

Software Defect Detection with ROCUS

Yuan Jiang (姜远), *Member, CCF*, Ming Li (黎铭)*, *Member, CCF, ACM, IEEE*
and Zhi-Hua Zhou (周志华), *Senior Member, CCF, IEEE, Member, ACM*

National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210093, China

E-mail: {jiangyuan, lim, zhoush}@nju.edu.cn

Received May 15, 2009; revised October 26, 2010.

Abstract Software defect detection aims to automatically identify defective software modules for efficient software test in order to improve the quality of a software system. Although many machine learning methods have been successfully applied to the task, most of them fail to consider two practical yet important issues in software defect detection. First, it is rather difficult to collect a large amount of labeled training data for learning a well-performing model; second, in a software system there are usually much fewer defective modules than defect-free modules, so learning would have to be conducted over an imbalanced data set. In this paper, we address these two practical issues simultaneously by proposing a novel semi-supervised learning approach named ROCUS. This method exploits the abundant unlabeled examples to improve the detection accuracy, as well as employs under-sampling to tackle the class-imbalance problem in the learning process. Experimental results of real-world software defect detection tasks show that ROCUS is effective for software defect detection. Its performance is better than a semi-supervised learning method that ignores the class-imbalance nature of the task and a class-imbalance learning method that does not make effective use of unlabeled data.

Keywords machine learning, data mining, semi-supervised learning, class-imbalance, software defect detection

1 Introduction

Enabled by technological advances in computer hardware, software systems have become increasingly powerful and versatile. However, the attendant increase in software complexity has made the timely development of reliable software systems extremely challenging. To make software systems reliable, it is very important to identify as many defects as possible before releasing the software. However, due to the complexity of the software systems and the tight project schedule, it is almost impossible to extensively test every path of the software under all possible runtime environment. Thus, accurately predicting whether a software module contains defects can help to allocate the limited test resources effectively, and hence, improve the quality of software systems. Such a process is usually referred to as *software defect detection*, which has already drawn much attention in software engineering community.

Machine learning techniques have been successfully applied to building predictive models for software defect detection^[1-7]. The static and dynamic code attributes

or software metrics are extracted from each software module to form an example, which is then labeled as “defective” or “defect-free”. Predictive models which learn from a large number of examples are expected to accurately predict whether a given module is defective.

However, most of these studies have not considered two practical yet important issues in software defect detection. First, although it is relatively easy to automatically generate examples from software modules using some standard tools, determining whether a module contains defect through extensive test usually consumes too much time and resource, since the number of program status grows exponentially as the complexity of software increases. With limited time and test resource, one can only obtain the labels for a small portion of modules. However, the predictive models that learn from such a small labeled training set may not perform well. Second, the data in software defect detection are essentially imbalanced. The number of defective modules is usually much less than that of the defect-free modules. Ignoring the imbalance nature of the problem, a learner that minimizes the prediction error can

Regular Paper

This work was supported by the National Natural Science Foundation of China under Grant Nos. 60975043, 60903103, and 60721002.

*Corresponding Author

©2011 Springer Science + Business Media, LLC & Science Press, China

often produce a useless predictive model that uniformly predicts *all* the modules as defect-free. Without taking these two issues into consideration, the effectiveness of software defect detection in many real-world tasks would be greatly reduced.

Some researchers have noticed the importance of these two issues in software defect detection and tried to tackle some of them based on machine learning techniques. For instance, Seliya and Khoshgoftaar^[8] employed *semi-supervised learning* to improve the performance achieved on a small amount of labeled data by exploiting the abundant unlabeled data; contrarily, Pelayo and Dick^[9] applied the resampling strategy to balance the skewed class distribution of the dataset before learning the predictive model for software defect detection. Although attempting to tackle one issue may gain performance improvement to some extent, both methods suffer the influence of the other issue that they have not considered. If conventional semi-supervised learning is used, assuming that the learner can accurately assign labels to the unlabeled data, the learner may be easily biased by the overwhelming number of newly-labeled defect-free modules, and hence the refined model would be less sensitive to the defect modules. The sensitivity drops fast as the iterative semi-supervised learning proceeds. On the other hand, resampling methods would become less effective if provided with only a few labeled examples, where overfitting is inevitable no matter when it replicates the small number of defective examples or reduces the number of overwhelming defect-free examples. Therefore, to achieve effective software defect detection, we need to consider these two important issues *simultaneously*. To the best of our knowledge, there is no previous work that has considered these two issues simultaneously in software defect detection.

In this paper, we address aforementioned two issues by proposing a novel semi-supervised learning method named ROCUS (RandOm Committee with Under-Sampling). This method incorporates recent advances in disagreement-based semi-supervised learning^[10] with under-sampling strategy^[11] for imbalanced data. The key idea is to keep the individual learner focusing on the minority-class during exploitation of the unlabeled data. Experiments on eight real-world software defect detection tasks show that ROCUS is effective for software defect detection. Its performance is better than both the semi-supervised learning method that ignores the class-imbalance nature of the tasks and the class-imbalance learning method that does not exploit unlabeled data.

The rest of the paper is organized as follows. Section 2 briefly reviews some related work. Section 3 presents

the ROCUS method. Section 4 reports the experiments over the software defect detection tasks. Finally, Section 5 concludes this paper.

2 Related Work

2.1 Software Defect Detection

Software defect detection, which aims to automatically identify the software module that contains certain defects, is essential to software quality insurance. Most of the software detection methods roughly fall into two categories. Methods in the first category leverage the execution information to identify suspicious program behaviors for defect detection^[12-14], while methods in the second category elaborate to extract static code properties, which are usually represented by a set of *software metrics*, for each module in the software system^[7,15-16]. Since it would be easier to measure the static code properties than measure the dynamic program behaviors, metric-based software defect detection has drawn much attention. Widely-used software metrics include LOC counts describing the module in terms of size, Halstead attributes measuring the number of operators and operands in the module as the reading complexity^[17] and McCabe complexity measures derived from the flow graph of the module^[18].

In the past decade, machine learning has been widely applied to construct predictive models based on the extracted software metrics to detect defects in the software modules. Typical methods include linear or logistic regression^[7,15], classification and regression trees^[3,19], artificial neural networks^[1,16], memory-based methods^[20-21] and Bayesian methods^[22-23]. In order to further increase the robustness to the outlier in the training data and improve the prediction performance, Guo *et al.*^[2] applied ensemble learning to the software defect detection and achieved better performance compared to other commonly-used methods such as logistic regression and decision tree. Recently, Lessmann *et al.*^[4] conducted an intensive empirical study, where they compared the predictive performance of 22 machine learning methods over the benchmark data sets of software defect detection.

Note that few previous studies have ever considered the characteristics of software defect detection (e.g., the defective module is difficult to collect). It has been showed that more accurate detection could be achieved even if only one of such characteristics is considered during learning^[8-9]. Since these characteristics are usually intertwined with each other, better performance could be expected if carefully considering them together during learning, which is what we do in this paper.

2.2 Semi-Supervised Learning

In many practical applications, many unlabeled data can be easily collected, while only a few labeled data can be obtained since much human effort and expertise are required. Semi-supervised learning^[24-25] is a machine learning technique where the learner automatically exploits the large amount of unlabeled data in addition to few labeled data to help improving the learning performance.

Generally, semi-supervised learning methods fall into four major categories, i.e., generative-model-based methods^[26-28], low density separation based methods^[29-31], graph-based methods^[32-34], and disagreement-based methods^[35-39].

Disagreement-based methods use multiple learners and exploit the disagreements among the learners during the learning process. If majority learners are much more confident of a disagreed unlabeled example than minority learner(s), then the majority will teach the minority of this example. Disagreement-based semi-supervised learning originates from the work of Blum and Mitchell^[35], where classifiers learned from two sufficient and redundant views teach each other using some confidently predicted unlabeled examples. Later, Goldman and Zhou^[36] proposed an algorithm which does not require two views but require two different learning algorithms. Zhou and Li^[38] proposed using three classifiers to exploit unlabeled data, where an unlabeled example is labeled and used to teach one classifier if the other two classifiers agree on its labeling. Later, Li and Zhou^[37] further extended the idea in [38] by collaborating more classifiers in training process. Besides classification, Zhou and Li^[39] also adapted the disagreement-based paradigm to semi-supervised regression. Disagreement-based semi-supervised learning paradigm has been widely applied to natural language processing (e.g., [40]), information retrieval (e.g., [41-42]), computer-aided diagnosis (e.g., [37]), etc.

Few researches applied semi-supervised learning to software defect detection, where the labeled training examples are limited while the unlabeled examples are abundant. Recently, Seliya and Khoshgoftaar^[8] applied a generative-model-based semi-supervised learning method to software defect detection and achieved performance improvement. Note that [8] adopted a generative approach for exploiting unlabeled data while the proposed method adopts a discriminative approach. Thus, we did not include it in our empirical study for the purpose of fair comparison.

2.3 Learning from Imbalanced Data

In many real-world applications such as software

defect detection, the class distribution of the data is imbalanced, that is, the examples from the minority class are (much) fewer than those from the other class. Since it is easy to achieve good performance by keeping the majority-class examples being classified correctly, the sensitivity of the classifiers to the minority class may be very low if directly learning from the imbalanced data. To achieve better sensitivity to the minority class, the class-imbalance problem should be explicitly tackled.

Popular class-imbalance learning techniques include sampling^[11,43-44] and cost-sensitive learning^[45-46]. Since sampling technique is used in this paper, we introduce sampling in more details.

Sampling attempts to achieve a balanced class distribution by altering the dataset. Under-sampling reduces the number of the majority-class examples while over-sampling increases the number of minority-class examples^[11], both of which have been shown to be effective to class-imbalance problems. Sophisticated methods can be employed to balance the class distribution, such as adding synthetic minority-class examples generated from the interpolation of neighboring minority-class examples^[43]; discarding the non-representative majority-class examples to balance the class distribution^[44]; combining different sampling methods for further improvement^[47]; using ensemble technique for exploratory under-sampling to avoid the removal of useful majority class examples^[48].

The class-imbalance learning method is seldom used in software defect detection. Recently, Pelayo and Dick^[9] studied the effectiveness of SMOTE^[43] over the software defect detection, and found that balancing the skewed class distribution is beneficial to software defect detection.

3 Proposed Approach

Let $L = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_{m_0}, y_{m_0})\}$ denote the set of labeled examples and let $U = \{\mathbf{x}_{m_0+1}, \mathbf{x}_{m_0+2}, \dots, \mathbf{x}_N\}$ denote the set of unlabeled examples, where \mathbf{x}_i is a d -dimensional feature vector, and $y_i \in \{-1, +1\}$ is the class label. Conventionally, $+1$ denotes the minority class (e.g., “defective” in software defect detection). Thereinafter, we refer to class $+1$ as the *minority-class* and -1 as the *majority-class*. Both L and U are independently drawn from the same unknown distribution \mathcal{D} whose marginal distributions satisfy $P_{\mathcal{D}}(y_i = +1) \ll P_{\mathcal{D}}(y_i = -1)$, and hence, L and U are imbalanced datasets in essence.

As mentioned in Section 1, directly applying semi-supervised learning to imbalanced data would be risky. Since L is imbalanced and usually small, very few examples of the minority-class would be used to initiate

the semi-supervised learning process. The resulting model may have poor sensitivity to the minority-class and hence can hardly identify the examples of the minority-class from the unlabeled set. In this case, learner would have to use little information from the minority-class and overwhelming information of the majority-class for model refinement, and this leads to even poorer sensitivity to the minority-class. As the iterative semi-supervised learning proceeds, the learned model would be biased to predict every example to the majority-class.

In order to successfully conduct iterative semi-supervised learning on the imbalanced data, the learner should have the following two properties. First, the learner should have strong generalization ability, such that even if provided with a small labeled training set with imbalanced class distribution, the learner would not have zero sensitivity to the minority-class examples during the automatically labeling process; second, the influence of overwhelming number of the newly labeled majority-class examples should be further reduced in order to improve the sensitivity of the learner to the minority examples after its refinement in each learning iteration. Based on these two considerations, we propose the ROCUS method to exploit the imbalanced unlabeled examples.

To meet the first requirement, we train multiple classifiers and then combine them for prediction. The reason behind this specific choice of the *ensemble learning* paradigm is that an ensemble of classifiers can usually achieve better generalization performance than a single classifier. Such superiority is more obvious when the training set is small^[37] and the class distribution is imbalanced^[48]. Thus, by exploiting the generalization power, the trained ensemble from L is able to identify some minority-class examples from U effectively.

Since multiple classifiers are used, we employ the disagreement-based semi-supervised learning paradigm^[10] to exploit the unlabeled examples in U . In detail, after the initial ensemble of classifiers $\{h_1, h_2, \dots, h_C\}$ are constructed, some individual classifiers select some examples in U to label according to a disagreement level, and then teach the other classifiers with the newly labeled examples. Here, similar to [37], we adopt a simple case where the classifiers $H_{-i} = \{h_1, \dots, h_{i-1}, h_{i+1}, \dots, h_C\}$ are responsible for selecting confidently labeled unlabeled examples in U for an individual classifier h_i . Given an unlabeled example, we first label this example using the majority voting of $C-1$ individual classifiers, and then estimate the labeling confidence using the degree of agreement on the current labeling among these classifiers. If the labeling confidence is greater than that of a pre-set threshold θ , we

feed h_i with this newly labeled example for its refinement. Inspired by [27], we associate a weight (between 0 and 1) with each unlabeled example according to its labeling confidence such that the contribution of those less confidently labeled examples will be reduced during the classifier refinement. To unify the representation, the weight of a labeled example is fixed to 1.

Note that even if the ensemble of classifiers can provide accurate prediction for each selected example, the sensitivity of the current classifier h_i may not be improved after its refinement with these newly labeled examples. Since U itself is imbalanced, h_i would still lose its sensitivity to the minority-class after learning many newly-labeled majority-class examples in U . Here, we employ under-sampling^[11], an efficient strategy for class-imbalance learning, to tackle this problem. Specifically, let \tilde{L}_{it} denote the newly labeled set in the t -th round, where the total weights of the minority-class examples in \tilde{L}_{it} are p_{it} ; let L'_{it} denote the corresponding under-sampled set, where the total weights of the minority-class examples are p_{it} and that of majority-class examples are p_{it}/γ . Here, γ specifies the expected ratio between the minority-class and the majority-class and is usually fixed to 1 to produce a balanced class distribution.

However, since under-sampling reduces the number of newly labeled examples, even a little misclassifications of unlabeled examples can greatly increase the noise rate of the newly labeled sets. Learning on noisy dataset may humble the performance of resulting classifier. Thus, we stop the iterative semi-supervised learning process if the datasets used for the classifier refinement become too noisy to improve the performance of the classifier. The relationship between the classifier's worst-case error and noise rate has been studied in [49], and has been applied to deriving the stopping criterion for some disagreement-based semi-supervised learning methods^[37-38,41]. Here, our derivation of stopping criterion is almost the same as that of these methods. For the self-containment of this paper, we include the derivation below.

W_0 and $W_{i,t}$ denote the total weights of examples in L and $L'_{i,t}$, respectively. Let $\hat{e}_{i,t}$ denote the estimated error rate of H_{-i} , the ensemble of classifiers excluding h_i . Assume that the noise rate of original labeled set L is very small, and hence the noise rate in the augmented training set, i.e., $L \cup L'_{i,t}$ for refining h_i in t -th iteration can be estimated by

$$\eta_{i,t} = \frac{\hat{e}_{i,t}W_{i,t}}{W_0 + W_{i,t}}. \quad (1)$$

We can define the utility function of the refinement

of classifier h_i in the t -th round based on the relationship between h_i 's worst-case error $\epsilon_{i,t}$, the (weighted) number of examples in the augmented training set ($W_0 + W_{i,t}$) and its noise rate $\eta_{i,t}$ as

$$u_{i,t} \equiv \frac{c}{\epsilon_{i,t}^2} = (W_0 + W_{i,t})(1 - 2\eta_{i,t})^2 \quad (2)$$

where c is a constant that makes the equality hold.

Since $u_{i,t}$ is inverse-proportional to the square of the worse case classification error $\epsilon_{i,t}$, by enforcing $u_{i,t} > u_{i,t-1}$ in the succeeded rounds, the performance of h_i will be improved after refinement in the t -th round. By comparing the right hand side of (2) in the t -th round and $(t-1)$ -th round, we have that $u_{i,t} > u_{i,t-1}$ holds if

$$0 < \frac{\hat{\epsilon}_{i,t}}{\hat{\epsilon}_{i,t-1}} < \frac{W_{i,t-1}}{W_{i,t}} < 1. \quad (3)$$

In some cases, (3) might not hold since $W_{i,t}$ may be much greater than $W_{i,t-1}$. To make (3) hold again, we randomly discard some examples of both the minority-class and majority-class according to γ .

Note that we do not apply undersampling to L . The rationality behind is that undersampling L will further decrease the number of the original labeled examples which are more reliable than the automatically labeled examples. The benefit from balancing the training data might be counteracted by discarding many reliably labeled data. However, in this case, the augmented training set for classifier refinement may be slightly imbalanced. In order to compensate for the effect caused by such imbalance, we rescale the output of each classifier h_i using the minority-majority ratio of current training examples r_i and then combine them for final prediction:

$$H^*(x) = \begin{cases} +1, & \frac{1}{C} \sum_i \frac{s(+1|h_i(x))}{s(+1|h_i(x)) + r_i s(-1|h_i(x))} > 0.5 \\ -1, & \text{otherwise,} \end{cases} \quad (4)$$

where $s(y|h(x)) \in [0, 1]$ gives a score for the prediction of x . If the score is greater than 0.5, x is predicted as “+1”, and “-1” otherwise.

As pointed out by Li and Zhou^[37], such a majority-teach-one style process may gradually reduce the diversity between individual classifiers. Although the performance of individual classifiers can be improved through the semi-supervised learning process, the performance of the ensemble may not be improved or even degrade due to the rapid decrease of diversity. The reason is that the “teachers” of two individual classifiers are quite similar, and thus the newly labeled set of these two classifiers would be similar. Refining over similar datasets makes these two classifiers more

similar. Following the suggestion of Li and Zhou^[37], we inject certain amount of randomness into the base learner such that even if newly labeled examples are similar, the learned classifiers can still be different. Here, we call the ensemble of randomized classifiers *random committee*. If we use random tree inducer that generates a randomized decision tree as the base learner, the ensemble is equivalent to what used in [37]. In this paper, we adopt another approach for randomness injection. Specifically, we project the data onto a set of randomly generated unit vectors and construct a classifier in this new space. We repeat this process to achieve a number of new classifiers. The dimensionality of the new space is usually smaller than the original one in order to achieve further diversity between different new spaces. Similar approach was used by Ho^[50] for constructing ensemble, where the random unit vectors are enforced to be parallel to the basis of the original space.

Table 1 presents the pseudo code of ROCUS. We firstly construct C classifiers from L by using Bagging^[51] with the base learning algorithm \mathcal{A} . Any learning algorithm that incorporates certain randomness may be used to instantiate \mathcal{A} . In this paper, \mathcal{A} injects randomness by conducting random projection before learning a classifier. In each semi-supervised learning iteration, each classifier h_i is refined using the newly labeled examples selected by H_{-i} , the ensemble

Table 1. Pseudo-Code of the ROCUS Algorithm

Algorithm. ROCUS

Input:

- the labeled set L , the unlabeled set U ,
- the confidence threshold θ , the minority-majority ratio γ ,
- the number of individual classifiers C ,
- the base learning algorithm \mathcal{A} of the random committee

Process:

1. Learn a random committee $\{h_1, \dots, h_C\}$ from L using Bagging and the base learning algorithm \mathcal{A}
2. Repeat Steps 3~9 until none of the classifier in the random committee changes
3. Set t ($t \in \mathbb{N}$) as the current iteration number
4. For each $i \in \{1, \dots, C\}$, do Steps 4~9
5. Estimate error $e_{i,t}$ of H_{-i} on L
6. Label all the unlabeled examples with H_{-i}
7. Add the unlabeled examples whose labeling confidence exceeds threshold θ to a newly labeled set $L'_{i,t}$
8. Undersample $L'_{i,t}$ such that the ratio of minority class over the majority class is no less than γ
9. If (3) holds, retrain h_i from $L \cup L'_{i,t}$ using the learning algorithm \mathcal{A}

Output: Compute $H^*(x)$ according to (4)

of classifiers other than h_i . Before the refinement, under-sampling is employed to tailor the newly labeled set such that its minority-majority ratio is roughly γ . We use the condition in (3) as the stopping criterion of the iterative learning process. As holding a separate validation set is infeasible in semi-supervised learning settings, the error rate $\hat{e}_{i,t}$ of H_{-i} is estimated on L under an assumption that the training data and test data have the same distribution.

Note that an alternative way to address the problems of “lack of sufficient labeled data” and “data imbalance” simultaneously by imposing a “class proportion” constraint over a special type of base learner, which can adjust the portion of labeling of unlabeled data according to the constraint, just as what TSVM^[31] does. However, such a strategy may exclude many good candidate base learners that have good performance over some particular defect detection problems but fail to adjust their labeling according to the constraints. In contrast, by incorporating under-sampling, disagreement-based semi-supervised learning method can be easily adapted to the exploitation of unlabeled data while the data are imbalanced. Since the requirement of the base learner in ROCUS is no more than the ability of injecting randomness, which can be easily achieved, we may choose different base learners according to specific application scenario, and hence applicability of ROCUS will be better.

4 Empirical Studies

We evaluate the effectiveness of ROCUS on eight software defect detection benchmark tasks. Each dataset corresponds to different software projects in NASA Metrics Data Program^[52]. Some of these software projects are developed for satellite flight control, while others are used for the ground-system. All the software projects are written in C/C++. Each project consists of a number of software modules, each of which is manually labeled as “defective” if one or more defects were reported during the test phase and “defect-

free” otherwise. Typical software metrics such as *LOC counts*, *McCabe complexity measures* derived from the pathway of modules, *Halstead attributes* measuring the number of operators and operands in the module as the reading complexity, are extracted from each software module using some standard code analysis tools^①. The detailed information of all the software metrics used in the current study can be found in [5]. The detailed information of the experimental datasets are tabulated in Table 2, where Ratio denotes the inverse of the minority-majority ratio of the datasets. It is obvious from the table that the datasets are imbalanced and the number of defective modules are smaller than that of the modules without any defect.

For each dataset, we randomly select 75% examples for training and keep the the remaining examples aside for test. Since all the examples in the training set are labeled, in order to simulate the case where only a small portion of training data are labeled, we randomly partition the training set into labeled and unlabeled sets according to a *labeled rate* μ . For example, if a training set consisting of 1000 examples and $\mu = 10\%$, 100 examples are put into the labeled set with their labels, while the remaining 900 examples are put into the unlabeled set without their labels. In the experiment, we use four different labeled rates: 10%, 20%, 30% and 40%.

In the experiments, the randomized base learner of ROCUS is instantiated as an ADABOOST^[53] preceded by a random projector. The random projector first randomly generates $2d/3$ random unit vectors and project all the examples onto these random vectors. Here, d is the number of features of the dataset. Following the suggestions of [37], we fix the size of the random committee to $C = 6$, and the confidence threshold θ is set to 0.75, which indicates an unlabeled example is regarded to be confidently labeled if more than 3/4 individual classifiers of the random committee agree on its labeling. We set the minority-majority ratio $\gamma = 1$ to enforce the newly labeled example set to be balanced.

Table 2. Software Defect Detection Datasets

Data	No. Attr.	No. Inst.	No. Min./No. Maj.	Ratio	Description
<i>jm1</i>	21	10 885	2106/8779	4.2	A real-time predictive ground system
<i>kc1</i>	21	2 109	326/1783	5.5	A storage management system for receiving and processing ground data
<i>kc2</i>	21	522	107/415	3.9	Another part of <i>kc1</i> project for science data processing
<i>mw1</i>	37	403	31/372	12.0	A zero gravity experiment system related to combustion
<i>pc1</i>	21	1 109	77/1032	13.4	Flight software for an earth orbiting satellite
<i>pc3</i>	37	1 563	160/1403	8.8	Flight software for an earth orbiting satellite
<i>pc4</i>	37	1 458	178/1280	7.2	Flight software for an earth orbiting satellite
<i>pc5</i>	38	17 186	516/16670	32.3	A safety enhancement of a cockpit upgrade system

^①Please refer to <http://www.locmetrics.com/alternatives.html> for some freely-accessed tools.

We compare the performance of ROCUS with the following methods.

- ROCA, standing for RandOm Committee using All labeled and unlabeled examples, is a disagreement-based semi-supervised learning method. ROCA is almost the same as ROCUS except it ignores the imbalanced class distribution. In each semi-supervised learning iteration, ROCA directly uses all the automatically labeled examples for classifier refinement.

- LABELED is a random committee learning only from the labeled set. It can be regarded as the initial status of ROCA before exploiting any unlabeled examples.

- UNDERSAMPLING is a class-imbalance learning method. It works in supervised settings. Although it does not exploit the unlabeled data to improve its performance of defect detection, it attempts to reduce the influence of the imbalanced labeled set using under-sampling. The minority-majority ratio after under-sampling (γ) is set to 1.

- SMOTE^[43] is another supervised class-imbalance learning method. Unlike UNDERSAMPLING reducing the number of the majority-class examples, SMOTE balance the labeled dataset by generating many virtual minority-class examples in the neighborhood of a minority-class example.

Note that we use ADABOOST as the base learning algorithm for the 4 compared methods, just like what is used in ROCUS. Additionally, we use an ADABOOST learning only on the labeled set as the baseline

for comparison. Since ADABOOST involves the training of multiple classifiers, to make the learning process of ADABOOST fast, we instantiate the base learner of ADABOOST as decision stumps which make decisions based on the value of only one feature. All the methods used in the experiment are implemented using WEKA^[54].

To compare performances of different methods of defect detection, we use two widely-used evaluation measures, namely F1-measure and AUC^[55]. F1-measure summarizes the precision and recall of the detection. High precision and recall result in high F1-measure. AUC measures the area under the ROC curve, which indicates how well the test examples is ordered according to the real-value output of the classifiers. AUC is large if the test examples of the minority class are placed to the top of the ordering. Note that, unlike F1-measure that directly reveals how well the prediction over the test examples is, AUC only shows the potential of a classifier to produce good classifications over the test examples. A classifier with high AUC can still produce bad classification if the learned threshold over the real-valued output is bad for the classification.

For each dataset at different labeled rates, we repeat the random partition of labeled/unlabeled/test sets for 50 times, the performances of all the compared methods are averaged over the 50 runs. The average F1-measures are tabulated in Tables 3~6 and the average AUCs are tabulated in Tables 7~10. In each table, the best performance among all the compared methods is boldfaced in each experimental dataset. The columns

Table 3. F1-Measures of the Compared Methods When $\mu = 0.1$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.097	0.150	0.369	0.416	0.109	-27.3	0.413	174.8
<i>kc1</i>	0.216	0.268	0.342	0.392	0.185	-31.0	0.419	56.3
<i>kc2</i>	0.484	0.444	0.482	0.492	0.433	-2.6	0.536	20.7
<i>mw1</i>	0.197	0.146	0.225	0.232	0.111	-24.5	0.238	62.8
<i>pc1</i>	0.133	0.167	0.213	0.191	0.081	-51.2	0.249	49.6
<i>pc3</i>	0.155	0.154	0.244	0.280	0.038	-75.3	0.267	74.0
<i>pc4</i>	0.316	0.238	0.434	0.417	0.050	-78.8	0.398	67.2
<i>pc5</i>	0.304	0.306	0.449	0.327	0.158	-48.4	0.487	59.3
Avg.	0.238	0.234	0.345	0.343	0.146	-42.4	0.376	70.6

Table 4. F1-Measures of the Compared Methods When $\mu = 0.2$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.059	0.126	0.387	0.426	0.098	-22.1	0.428	240.0
<i>kc1</i>	0.206	0.241	0.366	0.408	0.189	-21.7	0.425	76.4
<i>kc2</i>	0.471	0.463	0.481	0.502	0.456	-1.4	0.529	14.3
<i>mw1</i>	0.206	0.220	0.226	0.221	0.147	-33.2	0.266	21.0
<i>pc1</i>	0.081	0.156	0.215	0.199	0.053	-66.0	0.300	92.1
<i>pc3</i>	0.127	0.133	0.272	0.305	0.032	-75.7	0.326	144.8
<i>pc4</i>	0.248	0.293	0.505	0.481	0.081	-72.4	0.475	61.9
<i>pc5</i>	0.228	0.317	0.478	0.357	0.230	-27.4	0.452	42.5
Avg.	0.203	0.244	0.366	0.362	0.161	-4.0	0.400	86.6

Table 5. F1-Measures of the Compared Methods When $\mu = 0.3$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.045	0.087	0.400	0.429	0.064	-26.4	0.426	388.9
<i>kc1</i>	0.159	0.265	0.387	0.417	0.219	-17.2	0.441	66.3
<i>kc2</i>	0.468	0.504	0.487	0.535	0.489	-3.0	0.556	10.3
<i>mw1</i>	0.219	0.153	0.219	0.225	0.099	-35.1	0.254	66.0
<i>pc1</i>	0.103	0.159	0.212	0.228	0.042	-73.6	0.347	117.9
<i>pc3</i>	0.053	0.117	0.279	0.315	0.021	-82.1	0.367	213.2
<i>pc4</i>	0.288	0.316	0.536	0.504	0.144	-54.5	0.497	57.5
<i>pc5</i>	0.226	0.317	0.470	0.375	0.266	-16.1	0.434	36.9
Avg.	0.195	0.240	0.374	0.379	0.168	-38.5	0.415	119.6

Table 6. F1-Measures of the Compared Methods When $\mu = 0.4$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.042	0.087	0.396	0.432	0.074	-14.9	0.425	390.3
<i>kc1</i>	0.193	0.265	0.395	0.412	0.214	-19.4	0.443	67.1
<i>kc2</i>	0.473	0.480	0.524	0.520	0.496	3.2	0.562	17.1
<i>mw1</i>	0.195	0.175	0.246	0.202	0.090	-48.6	0.288	64.8
<i>pc1</i>	0.071	0.136	0.224	0.248	0.077	-43.4	0.353	159.5
<i>pc3</i>	0.040	0.098	0.275	0.328	0.024	-75.5	0.383	290.2
<i>pc4</i>	0.272	0.320	0.542	0.511	0.141	-56.0	0.504	57.4
<i>pc5</i>	0.228	0.330	0.480	0.378	0.291	-11.9	0.411	24.5
Avg.	0.189	0.236	0.385	0.379	0.176	-33.3	0.421	133.9

Table 7. AUCs of the Compared Methods When $\mu = 0.1$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.680	0.697	0.694	0.687	0.698	0.1	0.698	0.1
<i>kc1</i>	0.756	0.773	0.757	0.747	0.773	0.1	0.780	0.9
<i>kc2</i>	0.774	0.777	0.759	0.759	0.790	1.7	0.806	3.7
<i>mw1</i>	0.641	0.680	0.635	0.604	0.661	-2.7	0.708	4.1
<i>pc1</i>	0.716	0.741	0.715	0.666	0.727	-1.9	0.740	-0.1
<i>pc3</i>	0.735	0.725	0.732	0.712	0.716	-1.2	0.734	1.3
<i>pc4</i>	0.871	0.806	0.878	0.825	0.807	0.2	0.816	1.3
<i>pc5</i>	0.947	0.955	0.949	0.943	0.956	0.1	0.957	0.2
Avg.	0.765	0.769	0.765	0.743	0.766	-0.5	0.780	1.4

Table 8. AUCs of the Compared Methods When $\mu = 0.2$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.694	0.709	0.706	0.697	0.710	0.2	0.709	0.1
<i>kc1</i>	0.766	0.788	0.771	0.763	0.793	0.6	0.793	0.6
<i>kc2</i>	0.781	0.792	0.760	0.772	0.787	-0.7	0.806	1.8
<i>mw1</i>	0.694	0.705	0.686	0.648	0.686	-2.6	0.704	-0.1
<i>pc1</i>	0.758	0.798	0.770	0.705	0.794	-0.5	0.805	0.8
<i>pc3</i>	0.764	0.765	0.770	0.746	0.766	0.2	0.774	1.2
<i>pc4</i>	0.886	0.849	0.906	0.878	0.847	-0.3	0.851	0.3
<i>pc5</i>	0.950	0.963	0.955	0.953	0.963	0.0	0.962	-0.1
Avg.	0.787	0.796	0.791	0.770	0.793	-0.4	0.801	0.6

of “ROCA imprv.” and “ROCUS imprv.” report the performance improvement of ROCA and ROCUS, respectively, after exploiting the unlabeled data. Let a denote the initial performance before exploiting any unlabeled

data and b denote the final performance after the semi-supervised learning. The improvement is computed as $(b - a)/a$. The last rows of the tables report the average performances over all the experimental datasets.

Table 9. AUCs of the Compared Methods When $\mu = 0.3$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.698	0.707	0.713	0.696	0.706	-0.2	0.705	-0.3
<i>kc1</i>	0.769	0.800	0.780	0.773	0.800	0.0	0.800	-0.1
<i>kc2</i>	0.788	0.798	0.777	0.788	0.807	1.1	0.811	1.6
<i>mw1</i>	0.691	0.693	0.698	0.678	0.682	-1.6	0.704	1.7
<i>pc1</i>	0.774	0.837	0.796	0.747	0.834	-0.3	0.850	1.6
<i>pc3</i>	0.768	0.781	0.782	0.763	0.789	1.0	0.791	1.2
<i>pc4</i>	0.894	0.867	0.915	0.893	0.870	0.3	0.873	0.7
<i>pc5</i>	0.954	0.964	0.960	0.956	0.965	0.1	0.964	0.0
Avg.	0.792	0.806	0.803	0.787	0.807	0.1	0.812	0.8

Table 10. AUCs of the Compared Methods When $\mu = 0.4$

Dataset	ADABOOST	LABELED	SMOTE	UNDERSAMPLE	ROCA	ROCA Imprv. (%)	ROCUS	ROCUS Imprv. (%)
<i>jm1</i>	0.702	0.709	0.716	0.701	0.709	0.0	0.709	0.0
<i>kc1</i>	0.769	0.802	0.782	0.774	0.800	-0.2	0.800	-0.1
<i>kc2</i>	0.800	0.812	0.785	0.780	0.813	0.2	0.818	0.8
<i>mw1</i>	0.699	0.726	0.716	0.676	0.717	-1.2	0.717	-1.2
<i>pc1</i>	0.788	0.839	0.812	0.781	0.845	0.6	0.854	1.7
<i>pc3</i>	0.770	0.794	0.789	0.776	0.795	0.2	0.795	0.1
<i>pc4</i>	0.899	0.879	0.921	0.899	0.876	-0.3	0.882	0.3
<i>pc5</i>	0.954	0.965	0.960	0.956	0.966	0.0	0.965	0.0
Avg.	0.798	0.816	0.810	0.793	0.815	-0.1	0.818	0.2

Tables 3~6 show that ROCUS always outperforms the other compared methods in software defect detection tasks. The average F1-measure of ROCUS is always the highest at different label rates. By comparing ROCUS with its initial status LABELED, we can find that, after carefully exploiting the unlabeled examples, ROCUS is able to dramatically improve its initial performance. Even if only 10% training examples are labeled, the average performance improvement in terms of F1-measure still reaches 70.6%. As μ increases, such performance improvement can be even larger. Pairwise t -tests at 95% significance level indicate that the performance improvement on each dataset is of statistical significance. This fact suggests that the unlabeled examples are beneficial to constructing better predictive model in software defect detection.

We first compare ROCUS with another disagreement-based semi-supervised learning method ROCA. Although in the labeling strategy, the confidence estimation and the stopping criterion are exactly the same in both semi-supervised learning methods, the performance of the learned predictive models are quite different. In contrast to ROCUS which is able to improve the performance using the unlabeled data, the performance of the predictive model learned by ROCA degrades as the semi-supervised learning proceeds. It can be observed from the tables that, the final performance of ROCA is worse than LABELED, which is the initial

status shared by both semi-supervised learning methods, on almost all the datasets at different labeled rates. The only exception is on *kc2* when $\mu = 0.4$, where ROCA appears slightly better than LABELED. Pairwise t -tests at 95% significance level show that the performance degradation is significant. In fact, the performance of ROCA is the worst among all the compared methods, which is even worse than a simple ADABOOST directly applied only on the labeled set. This fact verifies our claim in previous sections that semi-supervised learning may not be effective on imbalance datasets, because the overwhelming majority-class examples in the newly-labeled set drive the focus of the predictive model away from the minority-class. Note that the only difference between these two semi-supervised learning algorithms is that ROCUS explicitly tackles the class-imbalance problem using the under-sampling technique during the exploitation of the unlabeled examples, while ROCA completely ignores the class-imbalance problem. Therefore, it can be concluded that under-sampling is essential to the effectiveness of semi-supervised learning on imbalanced datasets.

Second, we compare ROCUS with the two supervised learning methods SMOTE and UNDERSAMPLE, which is able to tackle the class-imbalance problem in learning. It can be observed from the tables that although SMOTE or UNDERSAMPLE may achieve the best performance on one or two datasets at a certain label rate, the average

F1-measures are always much less than that of ROCUS. For example, when $\mu = 0.1$, the average F1-measure of ROCUS is 0.376 while SMOTE and UNDERSAMPLE are 0.345 and 0.343, respectively. Such superiority can be more obvious when μ becomes larger. Moreover, Tables 3~6 also show that SMOTE, UNDERSAMPLE and ROCUS perform better than ADABOOST and LABELED, both of which fail to consider the imbalanced class distribution in the learning process. Therefore, it can be concluded that explicitly tackling the class-imbalance problem is helpful to software defect detection, and the performance of class-imbalance learning can be further improved if the unlabeled data are exploited in an appropriate way.

Tables 7~10 tabulate the AUC values of all the compared methods at 4 different label rates. The trends in the 4 tables are similar to that of F1-measure shown in Tables 3~6. Specifically, ROCUS always achieves the best performance in terms of average AUC among the compared methods. The exploitation of unlabeled data in ROCA can hardly lead to any improvement of the average AUC values and it even causes performance degeneration when μ is small. In contrast, after tackling the class-imbalance problem explicitly, unlabeled data become beneficial to ROCUS. Moreover, the performance of SMOTE is comparable to that of the methods

ignoring the class-imbalance problem (i.e., ADABOOST and LABELED), and the performance of UNDERSAMPLE is even worse than that of these two methods. In contrast, by exploiting available unlabeled examples, ROCUS performs better.

Although Tables 3~10 suggest that, in software defect detection, ROCUS can effectively exploit unlabeled examples to achieve better performance even if the class distribution is imbalanced, its performance may vibrate under different degrees of imbalance. In order to study the influence of the imbalance degree on ROCUS, we conduct additional experiments, where we alter the class distribution of the dataset. In detail, a dataset is tailored such that the number of the defective examples over the number of the non-defective examples is roughly γ . In the experiments, $1/\gamma \in \{1, 2, \dots, 10\}$. If the original ratio is larger than γ , we randomly discard some defective examples; otherwise, we randomly discard some defect-free examples. We repeat the experiment in each tailored dataset for 50 times, where the labeled/unlabeled/test datasets are generated according to the 4 label rates (i.e., 10%, 20%, 30%, 40%). Since F1-measure directly reports how well the defect detection is, we only illustrate the performances of the compared methods in terms of the F1-measure.

We plot the average F1-measure versus the inverse

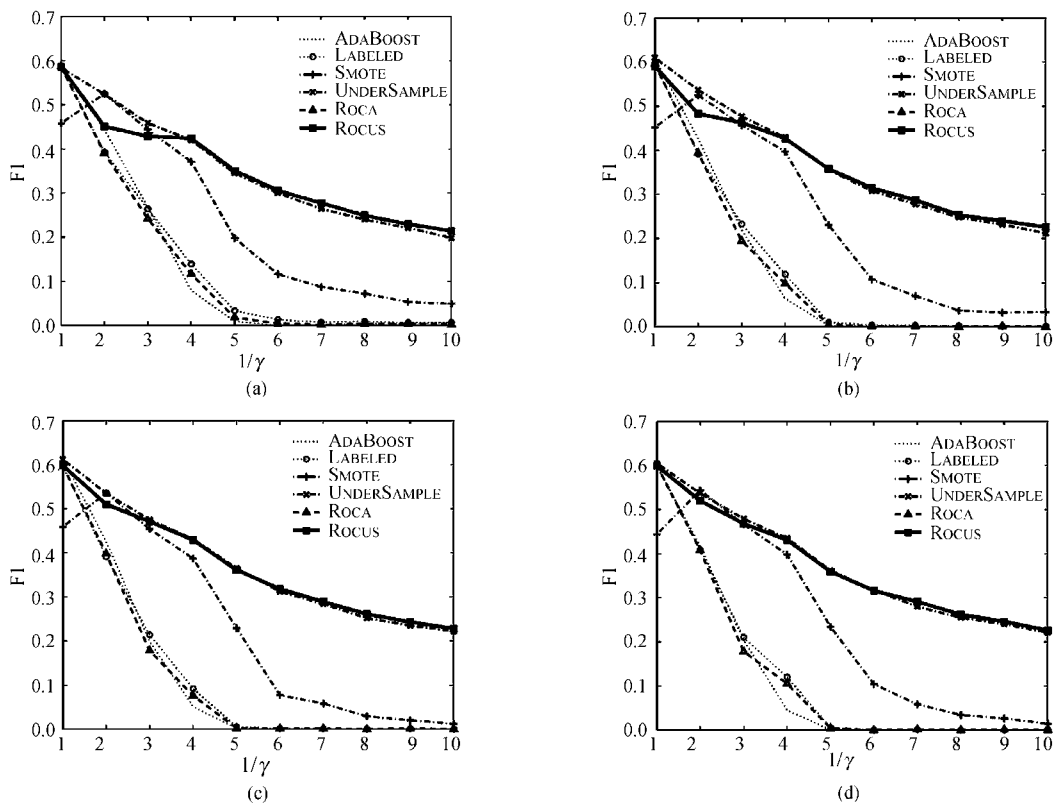


Fig.1. F1-measures of the compared methods on *jm1* at different minority-majority rates. (a) $\mu = 10\%$. (b) $\mu = 20\%$. (c) $\mu = 30\%$. (d) $\mu = 40\%$.

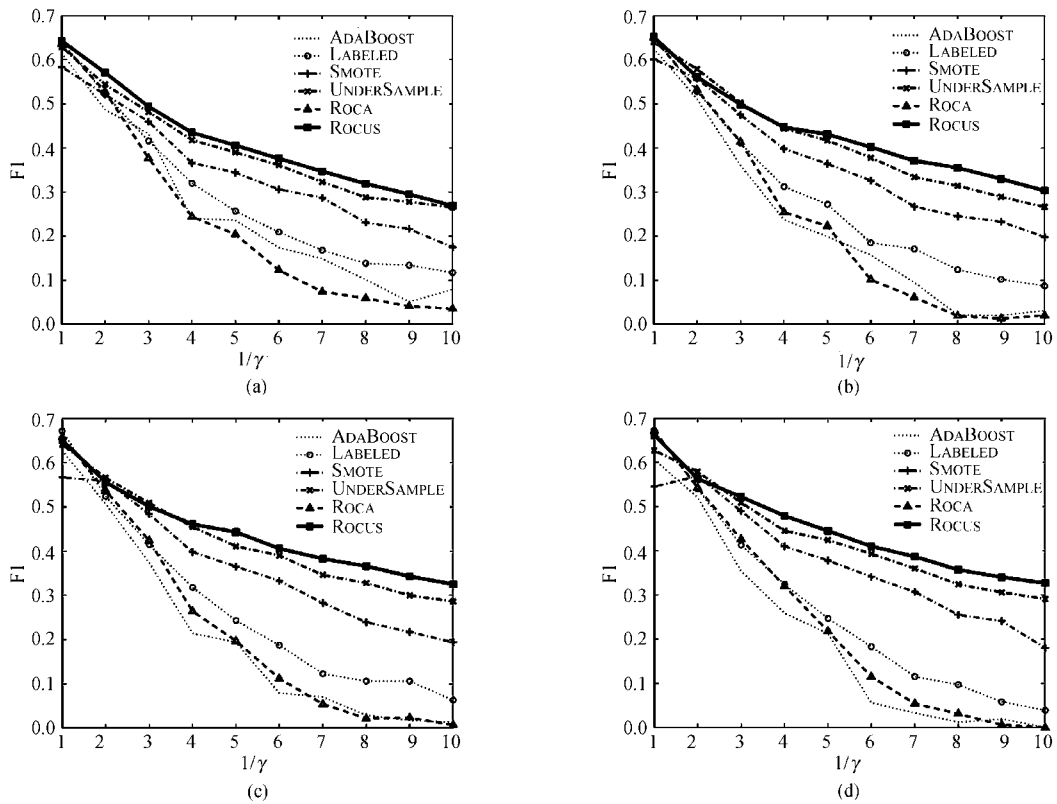


Fig.2. F1-measures of the compared methods on *kc1* at different minority-majority rates. (a) $\mu = 10\%$. (b) $\mu = 20\%$. (c) $\mu = 30\%$. (d) $\mu = 40\%$.

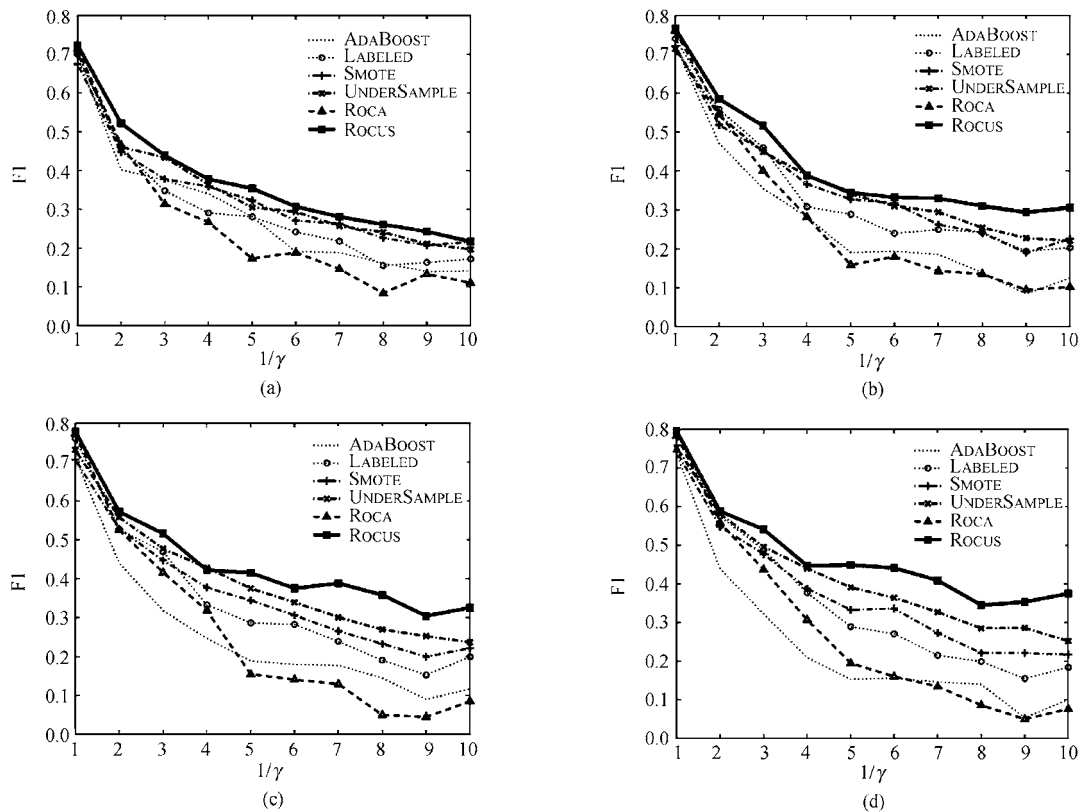


Fig.3. F1-measures of the compared methods on *pc1* at different minority-majority rates. (a) $\mu = 10\%$. (b) $\mu = 20\%$. (c) $\mu = 30\%$. (d) $\mu = 40\%$.

of the minority-majority ratio ($1/\gamma$) at different label rates on the experimental datasets. Due to the space limitation, we only report the results on datasets of small size (*pc1*), median size (*kc1*), and large size (*jm1*), as shown in Figs. 1~3. Similar trends are observed on the other datasets.

As expected, F1-measures of all the compared methods decrease as the dataset becomes more imbalanced, but the influence of the increase of class imbalance on ROCUS is the smallest among the compared methods. It can be observed from the figures that ROCUS almost always performs better than the other methods, and such superiority is more obvious when the class distribution becomes more imbalanced. We first consider the two semi-supervised learning methods. When the class distribution is balanced, ROCA and ROCUS appear to be comparable. As $1/\gamma$ increases, the performance of ROCA degrades rather fast, especially when label rate is small. For example, when only 10% of the training data from *kc1* are labeled, the curve of ROCA even drops below that of ADABOOST after $1/\gamma$ grows larger than 3. This fact suggests that the degree of imbalance has great influence on the semi-supervised learning method if it does not tackle the class-imbalance problem explicitly. Then, we compare ROCUS, SMOTE and UNDERSAMPLE, each of which explicitly considers the imbalanced class distribution in learning. The figures show that there exists a gap between the curves of ROCUS and the other two methods, respectively. Such a gap increases as $1/\gamma$ becomes larger, which indicates that exploiting unlabeled data can help to improve the performance of class-imbalance learning.

5 Conclusion

Detection of defects in software modules is important in improving the quality of a software system. However, many real-world data for software defect detection are imbalanced and only a small portion of examples are labeled as “defective” or “defect-free” in advance. In this paper, we address these two practical yet important issues simultaneously for software defect detection. We propose a disagreement-based semi-supervised learning method ROCUS to exploit the abundant unlabeled examples for better detection. ROCUS employs under-sampling to tailor the newly-labeled set, which effectively reduces the chance that the refined predictive model being less sensitive to the defective examples during the iterative semi-supervised learning process. Experimental results of real-world software

defect detection tasks show that ROCUS can achieve better detection performance than a semi-supervised learning method that ignores the class-imbalance nature of the tasks and a class-imbalance learning method that is not able to exploit unlabeled data.

Note that ROCUS maintains the diversity of the ensemble between individual classifiers by injecting randomness into the base learner. The diversity can also be achieved by introducing *selective ensemble*^[56]. Exploiting unlabeled data for improving the performance of individual classifiers and the diversity between them simultaneously in disagreement-based semi-supervised learning is an interesting future work.

In software defect detection, it would be more useful to predict the number of defects that may be contained in certain software modules. Such a problem can be formalized as an regression problem, just as done in [57]. Recently, Zhou and Li^[39] proposed a disagreement-based semi-supervised regression method, whose idea may be useful to extend ROCUS to predicting the number of defects in each module in future. Moreover, since misclassifying a defective module as defect-free may lead to worse consequence than misclassifying a defect-free module as defective, extending ROCUS to cost-sensitive scenario, where the overall cost rather than misclassification error is minimized during learning, would be another interesting future work. Additionally, in software defect detection settings, each module is represented as a single feature vector based on the metrics extracted from the module. It would be interesting to find appropriate metrics to represent each module with a set of feature vectors such that the defect detection problem may be solved within multi-instance learning^[58] framework, which will be investigated in the future.

Since the problems of “lack of sufficient labeled data” and “data imbalance” may be tangled together, exploiting the interaction between these two important issues in influencing the problems similar to software defect detection is expected to achieve better results, which is another interesting work to be done in future. Besides, it would be another interesting work to adapt ROCUS to other software engineering tasks or even to the tasks beyond software engineering, where the data distributions are essentially imbalanced and the labels for examples are difficult to obtain.

Acknowledgement The authors would like to thank the anonymous reviewers for their valuable

comments and suggestions to improve this paper, and thank Dr. Hongyu Zhang for his comments and Wei-Wei Tu for the proof-reading.

References

- [1] Dai Y S, Xie M, Long Q, Ng S H. Uncertainty analysis in software reliability modeling by Bayesian approach with maximum-entropy principle. *IEEE Transactions on Software Engineering*, 2007, 33(11): 781-795.
- [2] Guo L, Ma Y, Cukic B, Singh H. Robust prediction of fault-proneness by random forests. In *Proc. the 15th International Symposium on Software Reliability Engineering*, Nov. 2-5, 2004, pp.417-428.
- [3] Khoshgoftaar T M, Allen E B, Jones W D, Hudepohl J P. Classification-tree models of software-quality over multiple releases. *IEEE Transactions on Reliability*, 2000, 49(1): 4-11.
- [4] Lessmann S, Baesens B, Mues C, Pietsch S. Benchmarking classification models for software defect prediction: A proposed framework and novel findings. *IEEE Transactions on Software Engineering*, 2008, 34(4): 485-496.
- [5] Menzies T, Greenwald J, Frank A. Data mining static code attributes to learn defect predictors. *IEEE Transactions on Software Engineering*, 2007, 33(1): 2-13.
- [6] Zhang H, Zhang X. Comments on "Data mining static code attributes to learn defect predictors". *IEEE Transactions on Software Engineering*, 2007, 33(9): 635-637.
- [7] Zhou Y, Leung H. Empirical analysis of object-oriented design metrics for predicting high and low severity faults. *IEEE Transactions on Software Engineering*, 2006, 32(10): 771-789.
- [8] Seliya N, Khoshgoftaar T M. Software quality estimation with limited fault data: A semi-supervised learning perspective. *Software Quality Journal*, 2007, 15: 327-344.
- [9] Pelayo L, Dick S. Applying novel resampling strategies to software defect prediction. In *Proc. the 2007 Annual Meeting of the North American Fuzzy Information Processing Society*, San Diego, USA, Jun. 24-27, 2007, pp.69-72.
- [10] Zhou Z H, Li M. Semi-supervised learning by disagreement. *Knowledge and Information Systems*, 2010, 24(3): 415-439.
- [11] Drummond C, Holte R C. C4.5, class imbalance, and cost sensitivity: Why under-sampling beats over-sampling. In *Working Notes of the ICML'03 Workshop on Learning from Imbalanced Data Sets*, Washington DC, USA, Jul. 21, 2003.
- [12] Zheng A X, Jordan M I, Liblit B, Naik M, Aiken A. Statistical debugging: Simultaneous identification of multiple bugs. In *Proc. the 23rd International Conference on Machine Learning*, Pittsburgh, USA, Jun. 25-29, 2006, pp.1105-1112.
- [13] Andrzejewski D, Mulhern A, Liblit B, Zhu X. Statistical debugging using latent topic models. In *Proc. the 18th European Conference on Machine Learning*, Warsaw, Poland, Sept. 17-21, 2007, pp.6-17.
- [14] Chilimbi T M, Liblit B, Mehra K K *et al.* HOLMES: Effective statistical debugging via efficient path profiling. In *Proc. the 31st International Conference on Software Engineering*, Vancouver, Canada, May 16-24, 2009, pp.34-44.
- [15] Basili V R, Briand L C, Melo W L. A validation of object-oriented design metrics as quality indicators. *IEEE Transactions on Software Engineering*, 1996, 22(10): 751-761.
- [16] Khoshgoftaar T M, Allen E B. Neural Networks for Software Quality Prediction. *Computational Intelligence in Software Engineering*, Pedrycz W, Peters J F (eds.), World Scientific, Singapore, 1998, pp.33-63.
- [17] Halstead M H. Elements of Software Science, Elsevier, 1977.
- [18] McCabe T J. A complexity measure. *IEEE Transactions on Software Engineering*, 1976, 2(4): 308-320.
- [19] Gyimóthy T, Ferenc R, Siket I. Empirical validation of object-oriented metrics on open source software for fault prediction. *IEEE Transactions on Software Engineering*, 2005, 31(10): 897-910.
- [20] Ganesan K, Khoshgoftaar T M, Allen E B. Verifying requirements through mathematical modelling and animation. *International Journal of Software Engineering and Knowledge Engineering*, 2000, 10(2): 139-152.
- [21] Khoshgoftaar T M, Seliya N. Fault prediction modeling for software quality estimation: Comparing commonly used techniques. *Empirical Software Engineering*, 2003, 8(3): 255-283.
- [22] Fenton N E, Neil M. A critique of software defect prediction models. *IEEE Transactions Software Engineering*, 1999, 25(5): 675-689.
- [23] Pérez-Miñana E, Gras J J. Improving fault prediction using Bayesian networks for the development of embedded software applications. *Software Testing, Verification Reliability*, 2006, 16(3): 157-174.
- [24] Chapelle O, Schölkopf B, Zien A. *Semi-Supervised Learning*. Cambridge: MIT Press, MA, 2006.
- [25] Zhu X. Semi-supervised learning literature survey. Technical Report 1530, Department of Computer Sciences, University of Wisconsin at Madison, Madison, WI, 2006, http://www.cs.wisc.edu/~jerryzhu/pub/ssl_survey.pdf.
- [26] Miller D J, Uyar H S. A Mixture of Experts Classifier with Learning Based on Both Labelled and Unlabelled Data. *Advances in Neural Information Processing Systems 9*, Mozer M, Jordan M I, Petsche T (eds.), MIT Press, Cambridge, MA, 1997, pp.571-577.
- [27] Nigam K, McCallum A K, Thrun S, Mitchell T. Text classification from labeled and unlabeled documents using EM. *Machine Learning*, 2000, 39(2/3): 103-134.
- [28] Shahshahani B, Landgrebe D. The effect of unlabeled samples in reducing the small sample size problem and mitigating the Hughes phenomenon. *IEEE Transactions on Geoscience and Remote Sensing*, 1994, 32(5): 1087-1095.
- [29] Chapelle O, Zien A. Semi-supervised learning by low density separation. In *Proc. the 10th International Workshop on Artificial Intelligence and Statistics*, Barbados, Jan. 6-8, 2005, pp.57-64.
- [30] Grandvalet Y, Bengio Y. Semi-Supervised Learning by Entropy Minimization. *Advances in Neural Information Processing Systems*, Saul L K, Weiss Y, Bottou L (eds.), MIT Press, Cambridge, MA, 2005, pp.529-536.
- [31] Joachims T. Transductive inference for text classification using support vector machines. In *Proc. the 16th International Conference on Machine Learning*, Bled, Slovenia, Jun. 27-30, 1999, pp.200-209.
- [32] Belkin M, Niyogi P, Sindhvani V. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 2006, 7(11): 2399-2434.
- [33] Zhou D, Bousquet O, Lal T N, Weston J, Schölkopf B. Learning with Local and Global Consistency. *Advances in Neural Information Processing Systems 16*, Thrun S, Saul L, Schölkopf B (eds.), MIT Press, Cambridge, MA, 2004.
- [34] Zhu X, Ghahramani Z, Lafferty J. Semi-supervised learning using Gaussian fields and harmonic functions. In *Proc. the 20th International Conference on Machine Learning*, Washington, DC, USA, Aug. 21-24, 2003, pp.912-919.
- [35] Blum A, Mitchell T. Combining labeled and unlabeled data with co-training. In *Proc. the 11th Annual Conference on Computational Learning Theory*, Madison, USA, Jul. 24-26, 1998, pp.92-100.
- [36] Goldman S, Zhou Y. Enhancing supervised learning with unlabeled data. In *Proc. the 17th International Conference on*

- Machine Learning*, San Francisco, USA, Jun. 29-Jul. 2, 2000, pp.327-334.
- [37] Li M, Zhou Z H. Improve computer-aided diagnosis with machine learning techniques using undiagnosed samples. *IEEE Transactions on Systems, Man and Cybernetics — Part A: Systems and Humans*, 2007, 37(6): 1088-1098.
- [38] Zhou Z H, Li M. Tri-training: Exploiting unlabeled data using three classifiers. *IEEE Transactions on Knowledge and Data Engineering*, 2005, 17(11): 1529-1541.
- [39] Zhou Z H, Li M. Semi-supervised regression with co-training style algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 2007, 19(11): 1479-1493.
- [40] Steedman M, Osborne M, Sarkar A et al. Bootstrapping statistical parsers from small data sets. In *Proc. the 11th Conference on the European Chapter of the Association for Computational Linguistics*, Budapest, Hungary, Apr. 12-17, 2003, pp.331-338.
- [41] Li M, Zhou Z H. Semi-supervised document retrieval. *Information Processing & Management*, 2009, 45(3): 341-355.
- [42] Zhou Z H, Chen K J, Dai H B. Enhancing relevance feedback in image retrieval using unlabeled data. *ACM Transactions on Information Systems*, 2006, 24(2): 219-244.
- [43] Chawla N V, Bowyer K W, Hall L O, Kegelmeyer W P. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 2002, 16: 321-357.
- [44] Kubat M, Matwin S. Addressing the curse of imbalanced training sets: One-sided selection. In *Proc. the 14th Int. Conf. Machine Learning*, Nashville, USA, 1997, pp.179-186.
- [45] Domingos P. MetaCost: A general method for making classifiers cost-sensitive. In *Proc. the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Diego, USA, Aug. 15-18, 1999, pp.155-164.
- [46] Elkan C. The foundations of cost-sensitive learning. In *Proc. the 17th International Joint Conference on Artificial Intelligence*, Seattle, USA, Aug. 4-10, 2001, pp.973-978.
- [47] Batista G, Prati R C, Monard M C. A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations*, 2004, 6(1): 20-29.
- [48] Liu X Y, Wu J X, Zhou Z H. Exploratory under-sampling for class-imbalance learning. *IEEE Transactions on Systems, Man and Cybernetics — Part B: Cybernetics*, 2009, 39(2): 539-550.
- [49] Angluin D, Laird P. Learning from noisy examples. *Machine Learning*, 1988, 2(4): 343-370.
- [50] Ho T K. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1998, 20(8): 832-844.
- [51] Breiman L. Bagging predictors. *Machine Learning*, 1996, 24(2): 123-140.
- [52] Chapman M, Callis P, Jackson W. Metrics data program. NASA IV and V Facility, 2004, <http://mdp.ivv.nasa.gov/>.
- [53] Schapire R E. A brief introduction to Boosting. In *Proc. the 16th International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, Jul. 31-Aug. 6, 1999, pp.1401-1406.
- [54] Witten I H, Frank E. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2005.
- [55] Bradley A P. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 1997, 30(6): 1145-1159.
- [56] Zhou Z H, Wu J, Tang W. Ensembling neural networks: Many could be better than all. *Artificial Intelligence*, 2002, 137(1/2): 239-263.
- [57] Khoshgoftaar T M, Seliya N. Tree-based software quality estimation models for fault prediction. In *Proc. the 8th IEEE International Symp. Software Metrics*, Ottawa, Canada, Jun. 4-7, 2002, pp.203-214.
- [58] Dietterich T G, Lathrop R H, Lozano-Pérez T. Solving the Multiple instance problem with axis-parallel rectangles. *Artificial Intelligence*, 1997, 89(1/2): 31-71.



Yuan Jiang received the Ph.D. degree in computer science from Nanjing University, China, in 2004. Now she is an associate professor at the Department of Computer Science and Technology, Nanjing University. Her research interests include machine learning, information retrieval and data mining. In these areas, she has published more than 30 technical papers in refereed journals or conferences. She served as the publication chair of PAKDD'07, also served as program committee members for many conferences such as CCTA'07, BIBE'07, ICNC'08, ICNC'09, CCDM'09, NCIIP'09, CCML'10. She is now a committee member of Machine Learning Society of Chinese Association of Artificial Intelligence (CAAI), and is a committee member of Artificial Intelligence Society of Jiangsu Computer Association. She is a member of CCF. She was selected in the Program for New Century Excellent talents in University, Ministry of Education, in 2009.



Ming Li received the B.Sc. and Ph.D. degrees in computer science from Nanjing University, China, in 2003 and 2008 respectively. He is currently an assistant professor with LAMDA Group, the Department of Computer Sciences and Technology, Nanjing University. His major research interests include machine learning, data mining and information retrieval, especially on learning with labeled and unlabeled data. He has been granted various awards including the CCF Outstanding Doctoral Dissertation Award (2009), Microsoft Fellowship Award (2005). He served on the program committee of a number of important international conferences including KDD'10, ACML'10, ACML'09, ACM CKIM'09, IEEE ICME'10, AI'10, and served as reviewers for a number of refereed journals including IEEE Trans. KDE, IEEE Trans. NN, IEEE Trans. SMCC, ACM Trans. IST, Pattern Recognition, Knowledge and Information Systems, Journal of Computer Science and Technology. He is a committee member of the Machine Learning Society of the CAAI, member of ACM, IEEE, IEEE Computer Society, CCF and CAAI.



Zhi-Hua Zhou received the B.Sc., M.Sc. and Ph.D. degrees in computer science from Nanjing University, China, in 1996, 1998 and 2000, respectively, all with the highest honors. He joined the Department of Computer Science and Technology at Nanjing University as an assistant professor in 2001, and is currently Cheung Kong Professor

and Director of the LAMDA Group. His research interests are in artificial intelligence, machine learning, data mining, pattern recognition, information retrieval, evolutionary computation and neural computation. In these areas he has published over 70 papers in leading international journals or conference proceedings. Dr. Zhou has won various awards/honors including the National Science and Technology Award for Young Scholars of China (2006), the Award of National Science Fund for Distinguished Young Scholars of China (2003), the National Excellent Doctoral Dissertation Award of China (2003), the Microsoft Young Professorship Award (2006). He is an associate editor of *IEEE Transactions on Knowledge and Data Engineering*, associate editor-in-chief of *Chinese Science Bulletin*, and on the editorial boards of *Artificial Intelligence in Medicine*, *Intelligent Data Analysis*, *Science in China*. He is the founder of ACML, Steering Committee member of PAKDD and PRICAI, Program Committee Chair/Co-Chair of PAKDD'07, PRICAI'08 and ACML'09, vice Chair or area Chair of conferences including IEEE ICDM'06, IEEE ICDM'08, SIAM DM'09, ACM CIKM'09, and general chair/co-chair or program committee chair/co-chair of a dozen of native conferences. He is the chair of the Machine Learning Society of the CAAI, vice chair of the Artificial Intelligence and Pattern Recognition Society of the CCF, and chair of the IEEE Computer Society Nanjing Chapter. He is a fellow of IET, a member of AAAI and ACM, and a senior member of IEEE, IEEE Computer Society and IEEE Computational Intelligence Society, CCF.