



# Automating orthogonal defect classification using machine learning algorithms



Fábio Lopes<sup>a</sup>, João Agnelo<sup>a</sup>, César A. Teixeira<sup>a</sup>, Nuno Laranjeiro<sup>a,\*</sup>, Jorge Bernardino<sup>b,a</sup>

<sup>a</sup> CISUC, Department of Informatics Engineering, University of Coimbra, Portugal

<sup>b</sup> Polytechnic of Coimbra - ISEC, Portugal

## ARTICLE INFO

### Article history:

Received 27 March 2019

Received in revised form 1 August 2019

Accepted 6 September 2019

Available online 12 September 2019

### Keywords:

Software defects

Bug reports

Orthogonal defect classification

Machine learning

Text classification

## ABSTRACT

Software systems are increasingly being used in business or mission critical scenarios, where the presence of certain types of software defects, i.e., bugs, may result in catastrophic consequences (e.g., financial losses or even the loss of human lives). To deploy systems in which we can rely on, it is vital to understand the types of defects that tend to affect such systems. This allows developers to take proper action, such as adapting the development process or redirecting testing efforts (e.g., using a certain set of testing techniques, or focusing on certain parts of the system). Orthogonal Defect Classification (ODC) has emerged as a popular method for classifying software defects, but it requires one or more experts to categorize each defect in a quite complex and time-consuming process. In this paper, we evaluate the use of machine learning algorithms (k-Nearest Neighbors, Support Vector Machines, Naïve Bayes, Nearest Centroid, Random Forest and Recurrent Neural Networks) for automatic classification of software defects using ODC, based on unstructured textual bug reports. Experimental results reveal the difficulties in automatically classifying certain ODC attributes solely using reports, but also suggest that the overall classification accuracy may be improved in most of the cases, if larger datasets are used.

© 2019 Elsevier B.V. All rights reserved.

## 1. Introduction

Software systems are nowadays at the core of most businesses, where they are used to support critical operations, usually under the form of services [1]. In general, the size and complexity of critical software systems has grown throughout the years and the tendency is that this complexity and their importance will continue to increase. In this context, the presence of a software defect (i.e., a bug) may result in serious failures [2] that have several consequences on the business (e.g., lost business transactions, reputation losses, or even severe safety problems).

The development of dependable systems requires that developers carry out software defect detection activities (e.g., code inspections, static analysis, black-box testing) where it is vital to detect bugs early in the process, especially before production, otherwise the impact on product cost can become quite high [3]. However, from a software engineering perspective, what really makes the whole process more efficient is the execution of preventive actions. As it is not possible to have preventive actions for all different bugs, it is better to understand their main types and

then perform root cause analysis on the patterns found [4]. This allows for adjustments in the software engineering methodology, allowing to put more or less effort on certain software defect detection techniques, changing the development methodology (e.g., to use test-driven development), or redistributing team effort, just to name a few.

Orthogonal Defect Classification (ODC) is a set of analytical methods used for software development and test process analysis to characterize software defects that consists of eight orthogonal attributes [4]. Despite being a very popular framework for classifying software defects, the main problem associated with it is that it requires the presence of one or more experts, the classification process is relatively complex and particularly lengthy, and does not fit well with current time-to-market pressure put on developers. Previous works have tried to automate the classification process (e.g., [5–7]), but most of them are limited to one or two ODC attributes (e.g., Defect Type, [6,7], Impact [5,8]), require the addition of further information to bug reports, or end up using a simpler classification, which is placed on top of ODC, and where higher level families serve to group ODC attributes, which impacts the usefulness of ODC.

In this paper, we evaluate the applicability of a set of popular machine learning algorithms for performing automatic classification of software defects purely based on ODC and using unstructured text bug reports as input. The classifiers considered were:

\* Corresponding author.

E-mail addresses: [fadcl@student.dei.uc.pt](mailto:fadcl@student.dei.uc.pt) (F. Lopes), [jagnelo@student.dei.uc.pt](mailto:jagnelo@student.dei.uc.pt) (J. Agnelo), [cteixe@dei.uc.pt](mailto:cteixe@dei.uc.pt) (C.A. Teixeira), [cnl@dei.uc.pt](mailto:cnl@dei.uc.pt) (N. Laranjeiro), [jorge@isec.pt](mailto:jorge@isec.pt) (J. Bernardino).

k-Nearest Neighbors (k-NN), Support Vector Machines (SVM), Naïve Bayes (NB), Nearest Centroid (NC), Random Forest (RF), and Recurrent Neural Networks (RNNs). We illustrate our approach using a set of 4096 bug reports referring to ‘closed and resolved’ bugs that were extracted from the issue tracking systems of three popular NoSQL systems: MongoDB, Cassandra, and HBase. At the time of writing, this number of bugs represents about one fourth of all ‘closed and resolved’ bugs for these systems and were previously manually classified (and partially verified) by a researcher according to ODC. Two additional independent classifications were also produced to understand the reliability of the dataset.

The experimental results show the difficulties in automatically classifying bugs by solely using bug reports and disregarding the actual changes in the code. This is evident in some ODC attributes, such as Defect Type and Qualifier, which are highly related to the code corrections associated with the bug report. On the other hand, other attributes (e.g., Target) show high accuracy values as they do not depend on the actions taken by the developer to correct the bug. It also became obvious that, with exceptions, a relatively small dataset (which typically carries low number of occurrences of certain ODC classes) negatively impacts the overall results. Results suggest that, in general, larger and richer datasets may be very helpful in effectively automating the whole classification procedure.

We must emphasize that we are making our dataset publicly available and are also providing a Representational state transfer (RESTful) service [9] for performing automatic classification of bugs according to ODC. The service receives, as input, an unstructured textual description of the software defect (i.e., a bug report) and the machine learning algorithm to use. The dataset, Swagger API description, and RESTful service are available at <http://odc.dei.uc.pt>.

In summary, the main contributions of this paper are the following:

- A comprehensive evaluation of machine learning algorithms for automatic classification of bug reports using ODC;
- A machine learning system able to tackle the majority of ODC attributes, namely Activity, Trigger, Impact, Target, Defect Type, and Qualifier;
- The public availability of our dataset holding 4096 ODC annotated bug reports;
- A RESTful service allowing developers to effortlessly classify a bug report according to ODC.

The rest of this paper is organized as follows. Section 2 presents background concepts and related work on papers using ODC, with emphasis on research trying to automate software defect classification using machine learning algorithms. Section 3 presents the study designed to evaluate the applicability of machine learning algorithms to ODC and Section 4 discusses the results. Section 5 presents the threats to the validity of this work and finally Section 6 concludes this paper.

## 2. Background and related work

This section introduces the Orthogonal Defect Classification (ODC) [10] main concepts and overviews works related with the application of ODC. The section then presents a set of works that try to automate the classification of bug reports, particularly by using machine learning algorithms.

### 2.1. Orthogonal defect classification

Orthogonal Defect Classification is an analytical process used in software development and test process analysis to characterize software defects [10]. It allows extracting valuable information from defects and provides insights and aids diagnosis in software engineering processes. ODC is based on the definition of eight attributes grouped into two sections: open-report and closed-report. The *open-report* section refers to attributes for which there should be information at the moment of disclosure of a certain defect, in particular:

- Activity: refers to the activity being performed at the time the defect was disclosed (e.g., unit testing);
- Trigger: indicates what caused the defect to manifest (e.g., a blocked test);
- Impact: is the impact a user experienced when the defect surfaced.

The *closed-report* section includes the attributes that relate with the correction of a certain defect (and depend on that information). In particular:

- Target: the object that was the target of the correction (e.g., a set of build scripts);
- Defect Type: the type of change that was performed to correct the problem (e.g., changing a checking condition);
- Qualifier: describes the code point before the correction was made (e.g., missing, incorrect, or extraneous);
- Age: the moment in which the defect was introduced (e.g., during the correction of another defect);
- Source: refers to the origin of the defect, e.g., if it was detected in outsourced code or in-house developed code.

Several works have used ODC to classify sets of software defects to understand the distribution of the types of bugs found. Some works use the knowledge obtained from using ODC in empirical settings (e.g., fault injection experiments), such as [11] and [12]. Christmansson and Chillarege [11] inject errors instead of faults (i.e., an error is the consequence of a fault [13]) so as to skip the unknown waiting period needed for a certain fault to be activated. This requires some form of mapping between the set of faults that should be injected and their respective errors, which is obtained via using ODC classified bugs. Durães and Madeira [12] use the ODC classification distributions from a set of software bugs to categorize common programmer mistakes. The emerging patterns among these common mistakes are then used to define a set of software fault emulation operators. Works such as [14] and [15] propose workflow strategies consisting of multiple steps which aim to improve the effectiveness of software development processes. In the midst of these steps, the authors include the use of ODC to classify bugs found in the software product being developed. The final steps are mostly focused on extracting knowledge and lessons learned from the whole process, mostly based on defect distributions according to ODC attributes.

### 2.2. Automating ODC with machine learning

A common problem present in works that use ODC is that the bug classification is usually a time-consuming manual process that requires the presence of one (or ideally more than one) expert. In addition, the number of defects analyzed should be relatively large to enable further analysis, which aggravates the problem. This kind of task is not only time- and resource-consuming, but also error prone. Therefore, there is an obvious need to automate this process, either via rule-based approaches

(e.g., as proposed in [16]) or through the use of machine learning algorithms.

In [6], the authors propose an SVM-based automatic defect classification scheme. This scheme classifies defects into three families that essentially aggregate ODC attributes: (i) control and data flow; (ii) structural; and (iii) non-functional. The classification process applies a text mining solution, which is able to extract useful information from defect report descriptions and source code fixes, and uses this information as the main criteria to classify defects. To validate the proposal, 500 software defects were classified both manually by the authors and automatically by the resulting SVM algorithm. The authors reported being able to achieve a global accuracy of 77.8%. The fact that families of ODC attributes are used is a way of diminishing the error, which may be interesting in certain scenarios, but, at the same time, providing an automatic approach based purely in ODC is a quite different endeavor, as the diversity of attributes (and respective values) is much larger.

The authors of [7] propose an active learning and semi-supervised defect classification approach that targets the families of defects proposed in [6]. The approach resorts to minimal labeled data that includes diverse and informative examples (i.e., active learning) and uses both labeled and unlabeled defect examples in the learning process (i.e., semi-supervised learning). The main goal of this approach is to reduce the amount of time spent in categorizing defects to train machine learning algorithms. The authors validated the proposal by labeling only 50 ODC defects out of a set of 500 (unlabeled), and achieved a weighted precision of 65.1% (applied to the three mentioned families only and based on the special selection of the minimal set of defects).

In [17] the authors describe an automatic defect classification scheme based on feeding the Abstract Syntax Trees of the buggy code and its respective correction to a decision engine. This decision engine can classify defects according to four categories: (i) Data; (ii) Computational; (ii) Interface; and (iv) Control/Logic. With a dataset of 1174 defects extracted from defect-tracking repositories of three separate systems, the automatic defect classification scheme was able to classify defects with an accuracy ranging from 64% to 79%. Given the relatively low number of classes, automating such procedure tends to be an easier task (i.e., when compared with a more complex scheme, such as ODC), but, at the same time, less informative.

In [5], the authors describe AutoODC, which is an approach for automatically classifying software defects according to ODC and that views the case as a supervised text classification problem. AutoODC integrates ODC experience with domain knowledge and resorts to two text classification algorithms, namely Naïve Bayes and Support Vector Machines. The authors introduce the Relevance Annotation Framework, in which a person registers the reason why they believe a given defect should belong to a particular ODC attribute value, thus providing further information to the machine learning algorithm. The authors state that this additional approach is able to reduce the relative error of the classification by a factor between 14% and 29%. The accuracy of the proposed approach reaches 82.9% using the NB classifier and 80.7% using the SVM classifier on a total of 1653 industrial defect reports, and 77.5% and 75.2% using the NB and SVM classifiers respectively, on a set of open-source defect reports. Note, however that, trying to automate ODC exclusively relying on the typical information present in a bug report is a much harder task and, at the same time, a more realistic scenario.

Some works use machine learning algorithms to automate the process of bug triage, i.e., the process of deciding what to do with newly found software defects. These approaches do not focus on the classification of the defects in terms of their causes

or corrective measures, but typically on properties like severity level, so that defects can be separated into specific categories, allowing developers to, for instance, prioritize their corrections. The works [18–21] apply this approach to automatically assign new bugs to specific developers. To this end, the authors of [18] propose the use of a supervised Bayesian machine learning algorithm. The algorithm correctly assigns 30% of bug reports to the expected developers. Xuan et al. [21] describe a semi-supervised NB text classification algorithm which uses 'expectation-maximization' to take advantage of both categorized and uncategorized bug reports. The algorithm iteratively categorizes unlabeled data (which refers to the expectation step) and rebuilds a new classifier (which refers to the maximization step). This semi-supervised approach increases the overall accuracy by up to 6% when compared with its supervised counterpart.

Tian et al. [19] propose a framework, named DRONE, for predicting the severity of a software defect based on a set of factors (i.e., temporal, textual, author, related-report, severity, and product) that are extracted from bug reports via feature extraction. The framework includes a classification engine named Thresholding and Linear Regression to Classify Imbalanced Data (GRAY), which tries to enhance linear regression with a thresholding approach so that cases of imbalanced bug reports are handled. The approach is applied against a dataset of over 100,000 bug reports from Eclipse and it shows an improvement of 58.61% compared to results in Menzies and Marcus [22]. With the same bug triage goal, Dedík [20] proposes an algorithm based on SVM with Term Frequency-Inverse Document Frequency (TF-IDF). The achieved accuracy is of 57% for a dataset of Firefox bug reports, 53% for a Netbeans dataset and 53% for a defect dataset from a private company.

Finally, Hernández-González et al. [8] present an approach which focuses on automating the classification of software defects according to the ODC Impact attribute. The authors name this approach "learning from crowds" (it relies on the amount of data rather than quality), and it focuses on training a classifier with a large amount of potentially erroneous and noisy classifications from people with no particular expertise in ODC classification (i.e., partially reliable annotators). The work used the data produced by a group of 5 novice annotators, which classified 846 plus 598 bugs from the Compendium and Mozilla Firefox open-source projects, respectively. The authors resort to a Bayesian network classifier and conclude that learning from (unskilled) crowds can outperform the typical approach of using a reliable dataset to train the classifier. Still, the authors have used only 4 of the 13 values of the Impact attribute (Installability, Requirements, Usability, Other), reaching the whole 13 values with satisfactory accuracy is a complex task.

In summary, considering the prior developments on automatic defect classification, we essentially find two groups of works. One group uses machine learning to perform bug triage (including assigning bugs to developers and distinguishing severity). The second group specializes on ODC, but the different works either *focus on a single ODC attribute* (sometimes using a reduced set of values for the single attribute), or they *focus in families of attributes*, where several attributes are used to form higher-level groups. From the perspective of ODC these kinds of approaches are less useful, as we are trading the information richness that is strictly associated with the detailed ODC process with the information that can be obtained from a higher granularity classification (which is less informative), although it potentially allows accuracy metrics to reach higher values.

In this work, we opted to not discard less popular attributes nor consider families of attributes. In this sense, this work diverges from related work and results cannot be directly compared using the typical metrics (e.g., accuracy). These two particular

choices create greater difficulties to the machine learning algorithms, such as: (i) there is no additional information for each defect, other than what already exists in reality; (ii) the number of attributes (and classes, i.e., attribute values) are higher, thus the chance of error tends to be higher. Despite the additional difficulties, we consider that this case represents reality in a more accurate way and allows developers and researchers to obtain full benefits of applying ODC.

### 3. Study design

In this section, we describe the design of our study, which is based on the following set of steps:

- (1) Manual creation and preprocessing of the dataset;
- (2) Feature extraction;
- (3) Dimensionality reduction;
- (4) Configuration and execution of the machine learning classifiers;
- (5) Assessment of the classifiers performance.

The following subsections describe each of the above steps in detail, explaining the decisions taken throughout the work.

#### 3.1. Dataset creation and preprocessing

The creation and initial preprocessing of the dataset (step 1 of the study) used in this work involved going through the following set of substeps, which we describe in the next paragraphs:

- i. Selection of a set of NoSQL databases;
- ii. Selection of a set of defect descriptions from each one of the databases online bug tracking platforms;
- iii. Manual ODC classification of each software defect, carried out by one researcher (named *researcher1*);
- iv. Manual verification of the software defects classified in the previous step in two different verification activities:
  - (a) Internal verification carried out by the same researcher responsible for the original classification;
  - (b) External verification produced by two external researchers (named *researcher2* and *researcher3*).

We started by selecting three popular NoSQL databases, MongoDB, Cassandra, and HBase (selected according to popularity rankings in [db-engines.com](http://db-engines.com), [stackoverflow.com](http://stackoverflow.com), and [kdnuggets.com](http://kdnuggets.com)). One ODC-knowledgeable early stage researcher, named *researcher1* was selected to carry out the manual process of classifying ‘closed and resolved’ bug reports, according to ODC. *researcher1* first trained with a total of 300 defect descriptions (100 bug reports per each of the 3 NoSQL databases). During this process, and initially in an informal way, a few samples of the outcome of the classification were verified and discussed by two experienced researchers as a way of providing feedback to *researcher1*. These 300 bug descriptions were then discarded and, thus, are not part of the final dataset.

We then randomly selected 4096 ‘closed and resolved’ bug reports (1618, 1095 and 1383 defects for MongoDB, Cassandra, and Hbase, respectively), which, at time of writing, accounted for around one fourth of the total closed and resolved defects in each database. We divided the 4096 software defects dataset in five batches, with the first four batches holding 2048 bugs (512 defects each batch). *researcher1* then initiated with the manual classification of each defect.

The bug reports used are composed of a title, a description of the detected defect and several comments that end up describing what has been made to correct the defect. This information was

interpreted by *researcher1* and each defect was manually labeled using the six different labels, one for each ODC attribute: Activity, Trigger, Impact, Target, Defect Type, and Qualifier. We excluded the ‘Age’ and ‘Source’ ODC attributes from this work, as the bug tracking systems of the three systems under analysis did not provide sufficient information to classify these attributes for all the defects in our dataset.

To verify the classification of the defects and, most of all, to have some insight regarding the overall quality of the annotated dataset, we performed two verification activities against 40% of the defects (1640 out of 4096 defects), which we name Internal and External and that are depicted in Fig. 1.

During the internal verification activity, *researcher1* progressively re-inspected a total of 20% of the defects (820 of the 4096 defects) to check and correct any errors found. Thus, the internal verification was performed over the first four batches and immediately after the classification of a certain batch was concluded. In practice, by the end of each batch, the researcher double-checked a total of 40% of the already classified defects in that batch, using the following distributions: (i) a sequential selection of the first 20% of the defects in the batch; (ii) a random selection of another 20% from the remaining defects in the batch. After this was concluded and the dataset improved, we initiated the external verification activity, which involved two new and independent classifications carried out by two external researchers (designated *researcher2* and *researcher3*). These two researchers were provided with two sets of randomly selected and non-overlapping defects, which added up to 20% of the global dataset (i.e., 820 defects in total, 410 bugs per each one of the two researchers).

Table 1 shows the detailed outcome of the verification procedure, which is composed of the previously mentioned verification tasks, named Internal (carried out by *researcher1*), External1 (carried out by *researcher2*) and External2 (carried out by *researcher3*). The table shows the details only for the Defect Type attribute, due to the fact that it is the most widely used ODC attribute in related work (e.g., [23–25]). Notice that the defect counts for Internal, External 1, and External 2 are below the 20% and 10% marks (the exact numbers are 746, 379, and 382 defects, respectively), as we are considering only the defects where the Target attribute was confirmed to be ‘Code’ and, as such, can be marked with a defect type.

In each of the matrices, each cell holds the total number of bugs marked with a certain attribute value, the values read in the rows represent the outcome of the verification procedure. The highlighted cells in the diagonal refer to true positives (the bugs in which there is agreement between the original classification and the respective verification task).

We also analyzed the inter-rater agreement using Cohen’s Kappa ( $k$ ), which is able to measure the agreement between two raters that classify items in mutually exclusive categories [26]. In our case, we wanted to understand the agreement of *researcher2* (and *researcher3*) with *researcher1*. The definition of  $k$  is given by Eq. (1):

$$k = \frac{p_o - p_c}{1 - p_c} \quad (1)$$

where  $p_o$  is the relative observed agreement between raters (i.e., accuracy) and  $p_c$  is the probability of agreement by chance. If both raters fully agree, then  $k = 1$ , if there is no agreement beyond what is expected by chance, then  $k = 0$ , and, finally, a negative value reflects the cases where agreement is actually worse than random choice. The following terms are typically used for the quantitative values of  $k$ : *poor* when  $k < 0$ , *slight* when  $0 \leq k \leq 0.2$ , *fair* when  $0.21 \leq k \leq 0.40$ , *moderate* when  $0.41 \leq k \leq 0.60$ , *substantial* when  $0.61 \leq k \leq 0.80$ , and finally

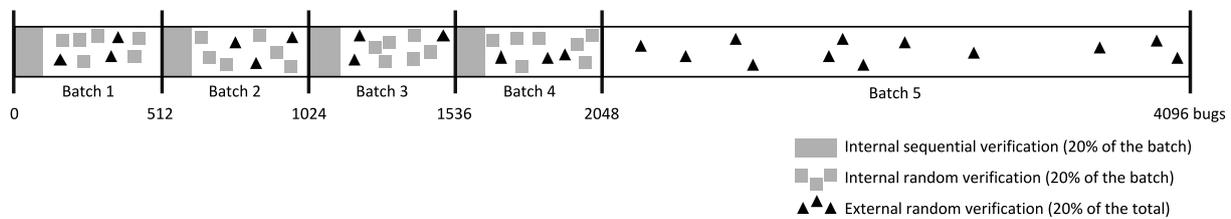


Fig. 1. Representation of the different verification tasks across the dataset.

Table 1

Detailed view of the verification tasks results.

	Defect Type	A/I	C	A/M	F/C/O	T/S	I/OOM	R
Internal (746 bugs)	Assignment/Initialization (A/I)	34		4				
	Checking (C)		71	6				
	Algorithm/Method (A/M)	4	4	394	13	1	4	
	Function/Class/Object (F/C/O)			6	126		1	1
	Timing/Serialization (T/S)			1		4		
	Interface/O-O Messages (I/OOM)	1		3	1		65	1
	Relationship (R)							1
External 1 (379 bugs)	Assignment/Initialization	20		2				
	Checking		41	3				
	Algorithm/Method		2	216	5			
	Function/Class/Object			3	44			
	Timing/Serialization			1		6		
	Interface/O-O Messages		2	1	3		29	
	Relationship							1
External 2 (382 bugs)	Assignment/Initialization	14		8			9	
	Checking		34	8				
	Algorithm/Method	5		219				
	Function/Class/Object			8	51			
	Timing/Serialization					2		
	Interface/O-O Messages			2			22	
	Relationship							

Table 2

Accuracy and Cohen's Kappa Agreement for researcher2 external verification.

ODC Attribute	Accuracy	k	Agreement
Activity	0.71	0.64	substantial
Trigger	0.61	0.59	moderate
Impact	0.82	0.81	almost perfect
Target	0.97	0.96	almost perfect
Defect Type	0.94	0.93	almost perfect
Qualifier	0.92	0.88	almost perfect

almost perfect when  $0.81 \leq k \leq 0.99$  [27]. The accuracy results of researcher2 are presented in Table 2, along with the Kappa value (referring to the agreement with researcher1).

Table 2 shows that we obtained almost perfect agreements for the majority of the ODC attributes. The exception is Activity with substantial agreement and Trigger with moderate agreement. In the former case, the most visible case of disagreement between the researchers is the case where code inspection bugs were marked by researcher2 as being one of the three possible Testing cases (i.e., unit testing, function testing, system testing). Indeed, researcher2 classified roughly one third of the 229 bugs originally marked with code inspection by researcher1 as being one of the three Testing activities (15 bugs attributed to unit testing, 30 to function testing, and 32 to system testing). This has direct impact in the Trigger classification, as, in this case, a wrong activity will lead to a wrong Trigger (this is due to the fact that some activities exclusively map into certain triggers). As such, we can expect that the accuracy values for Trigger are lower, which is actually the case. Due to time restrictions, researcher3 only classified the defect type attribute, for which it reached an accuracy of 0.90 which corresponds to a Kappa value of 0.88, also achieving an almost perfect agreement. Based on this, we consider that the

whole dataset is reliable, as we estimate the overall error is kept under similar values.

Table 3 holds a summary of the final outcome of the whole classification process and shows the details regarding the number of reports assigned to each label. As the Triggers directly map to the 5 Activities (i.e., Design Review/Code Inspection, Unit Test, Function Test, System Test), we also include the detailed values for the respective triggers. There is also a relation between the Defect Type and Qualifier attributes with the Target attribute, which only apply to reports classified with Target Code or Design (which means that not all of the 4096 defects qualify for being marked with a defect type/qualifier, in our case – 3846 defects are code defects). During the process, we had to exclude some classes (marked in gray in Table 3) due to their absence or relatively low number of instances in the dataset. We also identify (light blue cells) the smallest number of reports per attribute in Table 3 (explained in the next paragraph).

After observing the classified bug reports in detail, we noticed that the dataset was very imbalanced. The classifiers chosen (SVM, k-NN, NB, NC, RF and RNN) tend to classify a certain report using the most prevalent class if the dataset is imbalanced. An obvious solution is to use a sampling technique to balance the dataset. We decided to use an undersampling technique rather than oversampling because, in this context, oversampling would imply artificially generating bug reports (including mimicking conversations between developers), which is a very complex task and very prone to errors.

Undersampling consists of decreasing the number of reports so that their number will be the same for all classes. Thus, the class with less reports (marked in light blue in Table 3) keep all information while in the others, reports are randomly removed until their number reaches the number of reports in the less popular classes.

**Table 3**  
Number of documents per class in the dataset. Gray cells indicate removed classes due to small or nonexistent instances. Black cells indicate the smallest number of instances per attribute.

	Attribute	Classes	#	Classes	#
Opener Section	Activity	Design Review	1	System Test	658
		Code Inspection	2128	Unit Test	483
		Function Test	826	—	—
	Trigger - Code Inspection	Backward Compatibility	86	Internal Document	26
		Concurrency	210	Language Dependency	63
		Design Conformance	273	Lateral Compatibility	85
		Logic/Flow	960	Rare Situation	14
		Side Effects	503	—	—
	Trigger - Function Test	Coverage	535	Variation	233
		Sequencing	52	Interaction	6
	Trigger - System Test	Blocked Test	352	Software Configuration	88
		Recovery/Exception	26	Startup/Restart	68
		Workload/Stress	116	Hardware Configuration	7
	Trigger - Unit	Simple Path	413	Complex Path	70
	Impact	Capability	2165	Standards	98
		Integrity/Security	187	Usability	45
		Performance	199	Accessibility	0
		Reliability	988	Documentation	1
		Requirements	227	Installability	14
		Serviceability	156	Maintenance	1
Migration		15	—	—	
Closer Section	Target	Code	3846	Information Development	20
		Build/Package	230	Requirements	0
		Design	0	National Language	0
		Algorithm/Method	2365	Function/Class/Object	543
	Defect Type	Assignment/Initializatio	169	Interface O-O Messages	312
		Checking	407	Relationship	8
		Timing/Serialization	42	—	—
		Qualifier	Extraneous	158	Missing
		Incorrect	2529	—	—

### 3.2. Feature extraction

The goal of this step is to extract representative values from the original text reports that are supposed to be informative and non-redundant (as a way to help the machine learning algorithms). Several feature extraction techniques, described in the next paragraphs, can be applied to non-structured text data, namely: *bag-of-words* [28]; term frequency (TF) [28]; term frequency combined with inverse document frequency (TF-IDF) [28–30].

The *bag-of-words* technique, where each word counts as 1 when it appears in a document and 0 otherwise, does not allow to distinguish uncommon words from words that appear in the entire dataset, since all of them would carry the same value.

The TF technique, where an array of numbers represents how many times each word appears in a document, gives a higher value to the words that appear most often in a document. This may be problematic, given that the words that appear most often in a document are usually words which are widely used in a given context. Therefore, the classifiers would very likely show poor performance since all the classes would be defined by the same best words.

The TF-IDF computes a weight that discriminates the occurrence frequency of a given word in the entire dataset and allows us to understand its importance for a certain document. The TF part counts how many times it appears in the document and the IDF part is a ratio between the total number of documents and the number of documents where a certain term  $t$  appears. A higher number of occurrences of a word in the entire dataset decreases the word’s importance for the classification problem [29,30]. In this work we use TF-IDF.

The IDF value for a certain term  $t$  in a dataset  $D$  is given by Eq. (2):

$$IDF(t, D) = \log_{10} \left( \frac{N}{|\{d \in D : t \in d\}|} \right). \quad (2)$$

where  $N$  is the number of documents in the data set  $D$ , and the denominator is the number of documents where the given term  $t$  appears. If the denominator is 0 it is added one unit to make the division possible. When we multiply the term counter values with the IDF values we obtain a value (TF-IDF) that allow us to know the importance of a word in each document. The words that are the most important are those that have larger values, i.e. those that appear more times in each document and less times in the entire data set. The TF-IDF is given by:

$$TF-IDF(t, d, D) = TF(t, d) \times IDF(t, D). \quad (3)$$

where  $TF(t, d)$  represents the number of times that term  $t$  appears in a document. Therefore, at the end each report is represented by a vector  $\mathbf{x}$  with all the values that resulted from the application of TF-IDF.

Finally, the TF-IDF values are normalized using L2 Normalization represented by Expression (4):

$$L2 \text{ Normalization: } \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \dots \\ x_n \end{bmatrix} \rightarrow \begin{bmatrix} \frac{x_1}{\sqrt{x_1^2 + x_2^2 + \dots + x_n^2}} \\ \frac{x_2}{\sqrt{x_1^2 + x_2^2 + \dots + x_n^2}} \\ \dots \\ \frac{x_n}{\sqrt{x_1^2 + x_2^2 + \dots + x_n^2}} \end{bmatrix} \quad (4)$$

This normalization algorithm converts the vectors in normalized vectors where the sum of the squares of vector elements is 1, i.e., it is also known as euclidean norm.

The classification process included training (using part of the dataset) and testing (the execution of the classifiers using unknown data). Thus, we divided the dataset in these two parts (see Section 3.5 for further details). For performing feature extraction over the bug reports that belong to the training dataset, a dictionary was developed with all the words that appeared in the set and their respective number of occurrences. For this purpose, we relied on text feature extraction functions available in Scikit-learn: Machine Learning in Python [31]. For the testing set, the dictionary built with the training set was used to convert words into numbers using Eqs. (2) and (3). Although, the majority of state-of-art text classification models remove the words that appear too frequently, called stop words, we decided to keep them since these type of words are not useless in our case [5].

### 3.3. Dimensionality reduction

Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) are likely the most used Dimensionality Reduction methods in Machine Learning. On one hand, PCA is a non-supervised approach that tries to project features into a low-dimensional space, preserving the data variance as most as possible. On the other hand, LDA is supervised and the goal is to find a projection that maximizes data clusters separability and compactness [32].

Due to its simplicity, PCA was used to diminish the ratio between the number of samples and the number of features, known as Dimensionality Ratio [33]. High dimensionality ratios reduce the risk of overfitting (training data specialization), i.e., tend to improve the capacity to classify correctly new data. For this purpose, it was necessary to select a dimensionality ratio that would allow us to keep as much information as possible, while keeping the abovementioned ratio higher than one (values below one mean that we have more features than samples and accuracy in new data tend to be very low). In the case of this work, the value that we used was half of the number of documents in the

balanced dataset for each attribute. This leads to a dimensionality ratio of two, preserving around 70% of the data variance for all ODC categories, which is pointed as the minimum value of preserved variance in the literature [34]. After applying PCA for Dimensionality Reduction, we obtained new vectors  $\mathbf{x}$  with new features created based on the old best features, i.e, based on the best discriminative words.

### 3.4. Classifiers

Six classifiers that fit into six different types were considered: the  $k$ -Nearest Neighbor classifier ( $k$ -NN), known to be a lazy learner; a Bayesian classifier (Naïve Bayes), i.e., a parametric approach; the Support Vector Machine (SVM), which is a maximal-margin classifier; the Random Forest (RF), an ensemble classifier; the Nearest Centroid (NC), a minimum distance classifier; and the Recurrent Neural Network (RNN) which in this work was implemented as a deep learning architecture.

The  $k$  nearest neighbors ( $k$ -NN) classifier is named lazy learner because its training is just the storage of labeled data [35]. A new pattern is classified by finding the  $k$  closest neighbors, and the class of the new pattern will be the most prevalent class among the  $k$  closest neighbors. Thus, the  $k$  parameter should be odd to not allow the occurrence of draws. The best  $k$  value is usually obtained by experimentation, as we did in this work. Concerning the determination of the  $k$  closest neighbors, we opted to use the Euclidean distance, which is known to be the most used metric.

A Bayesian classifier is based on the Bayes rule, which is given by:

$$P(\omega_i|\mathbf{x}) = \frac{P(\omega_i)p(\mathbf{x}|\omega_i)}{p(\mathbf{x})}. \quad (5)$$

where  $\mathbf{x}$  is a feature vector corresponding to a given bug report that we want to classify,  $\omega_i$  is the  $i$ th class,  $p(\mathbf{x}|\omega_i)$  is the conditional probability density function (*pdf*) that describes the occurrence of a given vector  $\mathbf{x}$  when the class  $\omega_i$  is known,  $P(\omega_i)$  is the probability of occurrence of objects of class  $\omega_i$ ,  $p(\mathbf{x})$  is the *pdf* that define the probability of occurrence of a given pattern  $\mathbf{x}$  independently of the class, and  $P(\omega_i|\mathbf{x})$  is the probability that knowing a pattern  $\mathbf{x}$  its class is  $\omega_i$ . Thus in a classification problem we assign the class that corresponds to the biggest  $P(\omega_i|\mathbf{x})$  to  $\mathbf{x}$ . The classification rule is:

$$\mathbf{x} \in \omega_i \text{ if } P(\omega_i|\mathbf{x}) > P(\omega_j|\mathbf{x}) \forall i \neq j. \quad (6)$$

The training of a Bayesian classifier encompasses the determination of  $P(\omega_i)$  and  $p(\mathbf{x}|\omega_i)$  for all the classes. The former is estimated by computing the prevalence of patterns from a given class in the training set, for example, for a given class  $\omega_i$ :  $P(\omega_i) = n_i/N$ , being  $N$  the total number of patterns in the training set and  $n_i$  the number of patterns that belong to  $\omega_i$ . The latter involves the estimation of a *pdf* from data. Naïve Bayes (NB) is a particular case of the general Bayesian classifier, where independence between every pair of features is assumed, thus the estimation of  $p(\mathbf{x}|\omega_i)$  is simply given by:

$$p(\mathbf{x}|\omega_i) \approx \prod_{k=1}^F p(x_k|\omega_i), \quad (7)$$

where  $x_k$  is a given value of the  $k$ th feature and  $F$  the total number of features. Concerning the distribution type assumed for each  $p(x_k|\omega_i)$ , in this paper we considered the multinomial distribution, which is known to be a good choice for text classification [36].

In its native formulation the support vector machine (SVM) finds a decision hyperplane that maximizes the margin that separates two different classes [37]. Considering non-linearly separable training data pairs  $\{\mathbf{x}_i, y_i\}$  (with  $i = 1, \dots, N$ ), where

$y_i \in \{-1, 1\}$  are the class labels and  $\mathbf{x}_i \in R^d$  are corresponding features vectors, the function that defines the hyperplane is:

$$y_i (\mathbf{w}^T \mathbf{x}_i + b) \geq 1 - \xi_i; \quad y_i \in \{-1, 1\} \quad (8)$$

where  $\mathbf{w} = [w_1, \dots, w_d]^T$  is the vector normal to the hyperplane,  $b/(\|\mathbf{w}\|)$  is the hyperplane offset to the origin and  $\xi$  is a quantification of the degree of misclassification.  $\mathbf{w}$  and  $b$  are obtained by minimizing:

$$\Psi(\mathbf{w}) = \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^N \xi_i, \text{ subject to} \quad (9)$$

$$y_i (\mathbf{w}^T \mathbf{x}_i + b) \geq 1 - \xi_i; \quad i = 1, \dots, N.$$

$C$  defines the influence of  $\xi$  on the minimization criterion  $\Psi$ .

In a classification problem it may happen that data can be separated but by non-linear decision boundaries. SVM enable the definition of non-linear decision functions by applying a data transformation into a new space where data may be easier to classify by a linear decision function given by [38]:

$$d(\mathbf{x}) = w_1 f_1(\mathbf{x}) + \dots + w_k f_k(\mathbf{x}) + b, \quad (10)$$

where  $\mathbf{f} = [f_1, \dots, f_k]^T$  implements the non-linear transformation from the original feature space to a new space of dimension  $k$ . It can be demonstrated that by considering the dual optimization problem that the decision function is given by:

$$d(\mathbf{x}) = \sum_{SVs} \alpha_i y_i \mathbf{f}^T(\mathbf{x}_i) \mathbf{f}(\mathbf{x}) = \sum_{SVs} \alpha_i y_i (\mathbf{f}(\mathbf{x}_i) \cdot \mathbf{f}(\mathbf{x})) = \sum_{SVs} \alpha_i y_i K(\mathbf{x}_i, \mathbf{x}). \quad (11)$$

Where  $\alpha_i$  and  $y_i$  are the Lagrange multiplier and the  $\pm 1$  label associated with support vector  $\mathbf{x}_i$ , respectively.  $K(\mathbf{x}_i, \mathbf{x}) = \mathbf{f}(\mathbf{x}_i) \cdot \mathbf{f}(\mathbf{x})$  is the inner-product kernel, where the most popular and used one is the Gaussian RBF, given by:

$$K(\mathbf{x}_i, \mathbf{x}) = e^{-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2\sigma^2}} = e^{-\gamma \|\mathbf{x} - \mathbf{x}_i\|^2}. \quad (12)$$

$\gamma = \frac{1}{2\sigma^2}$  is a free parameter that controls the kernel function aperture, and should be subjected to a grid search procedure in conjunction with the  $C$  parameter.

The SVM is defined as a two-class classifier, but ODC classification is a multi-class problem. The usual approach for multi-class SVM classification is to use a combination of several binary SVM classifiers. Different methods exist but, in this work, we used the one-against-all multi-class approach [39]. This method transforms the multi-class problem, with  $c$  classes, into a series of  $c$  binary sub problems that can be solved by the binary SVM. The  $i$ th classifier output function  $\rho_i$  is trained taking the examples from a given class as 1 and the examples from all other classes as  $-1$ . For a new example  $\mathbf{x}$ , this method assigns to  $\mathbf{x}$  the class that were assigned in majority by all the binary classifiers [39].

The Random Forest (RF) classification is made by considering a set of random decision trees [40]. The trees training is usually performed with the “bagging” (Bootstrap aggregating) method [41], with the ambition that a combination of simple learning models increases the overall classification performance. The “random” designation comes from the fact that individual trees are based on a random selection of the available features. RF is less prone to overfitting than single decision trees because of its random nature. When compared to non-ensemble approaches, this type of algorithms tends to return more stable results, i.e., they present a reduced variance.

The first step in a Nearest Centroid (NC) classifier is the computation of the centroids based on a training dataset, i.e., the mean vector for each class. In the testing phase a new sample is classified by computing the distance between it and the centroids,

**Table 4**

Best parameter values that lead to the highest accuracy for each classifier and ODC attribute.

Attribute	Linear SVM	k-NN	Random Forest	RBF SVM		RNN	
	C	k	Number of Trees	C	Gamma	Hidden Vector Dimension	Dropout
Activity	1	17	1024	8	0.0625	8	20%
Trigger - Code Inspection	2	17	512	2	16	64	30%
Trigger - Function Test	2	5	512	4	0.25	64	30%
Trigger - System Test	1	11	256	4	0.03125	64	10%
Trigger - Unit Test	0.25	19	128	2	0.125	32	10%
Impact	1	21	512	8	0.25	64	50%
Target	16	15	512	32	0.0625	8	30%
Defect Type	8	7	256	8	0.125	16	30%
Qualifier	16	17	256	16	0.0625	16	20%

related to the different classes, then the class of the sample will be the class of the nearest centroid. In the text classification domain using TF-IDF feature vectors NC is also known as the Rocchio classifier [28].

Recurrent neural networks (RNNs) are popular neural networks used for text classification [42,43]. They allow to use past information for making better decisions in the present i.e., they are able to learn relations between words of each document instead of looking at each word independently. Therefore, they are classifiers that are aware of the sequence (previous, actual and next). However, simple RNN architecture suffers from vanishing gradient problem [44] and long term dependency problem as well [45]. In order to overcome such disadvantages, Hochreiter and Schmidhuber presented the long-short-term memory (LSTM) [46], a type of RNN that is able to control the flow of the information over the sequence steps (e.g. over the words of a text), and that may be composed by several LSTM units, forming a deep-learning architecture.

In this paper we used a RNN model that is based on LSTM units. For a better exploration of the text we used two LSTM layers, one which learns the texts from left to the right and other which does the inverse, named Bi-LSTM [47,48]. For training this architecture we used a batch size higher than one. Thus, we had to convert all the documents to the same length. We decided to use a threshold equal to the average length of all documents plus three times the standard deviation which allow us to maintain almost 99% of the information of all documents. Another method would be using the maximum length of the training set. However, this would lead to have longer documents just for more 1% of information. Then, we pad the smaller documents with the word “padding” and remove words from the documents longer than this threshold. Differently from the classical algorithms, for training RNN classifier we used word embeddings instead of TF-IDF feature. For training the word embedding model responsible for converting each word in a embedding vector we used Word2Vec algorithm [49] and all the reports presented on the MongoDB, Cassandra and HBase databases at the moment of the training of the algorithms (17,117 reports) allowing us to use non-annotated data to classify labeled reports. Our deep-learning model architecture is presented in Fig. 2. For the RNNs development we used the Keras API [50].

Configuring the classifiers mostly involves finding the best hyperparameters for our SVM, k-NN, RF and RNN. Thus, we performed a search, usually referred as Grid-Search, for the parameters that lead to the best classification, i.e., that leads to the best accuracy. The NB and NC classifiers have no free parameters and thus it was not necessary to perform this search. About the SVM we applied two classifiers, one with a linear kernel and other with Gaussian RBF kernel. Then, we searched for the C values for the linear one and for the C and  $\gamma$  values combinations for the other, for k-NN we searched for the number of neighbors ( $k$ ), for RF for the appropriate number of trees, and for RNN the appropriate hidden vector dimension and dropout percentage.

An application of the Grid-Search method is represented in Fig. 3, where the goal is to find both C values,  $\gamma$ ,  $k$ , the number of trees, the hidden vector dimension, and the dropout percentage that provides high accuracy (the best values, in the case of the example shown in Fig. 3, are marked in red). As presented in the Figure, we varied both C values and  $\gamma$  in powers of two between  $2^{-5}$  and  $2^5$ ,  $k$  assumed odd values between 1 and 21, the number of trees tested were  $[2^0, 2^1, \dots, 2^{10}]$ , the hidden vector dimension varied in powers of two between  $2^3$  and  $2^6$ , and dropout between 10% and 50% using a step of 10%. For both SVM models, we should choose a C value that leads to high accuracies but that at the same time is small. Larger C values lead to the definition of smaller separation margins based on the training data and the probability of overfitting is high. On the other hand, for k-NN we have the reverse situation, i.e., the number of neighbors should not be so small in order to prevent overfitting. RF is known to be very resistant to overfitting, even with a higher number of trees. The number of trees can result in a problem of complexity, and we should use a good compromise between performance and complexity. Regarding the RNN classifier, keeping the network with a low hidden vector dimension prevents overfitting [51]. Furthermore, the dropout regularization is also fundamental to make better relations between the features preventing overfitting as well [52]. Although Fig. 3 only shows the values for the Activity attribute, the procedure was exactly the same for all other ODC attributes and therefore we omit the graphical view of the remaining searches (all detailed values and graphics may be found at <http://odc.dei.uc.pt>).

Table 4 presents all the best values for C,  $\gamma$ ,  $k$ , number of trees, hidden vector dimension and dropout percentage. As previously mentioned, these values were selected by examining the best accuracies found with lower values of C,  $k$  and number of trees. We found out that the best C values for Linear SVM were confined to the range  $[2^{-2}, 2^4]$ , the best  $k$  values to the range  $[5, 21]$ , the best number of trees to the range  $[2^7, 2^{10}]$ , the best C values for RBF SVM to the range  $[2^0, 2^5]$  and the best  $\gamma$  values to the range  $[2^{-5}, 2^0]$ . Finally, for RNN classifier the best hidden vector dimension values were confined to the range  $[2^3, 2^4]$  and the best dropout percentage to the range  $[10\%, 20\%]$ .

### 3.5. Classifier assessment

As mentioned earlier, the classification process includes training (using part of the dataset) and testing (using unknown data, also known as out-of-sample data). It is important to mention that, for training the classifiers related to ‘open section’ ODC attributes we just used the title and description of the software defect. Thus, for these cases, we do not use the posterior developer comments, which many times discuss how a fix should be performed and would not fit the open section definition. For training the classifiers related with the ‘close section’ ODC attributes, we use all available information regarding each software defect.

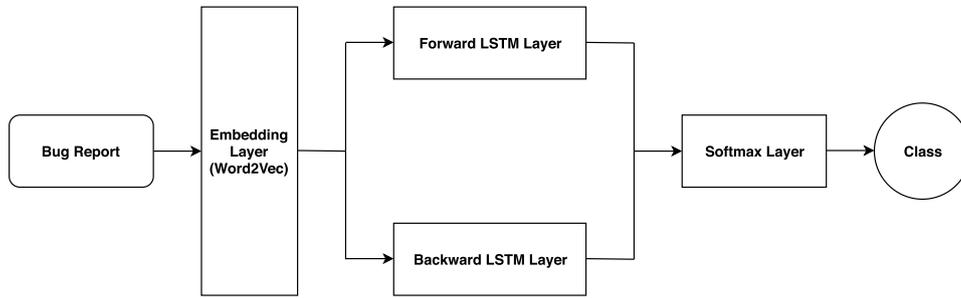


Fig. 2. RNN architecture used for classification of the bug reports.

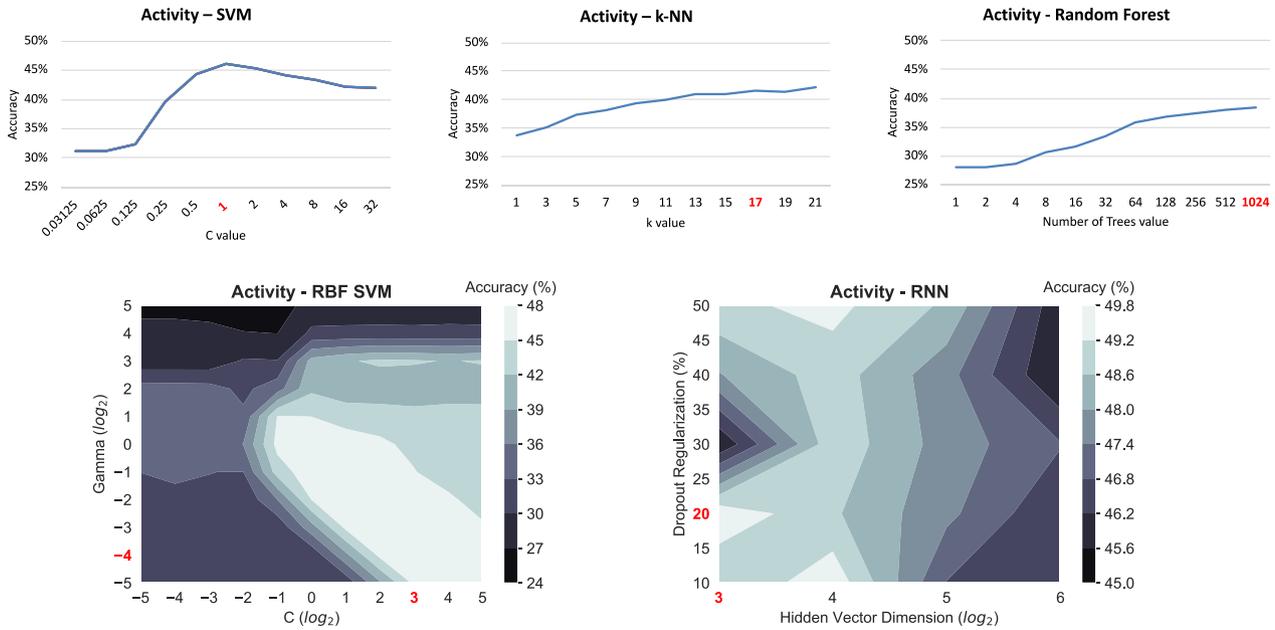


Fig. 3. Grid search for hyperparameters using Activity.

Table 5

Generic confusion matrix, where “ $\hat{c}_k$ ” indicates predictions provided by the algorithms.

		Predicted			
		$\hat{c}_1$	$\hat{c}_2$	...	$\hat{c}_K$
True	$c_1$	$n_{1,1}$	$n_{1,2}$	...	$n_{1,K}$
	$c_2$	$n_{2,1}$	$n_{2,2}$	...	$n_{2,K}$
	⋮	...	...	...	...
	$c_K$	$n_{K,1}$	$n_{K,2}$	...	$n_{K,K}$

In order to validate our classifiers, we used the K-Fold Cross Validation (CV) approach, with ten folds (i.e.,  $K = 10$ ) [53], which is known to be the most classic validation strategy. Supposing that the training data is composed by M bug reports (i.e., samples), the algorithm will split the data in 10 parts, where each part contains  $M/10$  samples. Afterwards the classifiers are trained 10 times for the same data. In each iteration, the classifiers are trained with nine parts and the part that is “left out” is used for testing. The algorithm output is compared to the real labeling by filling confusion matrices. Different ODC attributes present different numbers of classes, thus, confusion matrices with different sizes were filled. These can be represented by the generic confusion matrix illustrated in Table 5.

The diagonal terms  $n_{i,j}$ , where  $i = j$ , correspond to the instances where the algorithm’s output was consistent with the

real class label, i.e. the true positive patterns. The values  $n_{i,j}$ , where  $i \neq j$ , are the number of instances misclassified by the algorithm, which can be considered as false positives or false negatives depending on the class under analysis. Considering a given class k, four measurements can be taken:

- True positives ( $TP_k$ ): number of instances correctly classified as class k;
- False positives ( $FP_k$ ): number of instances classified as k when in fact they belong to other class;
- True negatives ( $TN_k$ ): number of instances correctly not classified as k;
- False negatives ( $FN_k$ ): number of instances assigned to other classes when in fact they belong to class k;

These measurements for a problem with K classes can be computed from the confusion matrix as follows:

$$\begin{aligned}
 n_{+,k} &= \sum_{i=1}^K n_{i,k}; \\
 n_{k,+} &= \sum_{j=1}^K n_{k,j}; \\
 TP_k &= n_{k,k}; \\
 FP_k &= n_{+,k} - n_{k,k}; \\
 FN_k &= n_{k,+} - n_{k,k}; \\
 TN_k &= N - TP_k - FP_k - FN_k.
 \end{aligned}
 \tag{13}$$

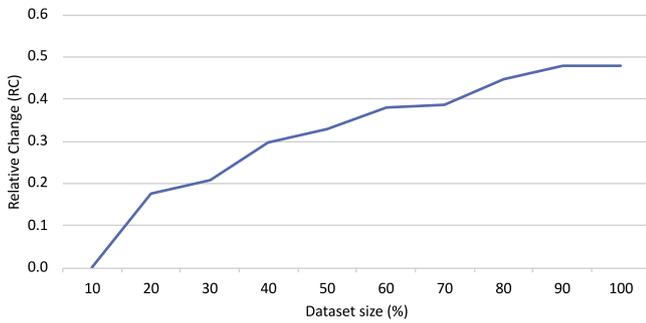


Fig. 4. Relative change for the Defect Type attribute.

where  $N$  is the number of patterns that do not belong to class  $k$ . We evaluated the algorithm's according to three metrics: accuracy, recall (or sensitivity), and precision. These metrics are respectively given by:

$$\text{accuracy}_k = \frac{TP_k + TN_k}{TP_k + FP_k + TN_k + FN_k} \quad (14)$$

$$\text{recall}_k = \frac{TP_k}{TP_k + FN_k} \quad (15)$$

$$\text{precision}_k = \frac{TP_k}{TP_k + FP_k} \quad (16)$$

In the next section we present the results obtained with each classifier and discuss results, based on the concepts presented in this section.

#### 4. Results and discussion

This section discusses the results obtained during this work in three different views. We first analyze how the size of the dataset affects the accuracy of the classifiers. We then overview the classification results according to the previously mentioned metrics (i.e., accuracy, recall, and precision) and using the whole dataset. Finally, we go through the detailed results obtained for each ODC attribute. The numbers shown in this section refer to the average of several runs, namely 25 runs using all algorithms, so that we could identify the best one per each ODC attribute. We report average performance indicators, as well as their variability in the testing data in the format *average value* ± *standard deviation*.

We begin by analyzing how the dataset size affects the overall accuracy of the classification. For this purpose, we carried out a set of tests using Linear SVM, which has a fast training and it was also one of the models that revealed the best accuracy values while classifying the complete dataset (see next paragraphs for details on the accuracy values obtained by each algorithm). We varied the size of the dataset using increments of 10% and computed the relative change of the accuracy across the different dataset sizes and using the smallest dataset size (i.e., 10%) as reference. The relative change is calculated as follows:

$$RC(a_i, a_{ref}) = \frac{a_i - a_{ref}}{a_{ref}} \quad (17)$$

where  $a_{ref}$  represents the reference accuracy value for the smallest dataset size (in this case, 10%) and  $a_i$  represents the accuracy obtained for a dataset with a certain size  $i$ .

Fig. 4 presents a graphical view of the relative change values for a single ODC attribute — defect type, which confirms our general expectation that a larger dataset size would lead to better accuracy values. Table 6 presents a more detailed vision of the relative change of accuracy for all ODC attributes per varying dataset sizes.

Table 6  
Relative change values for all ODC attributes.

ODC Attribute	Dataset Size										after balancing
	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%	
Activity	0.00	0.10	0.15	0.19	0.23	0.24	0.25	0.26	0.28	0.29	1932
Code Inspection	0.00	0.23	0.40	0.54	0.65	0.68	0.74	0.78	0.81	0.89	441
Function Test	0.00	-0.02	-0.05	-0.03	-0.08	-0.04	-0.04	0.00	-0.03	-0.04	156
System Test	0.00	-0.03	0.20	0.38	0.36	0.55	0.50	0.68	0.71	0.69	130
Unit Test	0.00	0.12	0.14	0.16	0.13	0.16	0.20	0.13	0.19	0.17	140
Impact	0.00	0.44	0.52	0.47	0.60	0.71	0.80	0.87	0.92	0.92	360
Target	0.00	0.08	0.12	0.17	0.20	0.20	0.22	0.23	0.24	0.25	460
Defect Type	0.00	0.18	0.21	0.30	0.33	0.38	0.39	0.45	0.48	0.48	845
Qualifier	0.00	0.04	0.01	0.05	0.10	0.09	0.10	0.12	0.10	0.10	474

In general, and as we can see in Table 6, all attributes (the exception is Function Test) benefit from increasing the dataset size. This is quite expectable, since machine learning algorithms tend to perform better and its performance is well estimated with increasing amounts of (good quality) information used during training [54]. In the particular case of the Function Test triggers, and after analyzing the data in further detail, we observed that the increasing amounts of data present in the different dataset sizes do not contribute towards the definition of a pattern that would allow distinguishing among the different triggers. It may be the case that a different dataset is required, or that the current sizes used are simply not sufficient for this particular case.

Additionally, we can observe that, for most of the attributes, the relative change of accuracy does not always progressively change with the increase of the dataset size. More specifically, Unit Test, System Test, Impact, Defect Type and Qualifier do not show the highest relative change value for the dataset size of 100%. Still, the differences between the highest values and those at 100% dataset size are relatively subtle.

Table 7 presents the overall results for all classifiers and per each ODC attribute, with the best values per attribute marked in orange (using the average of 25 runs). Fig. 5 visually shows the overall results for the best classifier and per each ODC attribute.

As we can see in Table 7 the Linear SVM and RNN classifiers seem to outperform all the others. Regarding the first one, this was expected since the data is sparse (a considerable number of features values are zero) and this type of classifier is known to handle these cases quite well considering it has a linear kernel, as discussed in [30]. It was also expected that RNN model got better results than the others since it is based on a powerful architecture capable of capturing the words sequence in each document and it is also able to use unannotated data for classifying the reports. Still, we wanted to know if the observed differences among the classifiers performance were statistically significant. Thus, we performed statistical tests which ended up by comparing the best classifier performance for each attribute, i.e., the model with the highest average accuracy for each attribute, with all the remaining classifiers for the same category. For each run we balance data, thus testing data is different between runs, which led us to consider classifier's performance as independent realizations. This allows us to use independent statistical tests. We decided to use t-test or the Mann–Whitney test, depending on the data following a normal distribution or not (respectively). To evaluate normality, we used two tests: Kolmogorov–Smirnov and Shapiro–Wilk. The t-test is only applied if both tests determine that data is normal.

Table 8 shows the cases where the model with the highest accuracy outperforms the remaining in a statistically significant manner (with an  $\alpha$  value of 0.05, i.e., 95% confidence). The cases marked with “X” refer to those where there are no statistical differences between the best model and a certain algorithm under analysis, e.g. in Activity attribute RNN model performs better than the other models while in Target Linear SVM outperforms all the others.

**Table 7**  
Overall classifier results per ODC attribute.

Attribute	Algorithm	Metric		
		Accuracy (%)	Recall (%)	Precision (%)
Activity	k-NN	41.3 ± 3.5	41.3 ± 3.5	41.4 ± 3.8
	NB	28.3 ± 2.7	28.3 ± 2.7	34.5 ± 12.6
	NC	40.5 ± 3.0	40.5 ± 3.0	45.8 ± 5.2
	RF	38.3 ± 3.0	38.3 ± 3.0	44.2 ± 7.4
	RNN	47.6 ± 3.5	47.6 ± 3.5	47.8 ± 3.4
	RBF SVM	45.8 ± 3.3	45.8 ± 3.3	46.1 ± 3.7
	Linear SVM	45.9 ± 3.3	45.9 ± 3.3	46.2 ± 3.6
Code Inspection	k-NN	38.6 ± 7.2	38.6 ± 7.2	39.0 ± 8.9
	NB	22.9 ± 4.4	22.9 ± 4.4	18.6 ± 10.3
	NC	37.4 ± 7.6	37.4 ± 7.6	40.4 ± 9.4
	RF	34.0 ± 6.8	34.0 ± 6.8	33.8 ± 10.0
	RNN	40.3 ± 7.0	40.3 ± 7.0	42.1 ± 8.2
	RBF SVM	36.1 ± 7.2	36.1 ± 7.2	42.3 ± 9.6
	Linear SVM	39.5 ± 7.0	39.5 ± 7.0	43.7 ± 7.9
Function Test	k-NN	35.7 ± 12.0	35.7 ± 12.0	35.0 ± 14.9
	NB	35.2 ± 8.9	35.2 ± 8.9	27.9 ± 17.3
	NC	35.7 ± 11.8	35.7 ± 11.8	35.7 ± 14.1
	RF	33.8 ± 10.9	33.8 ± 10.9	31.4 ± 16.9
	RNN	37.4 ± 12.8	37.4 ± 12.8	38.0 ± 15.7
	RBF SVM	36.1 ± 12.0	36.1 ± 12.0	35.6 ± 14.8
	Linear SVM	36.2 ± 11.4	36.2 ± 11.4	36.0 ± 13.7
System Test	k-NN	35.3 ± 12.3	35.3 ± 12.3	34.0 ± 15.2
	NB	26.2 ± 8.5	26.2 ± 8.5	17.1 ± 13.2
	NC	38.2 ± 12.2	38.2 ± 12.2	35.8 ± 15.9
	RF	29.0 ± 12.5	29.0 ± 12.5	25.6 ± 16.2
	RNN	41.4 ± 12.7	41.4 ± 12.7	40.7 ± 16.0
	RBF SVM	35.4 ± 12.4	35.4 ± 12.4	31.9 ± 15.9
	Linear SVM	39.0 ± 13.4	39.0 ± 13.4	37.5 ± 17.1
Unit Test	k-NN	57.0 ± 12.6	57.0 ± 12.6	57.6 ± 14.4
	NB	52.7 ± 9.3	52.7 ± 9.3	50.2 ± 22.2
	NC	57.0 ± 13.1	57.0 ± 13.1	57.7 ± 14.4
	RF	55.0 ± 12.5	55.0 ± 12.5	55.6 ± 15.2
	RNN	57.1 ± 12.7	57.1 ± 12.7	57.6 ± 14.9
	RBF SVM	56.3 ± 12.5	56.3 ± 12.5	57.6 ± 15.3
	Linear SVM	57.5 ± 12.3	57.5 ± 12.3	58.5 ± 14.7
Impact	k-NN	28.7 ± 7.1	28.7 ± 7.1	29.1 ± 8.9
	NB	13.8 ± 2.5	13.8 ± 2.5	6.6 ± 7.3
	NC	31.5 ± 7.3	31.5 ± 7.3	31.0 ± 9.1
	RF	20.3 ± 5.6	20.3 ± 5.6	21.4 ± 10.9
	RNN	33.3 ± 7.5	33.3 ± 7.5	34.4 ± 9.7
	RBF SVM	30.8 ± 8.3	30.8 ± 8.3	33.1 ± 10.3
	Linear SVM	31.6 ± 7.1	31.6 ± 7.1	34.2 ± 9.1
Target	k-NN	75.3 ± 6.3	75.3 ± 6.3	78.2 ± 6.2
	NB	51.9 ± 2.4	51.9 ± 2.4	55.1 ± 24.1
	NC	65.2 ± 8.1	65.2 ± 8.1	67.1 ± 8.5
	RF	74.0 ± 6.4	74.0 ± 6.4	79.7 ± 5.4
	RNN	84.7 ± 5.2	84.7 ± 5.2	85.3 ± 5.0
	RBF SVM	82.7 ± 5.5	82.7 ± 5.5	83.2 ± 5.5
	Linear SVM	85.6 ± 5.1	85.6 ± 5.1	86.0 ± 5.1
Defect Type	k-NN	28.3 ± 4.8	28.3 ± 4.8	27.8 ± 6.0
	NB	20.8 ± 2.0	20.8 ± 2.0	11.2 ± 9.8
	NC	28.8 ± 4.8	28.8 ± 4.8	30.3 ± 7.1
	RF	30.2 ± 4.4	30.2 ± 4.4	29.6 ± 6.1
	RNN	27.3 ± 4.8	27.3 ± 4.8	27.4 ± 5.1
	RBF SVM	29.7 ± 4.6	29.7 ± 4.6	29.7 ± 4.9
	Linear SVM	34.7 ± 5.1	34.7 ± 5.1	34.8 ± 5.5
Qualifier	k-NN	37.3 ± 7.4	37.3 ± 7.4	37.9 ± 7.8
	NB	34.1 ± 3.4	34.1 ± 3.4	22.6 ± 14.0
	NC	36.6 ± 7.2	36.6 ± 7.2	37.5 ± 8.5
	RF	37.7 ± 7.2	37.7 ± 7.2	39.2 ± 9.4
	RNN	38.9 ± 6.7	38.9 ± 6.7	39.1 ± 7.3
	RBF SVM	39.7 ± 6.8	39.7 ± 6.8	40.3 ± 7.2
	Linear SVM	39.7 ± 7.3	39.7 ± 7.3	40.0 ± 7.5

As we can see in Table 8, there are statistically significant differences between the performance of RNN and the remaining classifiers for Activity, Trigger – System Test and Impact while for Trigger – Code Inspection and Trigger – Function Test although RNN got higher average accuracy than Linear SVM there are no statistical differences between them. Regarding Linear SVM, there are statistical differences between it and all the others for Target and Defect Type while for Trigger – Unit Test there are only statistical differences between it, Naïve Bayes and Random Forest. About the Qualifier attribute, although RBF SVM got the highest

average performance there are no statistical differences between it and RNN and Linear SVM models.

Thus, we can conclude that RNN and Linear SVM models were the models with the best performances. Fig. 5 shows the best results obtained for each ODC attribute in a graphical manner.

As we can see in Fig. 5, we were able to reach an accuracy of  $47.6\% \pm 3.5\%$  for the Activity attribute, which essentially means that the 4 classes for the Activity attribute seem to be a little bit hard to separate using our data. For the Trigger attribute, which we have previously divided in parts according to the respective Activity (according to the Trigger-Activity mapping in [10]), we obtained accuracy values of  $40.3\% \pm 7.0\%$  for Code Inspection,  $37.4\% \pm 12.8\%$  for Function Test,  $41.4\% \pm 12.7\%$  for System Test and  $57.5\% \pm 12.3\%$  for Unit Test. Since the Trigger attribute is directly related with the Activity [10], these results were expected since the Activity results were not high. In Function Test, System Test and Unit Test, the standard deviation values are relatively high, suggesting that there is still high variability in the data in what concerns these three attributes.

The accuracy obtained for the Impact attribute was relatively low ( $33.3\% \pm 7.5\%$ ), which represents the worst classification performance since it has about half of its classes with no data and still has low accuracy. However, considering that we are still using a relatively large number of classes (i.e., 8 classes) and not a very large quantity of data, results certainly have space for improvement.

Regarding the Target attribute, we were able to obtain  $85.6\% \pm 5.1\%$  of accuracy (the highest of all attributes). Still, it is important to mention that we just used 2 out of the 6 classes (i.e., we used Code and also Build/Package), as we did not have enough data for the remaining 4 classes. The high accuracy for Target was expected since, in the case of this attribute, it is really easy to separate the respective classes just by looking for some keywords, as software defects which are not associated with Code (e.g., defects associated with build scripts) usually have that information explicitly mentioned in the bug reports.

Finally, for the Defect Type and Qualifier attributes we obtained  $34.7\% \pm 5.1\%$  and  $39.7\% \pm 6.8\%$  of accuracy, respectively. These low values obtained for Defect Type and Qualifier are expected, as these attributes are fully related to what is changed in the code, thus being the most difficult cases to classify solely based on the text of the bug reports. Thus, to improve the results, it is important to also analyze the code changes (i.e., using other means), which is something that falls outside of the scope of this paper and is left for future work.

Overall, we were able to easily automate the classification regarding the ODC Target attribute. The remaining cases revealed to be more difficult to handle. For these, and in general, the algorithm associated with the best performance showed to be correct between one third and half of the times. Although the values obtained are not high, we find them to be acceptable, mostly due to the relatively low number of reports available (and especially as a consequence of the data balancing performed during this study). As it is presented in Table 6, the classifiers tend to perform better if an adequate quantity of data is used for training, thus gathering more good quality data is a crucial step for a more effective classification procedure.

We now discuss the detailed results per ODC attribute, to understand in which classes the algorithms tend to perform better and in which ones the difficulties are higher. For this purpose, we analyze the 9 confusion matrices (one per ODC attribute, including the 4 groups of triggers), as previously explained in Section 3.5. It should be mentioned that the values shown in the confusion matrices are not integers since they represent an average of 25 runs.

Fig. 6a holds the detailed results for the Activity attribute. As we can see, for this classifier it is relatively easy to identify the

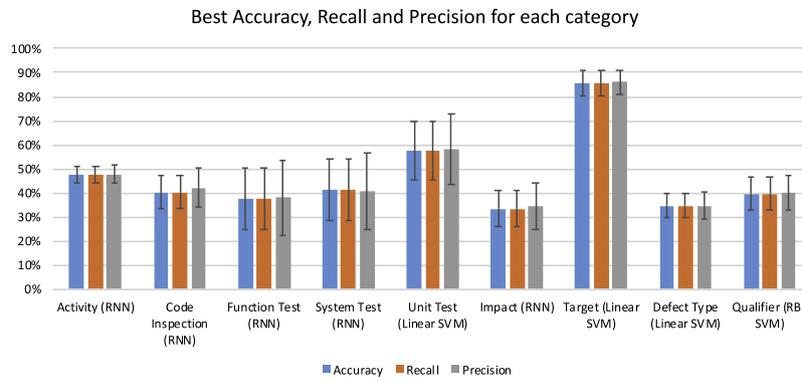


Fig. 5. Best classifier results per ODC attribute. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 8  
Statistical difference of SVM against the remaining algorithms.

Algorithm	Activity	Trigger Code Inspection	Trigger Function Test	Trigger System Test	Trigger Unit Test	Impact	Target	Defect Type	Qualifier
Best Classifier	RNN	RNN	RNN	RNN	Lin. SVM	RNN	Lin. SVM	Lin. SVM	RBF SVM
k-Nearest Neighbours			X		X				
Naïve Bayes									
Nearest Centroid			X		X				
Random Forest									
RNN					X				X
RBF SVM			X		X				
Linear SVM		X	X						X

System Test and also the Code Inspection classes. Conceptually, System Test (which refers to testing whole systems) should also be relatively easy to identify. It was also expectable that code inspection could be easily distinguishable from the remaining. On the other hand, we were expecting to see the classifier struggling with Unit Test and Function Test, which in fact was confirmed by the results. Unit Test refers to white-box testing of small units of code, while Function Test refers to testing functions and this separation is not always very clear (e.g., in some cases, a unit may be a function, the bug report may not be clear about this). In fact, the independent classification carried out by *researcher2* (see Section 3) resulted only in *moderate agreement*, which means that identifying the right activity used in the disclosure of a software defect is a somewhat difficult task.

Fig. 6b presents the Trigger – Code Inspection classifier results. Overall, the classifier experienced more difficulties when identifying Logic/Flow, which is also the one with the most generic definition [10] and, possibly more prone to being confused with some other. We can observe that quite often our classifier confuses this class with Design Conformance, Side Effects, and Language Dependency.

On the opposite side, the easiest class to identify is Concurrency. Although there are no overlaps in the definition of any of the classes, is it true that, conceptually, concurrency problems are quite easy cases to identify due to their own specificity. This leaves little error margin for human classifiers which in the end reflects on the performance of machine learning classifiers.

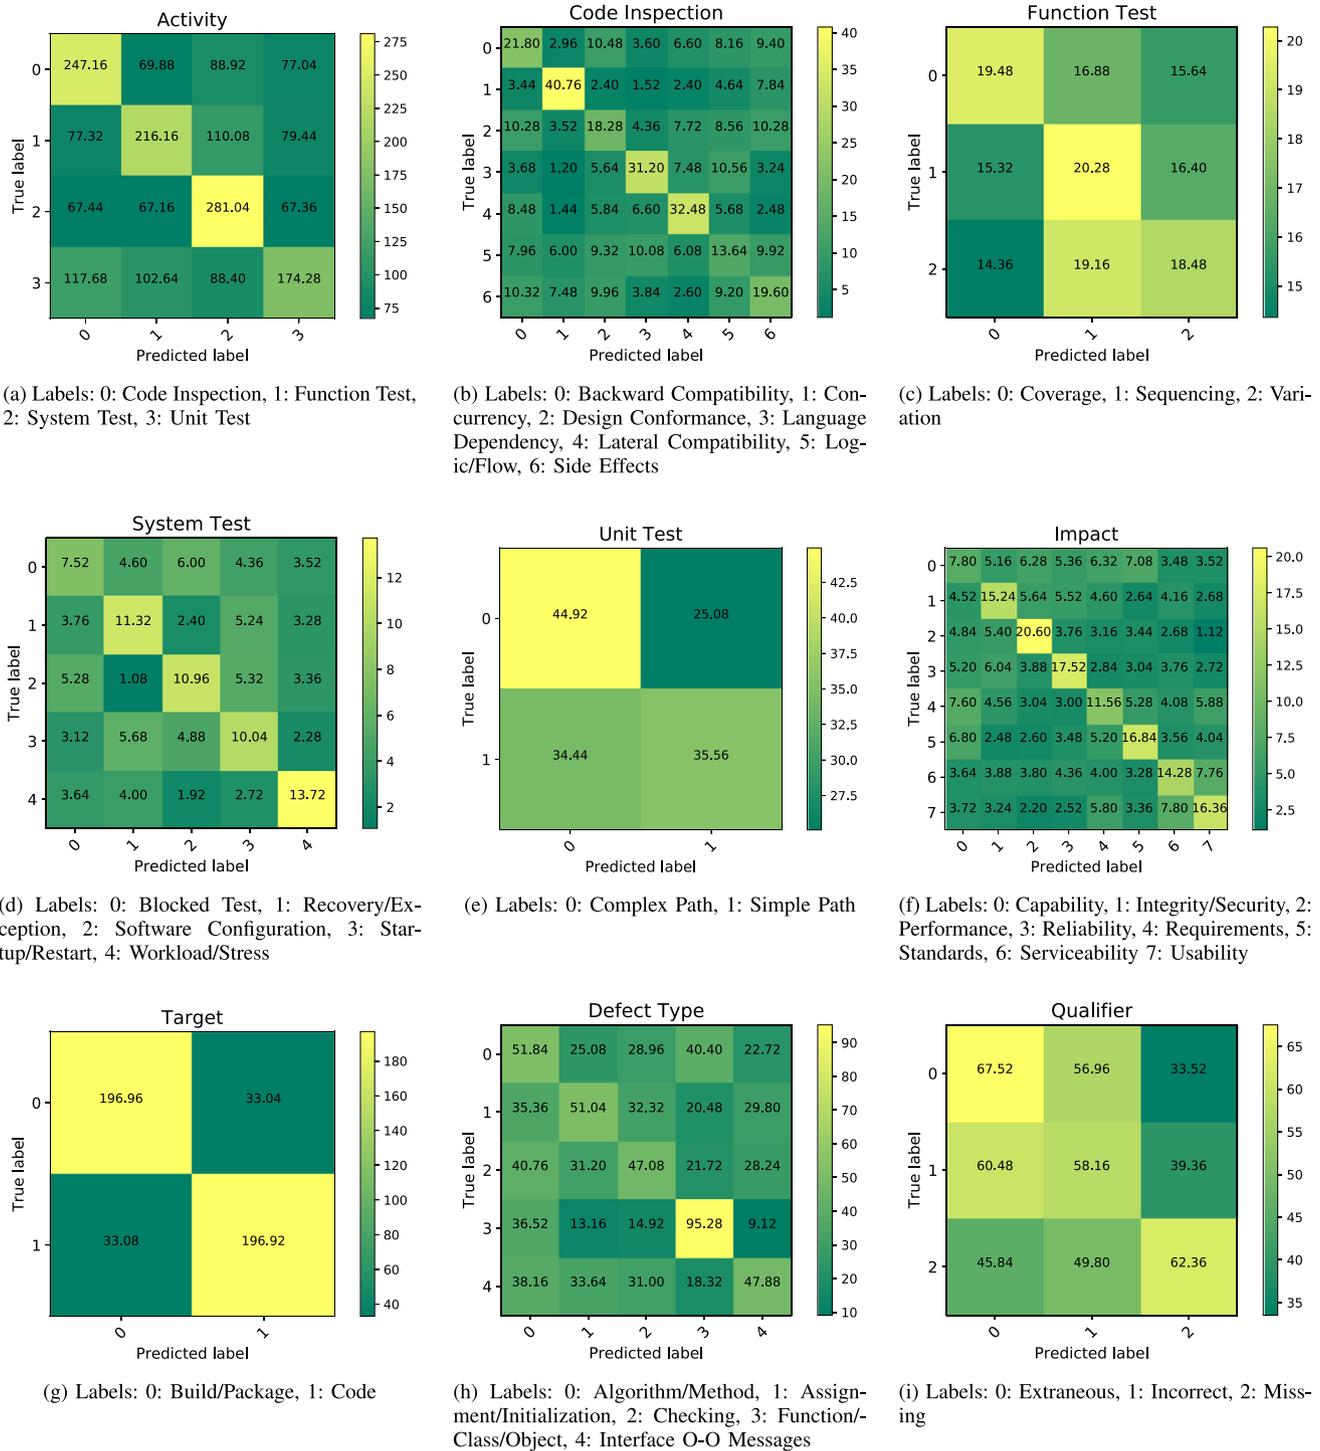
Regarding the Trigger – Function Test results depicted in Fig. 6c, we find that all three classes are confused quite often, although the Sequencing class shows the best true positive results. Function Test-labeled software defects are those which could have been caught through functionality testing (as suggested by the name), and although to the human reader there may be clear differences between the three presented classes (i.e., coverage, sequencing, variation), it is possible that the main keywords referring to each class end up by being quite similar. This was, in

fact, reported by *researcher1* that pointed out a general difficulty in distinguishing among the three classes.

The results for the Trigger – System Test, presented in Fig. 6d, were not quite what we expected since we find the definitions of the System Test triggers to be quite separate, which would make them relatively easy to tell apart, and we can see that several pairs of Triggers were confused by the classifier. A possible explanation for the confusion cases (e.g., mislabeling Blocked Test with Software Configuration) can be the relatively low number of samples (i.e., 130 samples) and, indeed, we have previously observed a general tendency for better results with larger datasets (see Table 6). Workload/Stress showed to be the easiest to identify class, mostly due to the specificities of the keywords related to this class.

In the case of Trigger – Unit Test the classifier shows difficulties in distinguishing both classes involved (see Fig. 6e). The definitions of these two classes (complex path and simple path) both refer to white-box testing and, although they are conceptually different, they do not largely differ from each other [10]). It is easy to imagine a bug report not being sufficiently clear about one of these particular cases. Despite the relatively similar output, the results show that it is easier to identify Complex Path Unit Tests than Simple Path.

Regarding the Impact attribute, shown in Fig. 6f, we can see that the Performance and Reliability classes are the easiest to identify and that, in general, the amount of defects that are confused are low, when compared to the amount of correctly labeled defects. The Capability class is associated with the worst Impact results, which may be explained by the fact that this is a generic class (i.e., it represents the ability of the software to perform its intended functions and satisfy known requirements) that is select when the user is not affected by any of the other Impact types [10]. We have seen that the human classifier used Capability far more than all other Impact classes (i.e., more than 50% of the bugs in the dataset were marked with Capability), and, when this happens, it is common for the machine classifier to be



**Fig. 6.** Detailed classification results per ODC attribute: (a) Activity; (b) Trigger – Code Inspection; (c) Trigger – Function Test; (d) Trigger – System Test; (e) Trigger – Unit Test; (f) Impact; (g) Target; (h) Defect Type; (i) Qualifier.

confused (a possible reason is that the keywords associated with more generic classes overlap with some of the keywords in more concrete classes).

As previously mentioned, the best results were obtained with the Target attribute, due to the fact that the respective classes Code and Build/Package are easily separable. Fig. 6g shows the results in detail.

As for the Defect Type classifier results, which are shown in Fig. 6h, we were actually not expecting good results because,

by definition, this attribute is about the nature of the change that was performed to fix the defect, thus mostly (if not completely) referring to the actual changes in the code. However, we can see that the results for Function/Class/Object are quite good, with the classifier seldom confusing this class with the remaining ones. This Defect Type class refers to large changes in the code, that would require formal design changes [10]. We can also see that, the higher number of errors tends to be related with the erroneous classification of a certain type of defect with

Algorithm/method (leftmost column). This can be explained by the fact that this is, by definition, the most generic class of the Defect Type attribute.

Finally, Fig. 6i shows the detailed results for the Qualifier attribute. The classes associated with the best results are Extraneous and Missing. It was expected that the three classes had similar results since their definitions do not overlap, with each Qualifier referring to a particular state in which the code was (i.e., present but wrong – Incorrect; present but unnecessary – Extraneous). The Incorrect class was the one which the algorithm confused the most, which we know to be a consequence of (as opposed to the other Qualifiers) this particular value having been used for most defects (nearly two thirds of the code defects). In the end, this results in the classifier not easily finding a distinguishable pattern for this particular class.

Overall, regardless of the accuracy values reported, we noted that there are a number of resemblances between the difficulties encountered by the classifier and what the main researcher involved in the creation of the dataset informally reported. The human difficulties would be mostly related with classes that, in practice, become not clearly separable (e.g., simple path and complex path), although they are different by definition. Also, the low quality of some of the reports creates difficulties to the human creation of the dataset, which ends up reflecting in the performance of the classifiers (which have to learn from an erroneous, many times inconsistent, classification). Thus, and although the focus of this work is not on the process taken to manually classify the set of bug reports, we would like to emphasize that the path to achieve higher accuracy values using machine learning classifiers starts with a high quality dataset, which implies a set of actions that, in a pragmatic way, allow us to reach such highly reliable dataset. The difficulties observed simply reflect the reality of handling such complex problem and indicate that there is still space for improvement, namely in the definition of new techniques (e.g., using a composition of different algorithms, mixing machine-learning with other techniques such as rule-based approaches) for automatic classification of bug reports using ODC.

## 5. Threats to validity

In this section, we briefly present threats to the validity of this work and discuss mitigation strategies. An obvious first observation refers to the relatively *small dataset size obtained after balancing*. In fact, the difficulties in obtaining large datasets (and especially considering that some classes will be rare) stem from the lengthy and arduous manual labeling of bugs according to the several ODC attributes. We are currently working on extending the dataset, so that we are able to further improve the classification results. The reduced number of reports lead to the *low dimensionality ratio* problem. We tried to control this issue by performing dimensionality reduction using Principal Component Analysis, while preserving as most as possible the retained variance.

An aspect that will be tackled in future work is related with the fact that some ODC attributes refer to code changes that are many times not properly described in the bug reports. As such, the fact that we are *only using the unstructured textual bug reports* is something that may influence the outcome of the classification (i.e., a code change is objective data, whereas a description of a code change may not be). Thus, in future work, we intend to use techniques that directly analyze the code (and especially, the code changes) as a way to provide additional information to the classifiers. Related with the previous threat, we must also mention that, the outcome of the classification process *depends on the quality of the bug reports* and especially on the consistency

of the text across reports. But the main aspect to retain here is that we are using what can be found in the field, and correcting or individually changing bug reports would not represent a realistic context.

Our *dataset was originally manually annotated (labeled) by one researcher*. Obviously, this activity may have introduced some errors in the process, especially inconsistent labeling of the same type of bugs. To reduce the likelihood of this threat, *researcher1*, the main researcher responsible for the classification, was first trained with 300 bug reports (which actually touched most of the ODC classes used later). In addition, throughout the classification process, we verified 40% of the classified bugs. 20% of the bugs were incrementally verified by *researcher1* and the remaining 20% by two independent researchers (10% each). The goal was to achieve a more consistent classification and also to gather some insight regarding the reliability of the dataset. Due to the amount of effort involved, in practice it becomes inviable to verify the complete dataset. Obviously, this does not remove the mislabeling problem (as two persons may make the same mistake), but it allows to mitigate it.

## 6. Conclusion

This study evaluated the applicability of a set of popular machine learning algorithms (k-Nearest Neighbors, Support Vector Machines, Recurrent Neural Networks, Naïve Bayes, Nearest Centroid, and Random Forest) for performing automatic classification of software defects based on ODC and using unstructured text bug reports as input. The experimental results show the difficulties in automatically classifying certain ODC attributes solely using unstructured textual reports, but also suggest that the overall classification accuracy may be improved if larger datasets are used. Based on our work, we are providing developers and researchers with a public and free RESTful service (available at <http://odc.dei.uc.pt>) for performing automatic classification of bugs according to ODC. The service receives the bug report and the machine learning algorithm to use as input and classifies the bug according the ODC scale. As future work, we intend to extend the dataset used, integrate code changes in the whole process, and also consider a combined approach using distinct machine learning algorithms and different voting schemes in the definition of a new automatic classification process, particularly tailored to handle the ODC specificities.

## 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 work has been partially supported by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 823788 and by the project METRICS, funded by the Portuguese Foundation for Science and Technology (FCT)–agreement no POCI-01-0145-FEDER-032504.

## References

- [1] A. Bouguettaya, M. Singh, M. Huhns, Q.Z. Sheng, H. Dong, Q. Yu, A.G. Neiat, S. Mistry, B. Benatallah, B. Medjahed, M. Ouzzani, F. Casati, X. Liu, H. Wang, D. Georgakopoulos, L. Chen, S. Nepal, Z. Malik, A. Erradi, Y. Wang, B. Blake, S. Dustdar, F. Leymann, M. Papazoglou, A service computing manifesto: The next 10 years, *Commun. ACM* 60 (4) (2017) 64–72, <http://dx.doi.org/10.1145/2983528>.

- [2] M. Grottke, D.S. Kim, R. Mansharamani, M. Nambiar, R. Natella, K.S. Trivedi, Recovery from software failures caused by mandelbugs, *IEEE Trans. Reliab.* 65 (1) (2016) 70–87, <http://dx.doi.org/10.1109/TR.2015.2452933>.
- [3] M. Grottke, K.S. Trivedi, Fighting bugs: remove, retry, replicate, and rejuvenate, *Computer* 40 (2) (2007) 107–109, <http://dx.doi.org/10.1109/MC.2007.55>.
- [4] R. Chillarege, Orthogonal defect classification, in: M.R. Lyu (Ed.), *Handbook of Software Reliability Engineering*, IEEE CS Press, 1996, pp. 359–399.
- [5] L. Huang, V. Ng, I. Persing, M. Chen, Z. Li, R. Geng, J. Tian, Autoodc: Automated generation of orthogonal defect classifications, *Autom. Softw. Eng.* 22 (1) (2015) 3–46.
- [6] F. Thung, D. Lo, L. Jiang, Automatic defect categorization, in: *Reverse Engineering (WCRE), 2012 19th Working Conference on*, IEEE, 2012, pp. 205–214.
- [7] F. Thung, X.-B.D. Le, D. Lo, Active semi-supervised defect categorization, in: *Proceedings of the 2015 IEEE 23rd International Conference on Program Comprehension*, IEEE Press, 2015, pp. 60–70.
- [8] J. Hernández-González, D. Rodríguez, I. Inza, R. Harrison, J.A. Lozano, Learning to classify software defects from crowds: a novel approach, *Appl. Soft Comput.* 62 (2018) 579–591.
- [9] R.T. Fielding, *Architectural Styles and the Design of Network-based Software Architectures* (Ph.D thesis), University of California, Irvine, 2000.
- [10] IBM. (2013, Sep.) Orthogonal Defect Classification v 5.2 for Software Design and Code. [Online]. Available: <https://researcher.watson.ibm.com/researcher/files/us-pasanth/ODC-5-2.pdf>.
- [11] J. Christmansson, R. Chillarege, Generation of an error set that emulates software faults based on field data, in: *Fault Tolerant Computing, 1996., Proceedings of Annual Symposium on*, IEEE, 1996, pp. 304–313.
- [12] J. Durães, H. Madeira, Definition of software fault emulation operators: a field data study, in: *2003 International Conference on Dependable Systems and Networks, 2003. Proceedings., San Francisco, CA, USA, 2003*, pp. 105–114.
- [13] A. Avizienis, J.-C. Laprie, B. Randell, C. Landwehr, Basic concepts and taxonomy of dependable and secure computing, *IEEE Trans. Depend. Secure Comput.* 1 (1) (2004) 11–33.
- [14] S. Kumaresh, R. Baskaran, Defect analysis and prevention for software process quality improvement, *Int. J. Comput. Appl.* 8 (2010).
- [15] L. Zhi-bo, H. Xue-mei, Y. Lei, D. Zhu-ping, X. Bing, Analysis of software process effectiveness based on orthogonal defect classification, *Procedia Environ. Sci.* 10 (2011) 765–770.
- [16] S. Bellucci, B. Portaluri, Automatic calculation of orthogonal defect classification (ODC) fields, U.S. Patent 8 214 798, Jul., 2012. [Online]. <http://www.freepatentsonline.com/8214798.html>.
- [17] C. Liu, Y. Zhao, Y. Yang, H. Lu, Y. Zhou, B. Xu, An ast-based approach to classifying defects, in: *Software Quality, Reliability and Security-Companion (QRS-C), 2015 IEEE International Conference on*, IEEE, 2015, pp. 14–21.
- [18] D. Cubranic, G. Murphy, Automatic bug triage using text categorization, in: *Proceedings of the Sixteenth International Conference on Software Engineering & Knowledge Engineering*, Citeseer, 2004.
- [19] Y. Tian, D. Lo, C. Sun, Drone: Predicting priority of reported bugs by multi-factor analysis, in: *Software Maintenance (ICSM), 2013 29th IEEE International Conference on*, IEEE, 2013, pp. 200–209.
- [20] B.V. Dedik, *Automatic Ticket Triage Using Supervised Text Classification*, Masaryk University, Faculty of Informatics, Brno, 2015.
- [21] J. Xuan, H. Jiang, Z. Ren, J. Yan, Z. Luo, Automatic Bug Triage using Semi-Supervised Text Classification, *CoRR* vol. abs/1704.04769, 2017.
- [22] T. Menzies, A. Marcus, Automated severity assessment of software defect reports, in: *Software Maintenance, 2008. ICSM 2008. IEEE International Conference on*, IEEE, 2008, pp. 346–355.
- [23] J.A. Durães, H.S. Madeira, Emulation of software faults: A field data study and a practical approach, *IEEE Trans. Softw. Eng.* 32 (11) (2006).
- [24] T. Sotiropoulos, H. Waeselynck, J. Guiochet, F. Ingrand, Can robot navigation bugs be found in simulation? an exploratory study, in: *2017 IEEE International Conference on Software Quality, Reliability and Security (QRS)*, 2017, pp. 150–159, <http://dx.doi.org/10.1109/QRS.2017.25>.
- [25] A. Rahman, S. Elder, F.H. Shezan, V. Frost, J. Stallings, L. Williams, Categorizing Defects in Infrastructure as Code, *arXiv:1809.07937 [cs]*, Sep. 2018, *arXiv: 1809.07937*. [Online]. Available: <http://arxiv.org/abs/1809.07937>.
- [26] J. Cohen, A coefficient of agreement for nominal scales, *Educ. Psychol. Meas.* 20 (1) (1960) 37–46, <http://dx.doi.org/10.1177/001316446002000104>.
- [27] J.R. Landis, G.G. Koch, The measurement of observer agreement for categorical data, *Biometrics* 33 (1) (1977) 159–174, [Online]. Available: <https://www.jstor.org/stable/2529310>.
- [28] C.D. Manning, P. Raghavan, H. Schütze, *Scoring, term weighting, and the vector space model*, in: *Introduction to Information Retrieval*, Cambridge University Press, 2008, pp. 100–123.
- [29] A. Rajaraman, J.D. Ullman, *Mining of Massive Datasets*, Cambridge University Press, New York, NY, USA, 2011.
- [30] T. Joachims, Text categorization with support vector machines: Learning with many relevant features, in: *Proceedings of the 10th European Conference on Machine Learning*, in: *ECML '98*, Springer-Verlag, London, UK, UK, 1998, pp. 137–142, <http://dl.acm.org/citation.cfm?id=645326.649721>.
- [31] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [32] T. Archanah, D. Sachin, Dimensionality reduction and classification through pca and lda, *Int. J. Comput. Appl.* 122 (17) (2015) 4–8, <http://dx.doi.org/10.5120/21790-5104>.
- [33] J. de Sa, *Pattern Recognition: Concepts, Methods, and Applications*, Springer, 2001, [Online]. Available: <https://books.google.pt/books?id=O5vwpjQqWlC>.
- [34] A. Rea, W. Rea, How Many Components should be Retained from a Multivariate Time Series PCA?, oct 2016. [Online]. Available: *arXiv:1610.03588*.
- [35] N.S. Altman, An introduction to kernel and nearest-neighbor nonparametric regression, *Amer. Statist.* 46 (3) (1992) 175–185.
- [36] A. McCallum, K. Nigam, et al., A comparison of event models for naive bayes text classification, in: *AAAI-98 Workshop on Learning for Text Categorization*, Vol. 752, Citeseer, 1998, pp. 41–48, no. 1.
- [37] C. Cortes, V. Vapnik, Support-vector networks, *Mach. Learn.* 20 (3) (1995) 273–297.
- [38] J.M. De Sa, *Pattern recognition: concepts, methods and applications*, Springer Science & Business Media, 2001, <http://dx.doi.org/10.1007/978-3-642-56651-6>.
- [39] K.-B. Duan, S. Keerthi, Which is the best multiclass SVM method? an empirical study, in: *International Workshop on Multiple Classifier Systems*, Springer, 2005, pp. 278–285.
- [40] L. Breiman, Random forests, *Mach. Learn.* 45 (1) (2001) 5–32, ch. 8.
- [41] L. Breiman, Bagging predictors, *Mach. Learn.* 24 (2) (1996) 123–140.
- [42] D. Tang, B. Qin, T. Liu, Document modeling with gated recurrent neural network for sentiment classification, in: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Lisbon, Portugal, 2015, pp. 1422–1432, <http://dx.doi.org/10.18653/v1/D15-1167>, <https://www.aclweb.org/anthology/D15-1167>.
- [43] P. Liu, X. Qiu, X. Huang, Recurrent neural network for text classification with multi-task learning, in: *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, in: *IJCAI'16, AAAI Press*, 2016, pp. 2873–2879, <http://dl.acm.org/citation.cfm?id=3060832.3061023>.
- [44] Y. Bengio, P. Simard, P. Frasconi, et al., Learning long-term dependencies with gradient descent is difficult, *IEEE Trans. Neural Netw.* 5 (2) (1994) 157–166.
- [45] I. Goodfellow, Y. Bengio, A. Courville, Optimization for training deep models, in: *Deep Learning*, MIT Press, 2016, pp. 267–320.
- [46] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780.
- [47] A. Graves, J. Schmidhuber, Framewise phoneme classification with bidirectional lstm and other neural network architectures, *Neural Netw.* 18 (5) (2005) 602–610, *IJCNN 2005*.
- [48] A. Graves, N. Jaitly, A.-r. Mohamed, Hybrid speech recognition with deep bidirectional LSTM, in: *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, IEEE, 2013, pp. 273–278.
- [49] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient Estimation of Word Representations in Vector Space, *arXiv preprint arXiv:1301.3781*, pp. 1–12, 2013.
- [50] F. Chollet, et al., Keras, <https://github.com/fchollet/keras>, 2015.
- [51] C. Zhang, S. Bengio, M. Hardt, B. Recht, O. Vinyals, Understanding Deep Learning Requires Rethinking Generalization, *arXiv preprint arXiv:1611.03530*, pp. 115, 2016.
- [52] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (1) (2014) 1929–1958.
- [53] R. Kohavi, A study of cross-validation and bootstrap for accuracy estimation and model selection, in: *Proceedings of the 14th International Joint Conference on Artificial Intelligence - Volume 2*, in: *IJCAI'95*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1995, pp. 1137–1143, <http://dl.acm.org/citation.cfm?id=1643031.1643047>.
- [54] C. Beleites, U. Neugebauer, T. Bocklitz, C. Krafft, J. Popp, Sample size planning for classification models, *Anal. Chim. Acta* 760 (2013) 25–33, <http://www.sciencedirect.com/science/article/pii/S0003267012016479>.



**Fábio Lopes** is a PhD student at DEI-Faculty of Sciences and Technology of the University of Coimbra. His Master's Thesis topic was in the field of Natural Language Processing, in particular Named Entity Recognition. His research interests include artificial intelligence, machine learning and natural language processing.



**João Agnelo** is an MSc student in Informatics Engineering at the Faculty of Sciences and Technology of the University of Coimbra. He concluded his BSc in Informatics Engineering at ISEC - Polytechnic Institute of Coimbra, in 2018. Alongside his studies, he is also currently working in a few research projects at CISUC – Centre for Informatics and Systems of the University of Coimbra. His interests include dependability, artificial intelligence and computer graphics.



**César A. Teixeira** is an associate professor at DEI-FCTUC, University of Coimbra, Portugal. He has a graduation (2003) in Systems and Computation Engineering and a PhD (2008) in Electronics Engineering and Informatics both from the University of the Algarve, Portugal. His expertise is on bio-signal processing, classification and modeling. More precisely on EEG-based epileptic seizure prediction and anesthesia monitorin. His publication record includes over 80 papers in refereed national and international journals and conferences. He participated in several

national (HeartSafe(PTDC EEI-PRO 2857 2012), and iCIS (CENTRO-07-0224-FEDER-002003)) and international projects (EPILEPSIAE (FP7-ICT-2007-211713), Welcome(FP7-611223), and Link (H2020-692023)).



**Nuno Laranjeiro** received the PhD degree in 2012 from the University of Coimbra, Portugal, where he currently is an assistant professor. His research focuses on robust software services as well as experimental dependability evaluation, web services interoperability, services security, and enterprise application integration. He has authored more than 50 papers in refereed conferences and journals in the dependability and services computing areas and participated in several national and international projects.



**Jorge Bernardino** is a Coordinator Professor at the Polytechnic of Coimbra - ISEC, Portugal. He received the PhD degree from the University of Coimbra in 2002. His main research interests are Big Data, NoSQL, Data Warehousing, Software Engineering, and experimental databases evaluation. He has authored over 200 publications in refereed journals, book chapters, and conferences. He was President of ISEC from 2005 to 2010 and President of ISEC Scientific Council from 2017 to 2019. During 2014 he was Visiting Professor at CMU - Carnegie Mellon University, USA. He participated in several national and international projects and is an ACM and IEEE member. Currently, he is Director of Applied Research Institute of the Polytechnic of Coimbra.