2010; Gyimothy et al., 2005). It is necessary, because the bugs are reported usually after releasing a version, thus at the time of the release there are too few bugs to construct a bug dataset. For the bug assignment, we used a heuristic method – similar to other studies (Gyimothy et al., 2005) – where we assigned each bug to the latest selected version before the bug was reported into the issue tracking system. This method leads to some uncertainty in the dataset, because it could happen that the bug is not yet present in the assigned version. Table 28 shows some characteristics of this uncertainty.

The second column is the average number of days elapsed between the date of the release and the date of the bugs reported. The maximum that could occur is 180 days, because we used intervals around 6 months long. We can see that these averages are quite high, the overall average is 85 days. The third column is the average number of commits contributed to the project between the release commit and the date of the bugs reported. These values vary for each project, because it depends on the developers' activity. For some projects (Elasticsearch, Hazelcast) it could mean thousands of modifications before the bug was reported. The more commits are performed, the higher the probability that the source code element became buggy after the release. The fourth column shows the average number of commits performed from the time when the bug was reported and when the fix was applied. These numbers are much smaller, which also demonstrates that bugs are fixed relatively fast. This fast corrective behavior causes before and after fix states to be less different for the BugHunter approach. Consequently, less difference in metric values makes building a precise prediction model more difficult.

The uncertainty is also visualized in Figs. 6 and 7 for the sake of better comprehension. The first timeline (traditional dataset) shows a case when the actual bug occurred after Release #1, then the bug was reported and finally fixed before Release #2. In this case, the source code elements touched by the bug are included in the dataset with a buggy label. This dataset is created with the state captured at the time of Release #1. However, these source code elements became buggy after Release #1, thus the dataset marked them buggy at that point incorrectly. This error comes from the methodology itself which could be leveraged by narrowing the time window (which is 6 months wide traditionally). On the other hand, 6 months interval is not an unwitting choice. If we narrow down the time window, we will also have less bugs for an interval, which results in a more unbiased dataset, thus only less powerful predictive models can be built. It would be important to see how many source code elements were marked as buggy incorrectly, however this cannot be easily measured since the exact time of the bug occurrence cannot be determined (we only know the time when a bug was reported).

The new BugHunter approach (see Fig. 7), however, is free from the uncertainty mentioned above because it uses only the buggy and the fully fixed states of the bug related source code elements. This way, the produced bug dataset is more precise, hence it is more appropriate for machine learning. Therefore, we cannot clearly state that the traditional dataset is better, even

Table 29

Significance test results for method level – Algorithms.

	NaiveBayes	NaiveBayes-Multinomial	Logistic	SGD	Simple-Logistic	Voted-Perceptron	Decision-Table	OneR	J48	Random-Forest
NaiveBayesMultinomial	0.900 (1.28)									
Logistic	0.129 (4.07)	0.007 (5.35)								
SGD	0.045 (4.60)	0.002 (5.88)	0.900 (0.53)							
SimpleLogistic	0.053 (4.52)	0.002 (5.80)	0.900 (0.45)	0.900 (0.08)						
VotedPerceptron	0.900 (1.88)	0.900 (0.60)	0.001 (5.95)	0.001 (6.48)	0.001 (6.41)					
DecisionTable	0.045 (4.60)	0.002 (5.88)	0.900 (0.53)	0.900 (0.00)	0.900 (0.08)	0.001 (6.48)				
OneR	0.302 (3.54)	0.027 (4.82)	0.900 (0.53)	0.900 (1.06)	0.900 (0.98)	0.006 (5.43)	0.900 (1.06)			
J48	0.023 (4.90)	0.001 (6.18)	0.900 (0.83)	0.900 (0.30)	0.900 (0.38)	0.001 (6.78)	0.900 (0.30)	0.900 (1.36)		
RandomForest	0.001 (7.61)	0.001 (8.89)	0.302 (3.54)	0.546 (3.02)	0.513 (3.09)	0.001 (9.50)	0.546 (3.02)	0.129 (4.07)	0.679 (2.71)	
RandomTree	0.019 (4.97)	0.001 (6.23)	0.900 (0.90)	0.900 (0.38)	0.900 (0.45)	0.001 (6.86)	0.900 (0.38)	0.900 (1.43)	0.900 (0.08)	0.712 (2.64)

Table 30

Significance test results for class level – Algorithms.

	NaiveBayes	NaiveBayes-Multinomial	Logistic	SGD	SimpleLo-gistic	VotedPer-cep-tron	Deci-sionTable	OneR	J48	Random-Forest
NaiveBayesMultinomial	0.900 (1.13)									
Logistic	0.169 (3.92)	0.646 (2.79)								
SGD	0.169 (3.92)	0.646 (2.79)	0.900 (0.00)							
SimpleLogistic	0.019 (4.97)	0.191 (3.84)	0.900 (1.06)	0.900 (1.06)						
VotedPerceptron	0.900 (1.88)	0.546 (3.02)	0.002 (5.80)	0.002 (5.80)	0.001 (6.86)					
DecisionTable	0.004 (5.58)	0.063 (4.45)	0.900 (1.66)	0.900 (1.66)	0.900 (0.60)	0.001 (7.46)				
OneR	0.878 (2.26)	0.900 (1.13)	0.900 (1.66)	0.900 (1.66)	0.679 (2.71)	0.113 (4.15)	0.406 (3.32)			
J48	0.369 (3.39)	0.878 (2.26)	0.900 (0.53)	0.900 (0.53)	0.900 (1.58)	0.009 (5.28)	0.900 (2.19)	0.900 (1.13)		
RandomForest	0.900 (0.75)	0.900 (0.38)	0.479 (3.17)	0.479 (3.17)	0.098 (4.22)	0.712 (2.64)	0.027 (4.82)	0.900 (1.51)	0.712 (2.64)	
RandomTree	0.900 (1.66)	0.900 (0.53)	0.878 (2.26)	0.878 (2.26)	0.406 (3.32)	0.302 (3.54)	0.169 (3.92)	0.900 (0.60)	0.900 (1.73)	0.900 (0.90)

Table 31

Significance test results for projected – Algorithms.

	NaiveBayes	NaiveBayes-Multinomial	Logistic	SGD	SimpleLo-gistic	VotedPer-cep-tron	Deci-sionTable	OneR	J48	Random-Forest
NaiveBayesMultinomial	0.900 (0.68)									
Logistic	0.003 (5.73)	0.001 (6.41)								
SGD	0.001 (7.01)	0.001 (7.69)	0.900 (1.28)							
SimpleLogistic	0.006 (5.43)	0.001 (6.11)	0.900 (0.30)	0.900 (1.58)						
VotedPerceptron	0.900 (1.81)	0.778 (2.49)	0.169 (3.92)	0.011 (5.20)	0.271 (3.62)					
DecisionTable	0.001 (7.69)	0.001 (8.37)	0.900 (1.96)	0.900 (0.68)	0.878 (2.26)	0.002 (5.88)				
OneR	0.005 (5.50)	0.001 (6.18)	0.900 (0.23)	0.900 (1.51)	0.900 (0.08)	0.242 (3.69)	0.900 (2.19)			
J48	0.001 (5.95)	0.001 (6.63)	0.900 (0.23)	0.900 (1.06)	0.900 (0.53)	0.113 (4.15)	0.900 (1.73)	0.900 (0.45)		
RandomForest	0.001 (7.46)	0.001 (8.14)	0.900 (1.73)	0.900 (0.45)	0.900 (2.04)	0.003 (5.65)	0.900 (0.23)	0.900 (1.96)	0.900 (1.51)	
RandomTree	0.005 (5.50)	0.001 (6.18)	0.900 (0.23)	0.900 (1.51)	0.900 (0.08)	0.242 (3.70)	0.900 (2.19)	0.900 (0.00)	0.900 (0.45)	0.900 (1.96)

Table 32

Significance test results for file level – Algorithms.

	NaiveBayes	NaiveBayes-Multinomial	Logistic	SGD	SimpleLo-gistic	VotedPer-cep-tron	Deci-sionTable	OneR	J48	Random-Forest
NaiveBayesMultinomial	0.900 (1.88)									
Logistic	0.129 (4.07)	0.900 (2.19)								
SGD	0.900 (0.15)	0.900 (2.04)	0.098 (4.22)							
SimpleLogistic	0.148 (4.00)	0.900 (2.11)	0.900 (0.08)	0.113 (4.15)						
VotedPerceptron	0.900 (1.21)	0.513 (3.09)	0.009 (5.28)	0.900 (1.06)	0.011 (5.20)					
DecisionTable	0.443 (3.24)	0.900 (1.36)	0.900 (0.83)	0.369 (3.39)	0.900 (0.75)	0.063 (4.45)				
OneR	0.148 (4.00)	0.900 (2.11)	0.900 (0.08)	0.113 (4.15)	0.900 (0.00)	0.011 (5.20)	0.900 (0.75)			
J48	0.513 (3.09)	0.900 (1.21)	0.900 (0.98)	0.443 (3.24)	0.900 (0.90)	0.085 (4.30)	0.900 (0.15)	0.900 (0.90)		
RandomForest	0.098 (4.22)	0.845 (2.34)	0.900 (0.15)	0.073 (4.37)	0.900 (0.23)	0.006 (5.43)	0.900 (0.98)	0.900 (0.23)	0.900 (1.13)	
RandomTree	0.098 (4.22)	0.845 (2.34)	0.900 (0.15)	0.073 (4.37)	0.900 (0.23)	0.006 (5.43)	0.900 (0.98)	0.900 (0.23)	0.900 (1.13)	0.900 (0.00)

despite the higher F-measure values. The difference between the values of the two dataset is around 0.10 at method level, 0.21 at class level, and 0.09 at file level. Projecting method level metrics

to class level achieved almost as high an F-measure value (0.7405) as in the traditional case (0.7831). The difference is only 0.04, yet it is on a much more precise dataset.

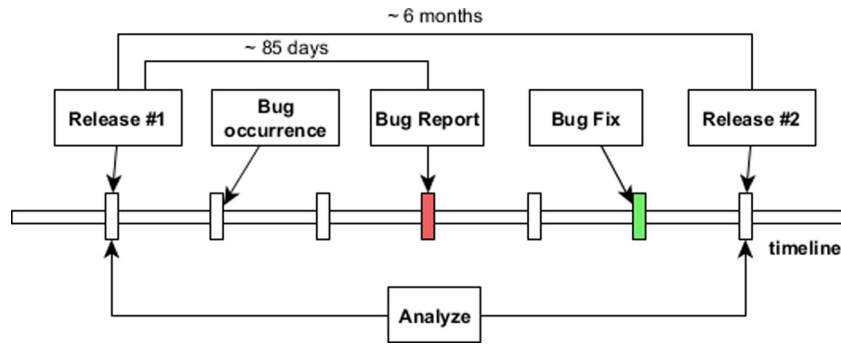


Fig. 6. Traditional approach.

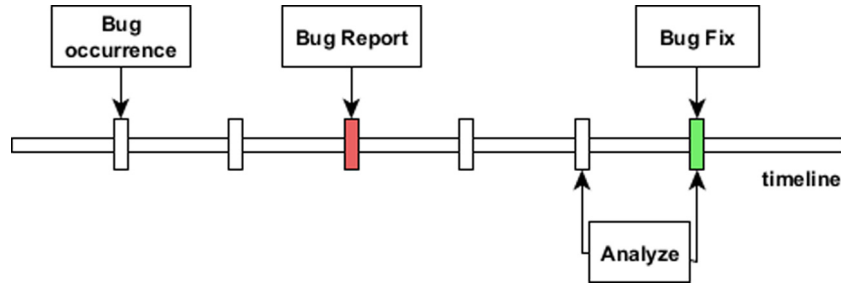


Fig. 7. BugHunter approach.

Answering RQ3: Traditional datasets include a high risk when labeling source code elements as buggy since the elements may become buggy after the release version. This injects false labeling into the training set, which might end up in deceptive machine learning results (as successfully predicting a bad label is not correct). Unfortunately, the number of incorrectly labeled source code elements cannot be determined since we only know the time when a bug was reported, we do not know the exact time when it was inserted into the system. These facts make it really hard to take one dataset and state that it is better for bug prediction.

but we do not know the issue at that time, which can also distort the measurements. Unfortunately, these uncertainties cannot be solved, since there are no further data to rely on.

7.2. Threats to internal validity

It would be meaningful to use multiple static source code analyzers in order to decrease the threats to internal validity caused by measuring source code element characteristics with only one tool. However, it would mean much more work, and even then, additional manual validations would be needed to decide which tool measures a given metric more precisely, which often depends on interpreting the conceptual definitions.

7.3. Threats to external validity

Currently, the constructed dataset consists of 15 projects which may limit the capabilities of the bug prediction models. Selecting more projects to be included in the dataset would increase the generalizability of the built bug prediction models. Considering additional source hosting platforms (SourceForge, Bitbucket, GitLab) would also increase the external validity of the dataset.

Widely used and accepted programming constructs and structures can vary from programming language to programming language. Using different constructs may have a significant result on the calculated metric values. Selecting projects written in different programming languages, not only Java software systems, could further strengthen the generalizability of our method.

8. Conclusions and future work

In this study, we developed a method that generates a bug dataset whose entries are source code elements touched by bug-fixes mined from GitHub. The entries represent before and after states of source code elements on which bug fixes were applied. The presented approach allows the simultaneous processing of several publicly available projects located on GitHub, thereby

7. Threats to validity

In this section, we briefly describe the threats to validity. Firstly, we present the construct validity, then the internal and external validity threats.

7.1. Threats to construct validity

When constructing a dataset in an automatic way, there are always some possible threats to validity. We validated our matching algorithm on JUnit, which was fair in size. However, investigating the validity of the matching in other systems could have revealed additional findings.

As we have seen, commit mismatches can occur during this process, which can distort the final bug dataset. However, manually validating all bugs and the corresponding commits would have been an enormous task.

Deciding which source code elements are faulty and which are not can also cause a construct validity threat. We consider a source code element faulty before the corresponding bug is fixed (the source code element had to be modified in that fix) and after the corresponding bug report is present. The source code element can already be faulty before the report and can have multiple changes in that period; and it can be faulty also after the last fix,

resulting in the production of a large – and automatically expandable – dataset. In contrast, previous studies have dealt with only a few large-scale datasets, which were created under strict individual management. Additionally, our dataset contains new source code metrics compared to other datasets, allowing the examination of the relationship between these metrics and software bugs. Furthermore, manual examinations showed the reliability of our approach, so the adaptation of project-specific labels to the presence of bugs remains the only non-automatic step. Our initial adaptation of 15 suitable Java projects lead to the construction of the current dataset, which is one of our main – publicly available – contributions.

During empirical evaluations, we showed that the dataset can be used for further investigations such as bug prediction. For this purpose, we used several machine learning algorithms at three different granularity levels (method, class, file) from which the method-level prediction achieved the highest F-measure values. As a novel kind of experiment, we also investigated whether the method-level metrics projected to the class level are better predictors than the class-level metrics themselves, and found a significant improvement in the results.

As potential future work, we are planning to expand the dataset with additional projects and even additional data sources, such as SourceForge and Bitbucket. Supporting different external bug tracking systems is another option for extending our approach. We will also dedicate more attention to the concrete prediction models we generate, as this study focused solely on showing the conceptual usability of our dataset.

CRedit authorship contribution statement

Rudolf Ferenc: Supervision, Conceptualization, Methodology, Writing – review & editing. **Péter Gyimesi:** Methodology, Software, Formal analysis, Investigation, Writing – original draft. **Gábor Gyimesi:** Software, Investigation. **Zoltán Tóth:** Validation, Writing – original draft. **Tibor Gyimóthy:** Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was supported by the EU-supported Hungarian national grant GINOP-2.3.2-15-2016-00037 titled “Internet of Living Things” and by grant TUDFO/47138-1/2019-ITM of the Ministry for Innovation and Technology, Hungary.

Appendix A

See Tables 29–32.

References

- Bakota, T., Hegedus, P., Kortvelyesi, P., Ferenc, R., Gyimothy, T., 2011. A probabilistic software quality model. In: Software Maintenance (ICSM), 2011 27th IEEE International Conference on. pp. 243–252. <http://dx.doi.org/10.1109/ICSM.2011.6080791>.
- Bán, D., Ferenc, R., 2014. Recognizing antipatterns and analyzing their effects on software maintainability. In: International Conference on Computational Science and Its Applications. Springer, pp. 337–352.
- Bán, D., Ferenc, R., Siket, I., Kiss, Á., Gyimóthy, T., 2018. Prediction models for performance, power, and energy efficiency of software executed on heterogeneous hardware. *J. Supercomput.* 1–25.
- Bangcharoensap, P., Ihara, A., Kamei, Y., Matsumoto, K., 2012. Locating source code to be fixed based on initial bug reports – A case study on the eclipse project. In: Empirical Software Engineering in Practice (IWESEP), 2012 Fourth International Workshop on. pp. 10–15. <http://dx.doi.org/10.1109/IWESEP.2012.14>.
- Basili, V.R., Briand, L.C., Melo, W.L., 1996. A validation of object-oriented design metrics as quality indicators. *IEEE Trans. Softw. Eng.* 22 (10), 751–761.
- Bellon, S., Koschke, R., Antoniol, G., Krinke, J., Merlo, E., 2007. Comparison and evaluation of clone detection tools. *IEEE Trans. Softw. Eng.* 33 (9), 577–591.
- Bird, C., Rigby, P., Barr, E., Hamilton, D., German, D., Devanbu, P., 2009. The promises and perils of mining git. In: Mining Software Repositories, 2009. MSR '09. 6th IEEE International Working Conference on. pp. 1–10. <http://dx.doi.org/10.1109/MSR.2009.5069475>.
- Boehm, B., Abts, C., Chulani, S., 2000. Software development cost estimation approaches – A survey. *Ann. Softw. Eng.* 10 (1), 177–205. <http://dx.doi.org/10.1023/A:1018991717352>.
- Booger, C., Moonen, L., 2008. Assessing the value of coding standards: An empirical study. In: Software Maintenance, 2008. ICSM 2008. IEEE International Conference on. IEEE, pp. 277–286.
- Bowes, D., Hall, T., Harman, M., Jia, Y., Sarro, F., Wu, F., 2016. Mutation-aware fault prediction. In: Proceedings of the 25th International Symposium on Software Testing and Analysis. ACM, pp. 330–341.
- Bruntink, M., Van Deursen, A., 2004. Predicting class testability using object-oriented metrics. In: Source Code Analysis and Manipulation, 2004. Fourth IEEE International Workshop on. IEEE, pp. 136–145.
- Catal, C., Diri, B., 2009. A systematic review of software fault prediction studies. *Expert Syst. Appl.* 36 (4), 7346–7354.
- Chidamber, S.R., Kemerer, C.F., 1994. A metrics suite for object oriented design. *IEEE Trans. Softw. Eng.* 20 (6), 476–493.
- Cohen, J., 1992. A power primer. *Psychol. Bull.* 112 (1), 155.
- Coles, H., 2018. Pitest mutation framework. <http://pitest.org/>, accessed: 2018-01-25.
- Couto, C., Silva, C., Valente, M., Bigonha, R., Anquetil, N., 2012. Uncovering causal relationships between software metrics and bugs. In: Software Maintenance and Reengineering (CSMR), 2012 16th European Conference on. pp. 223–232. <http://dx.doi.org/10.1109/CSMR.2012.31>.
- Dallmeier, V., Zimmermann, T., 2007a. Automatic Extraction of Bug Localization Benchmarks from History. Tech. rep., Universität des Saarlandes and Saarbrücken and Germany.
- Dallmeier, V., Zimmermann, T., 2007b. Extraction of bug localization benchmarks from history. In: Proceedings of the Twenty-Second IEEE/ACM International Conference on Automated Software Engineering. ACM, pp. 433–436.
- D'Ambros, M., Lanza, M., Robbes, R., 2010. An extensive comparison of bug prediction approaches. In: Proceedings of MSR 2010 (7th IEEE Working Conference on Mining Software Repositories). pp. 31–41, URL <http://bug.inf.usi.ch/>.
- Davies, S., Roper, M., Wood, M., 2012. Using bug report similarity to enhance bug localisation. In: Reverse Engineering (WCRE), 2012 19th Working Conference on. IEEE, pp. 125–134.
- Dyer, R., Nguyen, H.A., Rajan, H., Nguyen, T.N., 2013. Boa: A language and infrastructure for analyzing ultra-large-scale software repositories. In: Proceedings of the 2013 International Conference on Software Engineering. IEEE Press, pp. 422–431.
- Fischer, M., Pinzger, M., Gall, H., 2003. Populating a release history database from version control and bug tracking systems. In: Software Maintenance, 2003. ICSM 2003. Proceedings. International Conference on. IEEE, pp. 23–32.
- Friedman, M., 1940. A comparison of alternative tests of significance for the problem of m rankings. *Ann. Math. Stat.* 11 (1), 86–92.
- Ghotra, B., McIntosh, S., Hassan, A.E., 2015. Revisiting the impact of classification techniques on the performance of defect prediction models. In: Proceedings of the 37th International Conference on Software Engineering-Volume 1. IEEE Press, pp. 789–800.
- Gyimesi, P., Gyimesi, G., Tóth, Z., Ferenc, R., 2015. Characterization of source code defects by data mining conducted on github. In: International Conference on Computational Science and Its Applications. Springer, pp. 47–62.
- Gyimothy, T., Ferenc, R., Siket, I., 2005. Empirical validation of object-oriented metrics on open source software for fault prediction. *IEEE Trans. Softw. Eng.* 31 (10), 897–910.
- Hall, T., Zhang, M., Bowes, D., Sun, Y., 2014. Some code smells have a significant but small effect on faults. *ACM Trans. Softw. Eng. Methodol. (TOSEM)* 23 (4), 33.
- Harman, M., Islam, S., Jia, Y., Minku, L.L., Sarro, F., Srivisut, K., 2014. Less is more: Temporal fault predictive performance over multiple hadoop releases. In: International Symposium on Search Based Software Engineering. Springer, pp. 240–246.
- He, H., Garcia, E., et al., 2009. Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.* 21 (9), 1263–1284.
- Herbold, S., Trautsch, A., Grabowski, J., 2018. A comparative study to benchmark cross-project defect prediction approaches. *IEEE Trans. Softw. Eng.* 44 (9), 811–833. <http://dx.doi.org/10.1109/TSE.2017.2724538>.

2018. Jhawk. <http://www.virtualmachinery.com/jhawkprod.htm>, Accessed: 2018-01-25.
- Johnson, B., Song, Y., Murphy-Hill, E., Bowdidge, R., 2013. Why don't software developers use static analysis tools to find bugs? In: 2013 35th International Conference on Software Engineering (ICSE). IEEE, pp. 672–681.
- Jureczko, M., Madeyski, L., 2010. Towards identifying software project clusters with regard to defect prediction. In: Proceedings of the 6th International Conference on Predictive Models in Software Engineering. ACM, p. 9.
- Kalliamvakou, E., Gousios, G., Blincoe, K., Singer, L., German, D.M., Damian, D., 2014. The promises and perils of mining github. In: MSR 2014 Proceedings of the 11th Working Conference on Mining Software Repositories. pp. 92–101, URL <http://www.gousios.gr/pub/promises-perils-github.pdf>.
- Kim, S., Zimmermann, T., Pan, K., Whitehead Jr, E.J., 2006. Automatic identification of bug-introducing changes. In: Automated Software Engineering, 2006. ASE'06. 21st IEEE/ACM International Conference on. IEEE, pp. 81–90.
- Kohavi, R., 1995. The power of decision tables. In: 8th European Conference on Machine Learning. Springer, pp. 174–189.
- Landwehr, N., Hall, M., Frank, E., 2005. Logistic model trees. *Mach. Learn.* 95 (1–2), 161–205.
- Li, Z., Jing, X., Zhu, X., 2018. Progress on approaches to software defect prediction. *IET Softw.* 12 (3), 161–175. <http://dx.doi.org/10.1049/iet-sen.2017.0148>.
- Li, Z., Jing, X., Zhu, X., Zhang, H., Xu, B., Ying, S., 2019. On the multiple sources and privacy preservation issues for heterogeneous defect prediction. *IEEE Transactions on Software Engineering* 45 (4), 391–411.
- Ma, Y., Guo, L., Cukic, B., 2006. A statistical framework for the prediction of fault-proneness. In: Advances in Machine Learning Application in Software Engineering. Idea Group Inc, pp. 237–265.
- Mockus, A., Votta, L.G., 2000. Identifying reasons for software changes using historic databases. In: Software Maintenance, 2000. Proceedings. International Conference on. IEEE, pp. 120–130.
- Nemenyi, P., 1963. Distribution-Free Multiple Comparison (PhD dissertation). Princeton University.
- Ostrand, T.J., Weyuker, E.J., Bell, R.M., 2005. Predicting the location and number of faults in large software systems. *IEEE Trans. Softw. Eng.* 31 (4), 340–355.
- Ostrand, T.J., Weyuker, E.J., Bell, R.M., 2007. Automating algorithms for the identification of fault-prone files. In: Proceedings of the 2007 International Symposium on Software Testing and Analysis. ACM, pp. 219–227.
- Petrić, J., Bowes, D., Hall, T., Christianson, B., Baddoo, N., 2016. The jinx on the NASA software defect data sets. In: Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering. pp. 1–5.
- Porter, A., Selby, R.W., et al., 1990. Empirically guided software development using metric-based classification trees. *IEEE Softw.* 7 (2), 46–54.
- Quinlan, R., 1993. C4.5: Programs for Machine Learning. Morgan Kaufmann Publishers, San Mateo, CA.
- Rosenthal, R., Cooper, H., Hedges, L., 1994. Parametric measures of effect size. In: *The Handbook of Research Synthesis*, Vol. 621. pp. 231–244.
- Roy, C.K., Cordy, J.R., Koschke, R., 2009. Comparison and evaluation of code clone detection techniques and tools: A qualitative approach. *Sci. Comput. Program.* 74 (7), 470–495.
- Saha, R.K., Lease, M., Khurshid, S., Perry, D.E., 2013. Improving bug localization using structured information retrieval. In: Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on. IEEE, pp. 345–355.
- Shepperd, M., Song, Q., Sun, Z., Mair, C., 2013. Data quality: Some comments on the nasa software defect datasets. *IEEE Trans. Softw. Eng.* 39 (9), 1208–1215.
- Shippey, T., Hall, T., Counsell, S., Bowes, D., 2016. So you need more method level datasets for your software defect prediction?: Voilà! In: Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. In: ESEM '16, ACM, New York, NY, USA, pp. 12:1–12:6. <http://dx.doi.org/10.1145/2961111.2962620>, URL <http://doi.acm.org/10.1145/2961111.2962620>.
- Shirabad, J.S., Menzies, T.J., 2005. The PROMISE Repository of Software Engineering Databases, Vol. 24. School of Information Technology and Engineering, University of Ottawa, Canada, URL <http://promise.site.uottawa.ca/SERepository>.
- Śliwowski, J., Zimmermann, T., Zeller, A., 2005. When do changes induce fixes? In: ACM Sigsoft Software Engineering Notes, Vol. 30. ACM, pp. 1–5.
- Tóth, Z., Gyimesi, P., Ferenc, R., 2016. A public bug database of github projects and its application in bug prediction. In: International Conference on Computational Science and Its Applications. Springer, pp. 625–638.
- Toth, Z., Novak, G., Ferenc, R., Siket, I., 2013. Using version control history to follow the changes of source code elements. In: Software Maintenance and Reengineering (CSMR).
- Wang, D., Lin, M., Zhang, H., Hu, H., 2010. Detect related bugs from source code using bug information. In: Computer Software and Applications Conference (COMPSAC). URL <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5676263&abstractAccess=no&userType=inst>.
- Wang, S., Lo, D., 2014. Version history, similar report, and structure: Putting them together for improved bug localization. In: Proceedings of the 22nd International Conference on Program Comprehension. ACM, pp. 53–63.
- Wang, S., Yao, X., 2013. Using class imbalance learning for software defect prediction. *IEEE Trans. Reliab.* 62 (2), 434–443.
- Wilcoxon, F., 1945. Individual comparisons by ranking methods. *Biom. Bull.* 1 (6), 80–83.
- Williams, C., Spacco, J., 2008. Szz revisited: verifying when changes induce fixes. In: Proceedings of the 2008 Workshop on Defects in Large Software Systems. ACM, pp. 32–36.
- Wu, R., Zhang, H., Kim, S., Cheung, S.-C., 2011. Relink: recovering links between bugs and changes. In: Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering. ACM, pp. 15–25.
- Zhou, Y., Leung, H., 2006. Empirical analysis of object-oriented design metrics for predicting high and low severity faults. *IEEE Trans. Softw. Eng.* 32 (10), 771–789.
- Zhou, J., Zhang, H., Lo, D., 2012. Where should the bugs be fixed? more accurate information retrieval-based bug localization based on bug reports. In: 2012 34th International Conference on Software Engineering (ICSE). URL <http://www.mysmu.edu/faculty/davidlo/papers/icse12-localization.pdf>.
- Zimmermann, T., Premraj, R., Zeller, A., 2007. Predicting defects for eclipse. In: Predictor Models in Software Engineering, 2007. PROMISE'07: ICSE Workshops 2007. International Workshop on. IEEE, p. 9.

Dr. Habil. Rudolf FERENC is the head of the Department of Software Engineering and an associate professor at the University of Szeged, Hungary. His research interests include static code analysis, metrics, quality assurance, design pattern and antipattern mining, and bug detection. He has more than 100 publications in these fields with over 2000 citations. He is leading several R&D projects, which are related to quality assessment, improvement and architecture reconstruction of software systems for major banks and software development companies. He has been serving as Program Co-Chair and Program Committee member at the major conferences in this field (ICSE, ICSME, ESEC/FSE, SANER, CSMR, WCRE, ICPC, SCAM, FASE, etc.) since 2005.

Péter Gyimesi is a Ph.D. student at the University of Szeged, who is currently working on his Ph.D. thesis. He received his MSc degree in Computer Science from the University of Szeged in 2015. His research interests include software quality assurance, software metrics and bug prediction. He has participated in several R&D projects related to software quality assurance.

Gábor Gyimesi is software developer associate at the University of Szeged. He received his MSc degree in Computer Science from the University of Szeged in 2015.

Zoltán Tóth is assistant lecturer at the University of Szeged. He is a Ph.D. candidate, who has submitted his thesis to the doctoral committee for review. He received his MSc degree in Computer Science from the University of Szeged in 2014. His main research interests include static code analysis, software metrics, bug prediction, quality assurance. He has participated in several R&D projects related to software quality assurance. He took part in the development of static source code analyzers for RPG and JavaScript software systems.

Tibor Gyimóthy is professor at the Department of Software Engineering at the University of Szeged in Hungary and the head of the MTA-SZTE Research Group on Artificial Intelligence. His research interests are concentrated on high quality software development and maintenance. His main contributions are methods and tools for identifying and predicting harmful program elements in the source code. He published over 210 research papers with more than 2600 independent citations. His scientific contributions were awarded with the Széchenyi Professor Grant, László Kalmár Award, HAS Academic Prize, Gábor Dénes Award and Szent-Györgyi Albert Prize. Tibor Gyimóthy is four-time program committee member of the International Conference on Software Engineering (ICSE), which is the most significant software engineering conference in the world. In 2011, Tibor Gyimóthy was elected the Conference Chair of the European Software Engineering Conference (ESEC)/ ACM Foundations of Software Engineering (FSE).