

Poster: Identification of Methods with Low Fault Risk

Rainer Niedermayr
University of Stuttgart, CQSE GmbH
Garching b. München, Germany
niedermayr@cqse.eu

Tobias Röhm
CQSE GmbH
Garching b. München, Germany
roehm@cqse.eu

Stefan Wagner
University of Stuttgart
Stuttgart, Germany
stefan.wagner@informatik.uni-stuttgart.de

ABSTRACT

Test resources are usually limited and therefore it is often not possible to completely test an application before a release. Therefore, testers need to focus their activities on the relevant code regions. In this paper, we introduce an inverse defect prediction approach to identify methods that contain hardly any faults. We applied our approach to six Java open-source projects and show that on average 31.6% of the methods of a project have a low fault risk; they contain in total, on average, only 5.8% of all faults. Furthermore, the results suggest that, unlike defect prediction, our approach can also be applied in cross-project prediction scenarios. Therefore, inverse defect prediction can help prioritize untested code areas and guide testers to increase the fault detection probability.

ACM Reference Format:

Rainer Niedermayr, Tobias Röhm, and Stefan Wagner. 2018. Poster: Identification of Methods with Low Fault Risk. In *ICSE '18 Companion: 40th International Conference on Software Engineering Companion, May 27-June 3, 2018, Gothenburg, Sweden*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3183440.3195022>

1 INTRODUCTION

Software testing can be very time consuming and test resources are usually scarce. Therefore, it is often not possible to completely test the whole code base before each release. Consequently, development teams must limit their testing scope and focus on code regions that have the best cost-benefit ratio regarding test resources [2].

To support development teams in this activity, defect prediction has been developed and studied extensively. Defect prediction identifies code regions that are likely to contain a fault and should therefore be tested [5, 7]. However, after several decades of research on defect prediction, it is still hardly used in practice. Defect prediction models need to be trained with precise historical fault data from the project, which is often not available—especially in new projects. Cross-project predictions, which use models trained from data of other projects, are still considered as a difficult task [6, 8].

This paper suggests an alternative approach for prioritizing code regions: inverse defect prediction (IDP). The idea behind IDP is to identify code regions with *low* fault risk, which can be deferred when writing automated tests if none yet exist. The main difference to traditional defect prediction lies in the predicted classes and in the optimization target. While defect prediction classifies an artifact either as buggy or non-buggy, IDP identifies artifacts that

Table 1: Computed metrics for each method (excerpt).

	Metric Name	Type
M1	Source Lines of Code (SLOC)	length
M2	Cyclomatic Complexity (CC)	complexity
M3	Max. Nesting Depth	max. value
M9	Array Accesses	count
M11	Value Assignments	count
M16	If Conditions	count
M20	Loops	count
M21	Method Invocations	count
M23	Null Literals	count
M24	Return Statements	count
M29	Ternary Operations	count
M30	Throw Statements	count
M34	Is Constructor	category
M36	Is Getter	category
...		

exhibit a low fault risk and does not make an assumption about the remaining artifacts. Instead of predicting all non-faulty artifacts, IDP aims to identify only those that contain *no faults with high certainty*. Therefore, IDP strives to achieve a high precision for the identified artifacts, while recall is less important. In contrast, defect prediction aims at a high recall to detect as many faults as possible and at a high precision such that only few false positives occur.

We applied IDP on the Defects4J dataset [3] at the method level. We evaluated how many faults the identified low-fault-risk methods contain and how much savings potential can be gained by ignoring them during testing. Our results show that IDP can successfully identify low-fault-risk methods, which contain considerably less faults than an arbitrary method, and are frequent enough to provide a worthwhile savings potential for QA activities. Moreover, our results indicate that IDP can be an alternative to traditional defect prediction in cross-project prediction scenarios.

2 APPROACH

The IDP approach to identify low-fault-risk methods comprises the computation of source-code metrics for each method, the data pre-processing before the mining, and the creation of a classifier using association rule mining.

Like defect prediction models, IDP needs metrics to train a classifier and an indication whether a method was faulty at least once. We computed for each method common software-analysis metrics, metrics that count occurrences of Java language constructs, and categories to which a method can belong to. Table 1 presents an excerpt of the 39 used metrics.

As association rule mining computes frequent itemsets from categorical attributes, our next step was to discretize the numerical metrics. For the numeric software-analysis metrics, such as SLOC and the cyclomatic complexity, we inspected the value distribution, computed tertiles, and created a new binary metric for each tertile. For all count metrics, we created a binary “has-no”-metric, which is true if the value is zero, e.g., $CountLoops = 0 \Rightarrow NoLoops = true$. For the method categories (setter, getter, ...), no transformation was necessary since they are already binary.

Defect datasets are often highly imbalanced and contain only a small portion of faulty methods [4]. Therefore, we applied SMOTE¹, a well-known algorithm for over- and under-sampling, to the training data to address the imbalance.

To identify low-fault-risk methods, we first mined association rules of the type $\{Metric1, Metric2, Metric3, \dots\} \rightarrow \{NotFaulty\}$ by applying the *Apriori* algorithm [1] to the training data. Next, we ordered the obtained rules descending by their confidence value. The *confidence* of a rule expresses the proportion of methods that are non-faulty and satisfy the metric predicates out of all methods that satisfy the metric predicates. Thus, it can be considered as the precision of a rule. Finally, to build the low-fault-risk classifier, we combined the top n association rules with the highest confidence values using the logical-or operator. Hence, we considered a method to have a low fault risk if at least one of the top n rules matched. To determine n , we computed the maximum number of rules until the share of faulty methods in low-fault-risk methods of the training set exceeded a certain threshold.

3 EVALUATION

We evaluated our approach on the Defects4J dataset [3], which contains real faults for six open-source projects. We performed both within- and cross-project predictions and evaluated the identified methods. For within-project predictions, we applied 10-fold cross-validation. We randomly sampled the dataset of each project into ten stratified partitions of equal size and used each sample once for testing the classifier, which is trained on the remaining nine samples. Table 2 exemplarily presents resulting association rules. For cross-project prediction, we evaluated each of the six projects with a classifier trained on data from the respective other five projects.

The results of the within-project predictions show that IDP classified between 16% and 75.3% of the methods as “low fault risk” (median: 31.6%); the identified methods comprise between 5% and 68.5% of the SLOC (median: 14.2%). These methods are on average six times less likely to contain a fault than an arbitrary method. Low-fault-risk methods are usually short, but their source code lines of code (SLOC) are still 3.4 times less likely to contain a fault. Figure 1 illustrates the within-project results.

Cross-project IDP found between 18.4% and 31.2% of the methods to have a low fault risk (median: 23.2%); they contain between 6.3% and 10.1% of the SLOC (median: 7.8%). Surprisingly, their fault density is lower (compared to within-project predictions), because they are nearly 11 times less likely to contain a fault than an arbitrary method. Based on SLOC, they are 3.7 times less likely to contain a fault.

¹Synthetic Minority Over-sampling Technique

Table 2: Top three association rules for the project *Apache Commons Lang* (within-project prediction, fold 1).

#	Rule
1	$\{UniqueVariableIdentifiersLessThan2, NoMethodInvocations\} \Rightarrow \{NotFaulty\}$
2	$\{SlocLessThan4, NoAnyArithmeticOp, NoMethodInvocations\} \Rightarrow \{NotFaulty\}$
3	$\{SlocLessThan4, MaxMethodChainingLessThan2, UniqueVariableIdentifiersLessThan2\} \Rightarrow \{NotFaulty\}$




Figure 1: ● Proportion of low-fault-risk methods classified by within-project IDP and ● their share in faults of all faults. In the project *Google Closure Compiler*, IDP classified 29.9% of the methods to have a low fault risk; they contain only 4.7% of all faults.

4 CONCLUSION

Inverse defect prediction using association rule mining can be used to identify low-fault-risk methods. The identified methods are indeed considerably less likely to contain a fault and can provide a worthwhile savings potential for QA activities. This applies to both within- and cross-project predictions.

ACKNOWLEDGMENTS

This work was partially funded by the German Federal Ministry of Education and Research (BMBF), grant “Q-Effekt, 01IS15003A”. The responsibility for this article lies with the authors.

REFERENCES

- [1] R Agrawal, R Srikant, and others. 1994. Fast Algorithms for Mining Association Rules. In *VLDB'94*.
- [2] T. Hall, S. Beecham, D. Bowes, D. Gray, and S. Counsell. 2012. A Systematic Literature Review on Fault Prediction Performance in Software Engineering. *TSE* 38, 6 (2012).
- [3] R. Just, D. Jalali, and M. Ernst. 2014. Defects4J: A Database of Existing Faults to Enable Controlled Testing Studies for Java Programs. In *ISSTA'14*. Tool demo.
- [4] T. Khoshgoftaar, K. Gao, and N. Seliya. 2010. Attribute Selection and Imbalanced Data: Problems in Software Defect Prediction. In *ICTAI'10*.
- [5] T. Menzies, J. Greenwald, and A. Frank. 2007. Data Mining Static Code Attributes to Learn Defect Predictors. *TSE* 33, 1 (2007).
- [6] B. Turhan, T. Menzies, A. Bener, and J. Di Stefano. 2009. On the Relative Value of Cross-Company and Within-Company Data for Defect Prediction. *Empir. Softw. Eng.* 14, 5 (2009).
- [7] E. Weyuker and T. Ostrand. 2008. What Can Fault Prediction Do for YOU?. In *TAP'08*.
- [8] T. Zimmermann, N. Nagappan, H. Gall, E. Giger, and B. Murphy. 2009. Cross-Project Defect prediction: A Large Scale Experiment on Data vs. Domain vs. Process. In *ESEC/FSE'09*. ACM.