

Predicting Defects for Eclipse

Thomas Zimmermann
Saarland University
tz@acm.org

Rahul Premraj
Saarland University
premrj@cs.uni-sb.de

Andreas Zeller
Saarland University
zeller@acm.org

Abstract

We have mapped defects from the bug database of Eclipse (one of the largest open-source projects) to source code locations. The resulting data set lists the number of pre- and post-release defects for every package and file in the Eclipse releases 2.0, 2.1, and 3.0. We additionally annotated the data with common complexity metrics. All data is publicly available and can serve as a benchmark for defect prediction models.

1. Introduction

Why is it that some programs are more failure-prone than others? This is one of the central questions of software engineering. To answer it, we must first know *which* programs are more failure-prone than others. With this knowledge, we can search for properties of the program or its development process that commonly correlate with defect density; in other words, once we can measure the effect, we can search for its causes.

One of the most abundant, widespread, and reliable sources for failure information is a *bug database*, listing all the problems that occurred during the software lifetime. Unfortunately, bug databases frequently do not directly record how, where, and by whom the problem in question was fixed. This information is hidden in the *version database*, recording all changes to the software source code.

In recent years, a number of techniques have been developed to relate bug reports to fixes [3, 6, 17]. Since we thus can relate bugs to fixes, and fixes to the locations they apply to, we can easily determine the *number of defects* of a component—simply by counting the applied fixes.

We have conducted such a work on the code base of the Eclipse programming environment. In particular, we have computed the mapping of packages and classes to the number of defects that were reported in the first six months *before* and *after* release. In previous work, we made our *Eclipse bug data set* freely

Project:	Eclipse (eclipse.org)
Content:	Defect counts (pre- and post-release) Complexity metrics
Releases:	2.0, 2.1, and 3.0
Level:	Packages and files
URL:	http://www.st.cs.uni-sb.de/softevo/bug-data/eclipse
More data:	Eclipse source code (for archived releases): http://archive.eclipse.org/eclipse/downloads/

Figure 1. Summary of our data set.

available [15]. For this paper, we extended our data with common complexity metrics and the counts of syntactic elements (obtained from abstract syntax trees). With this new data, many predictor models can be built out of the box which we demonstrate in this paper.

We invite readers to use our data for research purposes and to build their own models. We hope that the public availability of data sets like ours will foster empirical research in software engineering, just like the public availability of open source programs fostered research in program analysis.

2. State of the art

Predicting which components are more failure-prone than others has been addressed by a number of researchers in the past. This work, discussed below, used either complexity metrics or historical data to predict failures.

2.1. Complexity metrics

Typically, research on defect-proneness defines metrics to capture the complexity of software and builds models that relate these metrics to defect-proneness [4]. Basili et al. [1] were among the first to validate that OO metrics are useful for predicting defect density. Subramanyam and Krishnan [18] presented a survey on eight more empirical studies, all showing that OO metrics are significantly associated with defects. Post-release defects are the defects that actually matter for the end-users of a program. Only few studies ad-

Table 1. Metrics in the Eclipse data set.

		Metric	File level	Package level
methods	FOUT	Number of method calls (fan out)	avg, max, total	avg, max, total
	MLOC	Method lines of code	avg, max, total	avg, max, total
	NBD	Nested block depth	avg, max, total	avg, max, total
	PAR	Number of parameters	avg, max, total	avg, max, total
	VG	McCabe cyclomatic complexity	avg, max, total	avg, max, total
classes	NOF	Number of fields	avg, max, total	avg, max, total
	NOM	Number of methods	avg, max, total	avg, max, total
	NSF	Number of static fields	avg, max, total	avg, max, total
	NSM	Number of static methods	avg, max, total	avg, max, total
files	ACD	Number of anonymous type declarations	value	avg, max, total
	NOI	Number of interfaces	value	avg, max, total
	NOT	Number of classes	value	avg, max, total
	TLOC	Total lines of code	value	avg, max, total
packages	NOCU	Number of files (compilation units)	N/A	value

dressed post-release defects so far: Binkley and Schach [2] developed a coupling dependency metric and showed that it outperforms several other metrics; Ohlsson and Alberg [13] investigated a number of metrics to predict modules that fail during test or operation. Schröter et al. [16] showed that design data such as import relationships also can predict post-release failures.

The MetriZone project at Microsoft Research investigates how to make early estimates of software quality to predict post-release failures. Nagappan and Ball [11] showed that relative code churn predicts software defect density; (absolute) code churn is the number of lines added or deleted between versions. Additionally, Nagappan et al. [12] carried out the largest study on commercial software so far: Within five Microsoft projects, they identified metrics that predict post-release failures and reported how to systematically build predictors for post-release failures from history. Nagappan and Ball [10] also showed that the ratio between the number of dependencies within a component and the number of dependencies across a component can predict post-release failures.

2.2. Historical data

Several researchers used historical data without taking bug databases into account. Khoshgoftaar et al. [9] classified modules as defect-prone whenever the number of lines of code added or deleted exceeded a threshold. Graves et al. [7] used the sum of contributions to a module in its history to predict defect density. Ostrand et al. [14] used historical data from up to 17 releases to predict the files with the highest defect density in the next release. Hudepohl et al. [8] predicted whether a module would be defect-prone by combining

metrics and historical data. From several software metrics, Denaro et al. [5] learned logistic regression models for Apache 1.3 and verified them against Apache 2.0.

3. Eclipse data set

This section presents details on how we computed the *Eclipse bug data set* and a description of its contents. Additionally, we point out research problems that can be investigated with our data set.

3.1. Data collection

How do we know which components failed and which did not? This data can be collected from version archives like CVS and bug tracking systems like BUGZILLA in two steps:

1. We identify corrections (or fixes) in version archives: Within the messages that describe changes, we search for references to bug reports such as “Fixed 42233” or “bug #23444”. Basically every number is a potential reference to a bug report; however such references have a low trust at first. We increase the trust level when a message contains keywords such as “fixed” or “bug” or matches patterns like “# and a number”. This approach was previously used in research [3, 6, 17].
2. We use the bug tracking system to map bug reports to releases. Each bug report contains a field called “version” that lists the release for which the bug was reported; however, since the values of this field may change during the life cycle of a bug (e.g., when a bug is carried over to the next release), we only use the first reported release.

We distinguish two different kinds of defects: *pre-release defects* are observed during development and testing of a program, while *post-release defects* are observed after the program has been deployed to its users.

Since we know the location of every defect that has been fixed, it is easy to count the number of defects per location and release.

For the computation of complexity metrics, we used the Java parser of Eclipse. We implemented *visitors* that compute standard metrics (see Table 1) for methods, classes, and files (compilation units) and *aggregators* that combine the metric values into single values for the levels we were interested in (files and packages). For aggregation we used the average, total, and maximum values of the metrics; we omitted minimum values because they are zero in most cases. The source code metrics were computed on the archived builds of Eclipse (see the URL in Figure 1). Note that file level is different from class level since in Java one file can contain several classes.

3.2. Data description

Our data consists of six files in total—one file for each level (files, packages) and release (2.0, 2.1, 3.0). Table 2 summarizes the total number of cases per file. Each case contains the following information:

- **name:** The name of the file or package, respectively, to which this case corresponds. It can be used to identify the source code in the release and may be needed for additional data collection.
- **pre-release defects:** The number of non-trivial defects that were reported in the *last six months before release*.
- **post-release defects:** The number of non-trivial defects that were reported in the *first six months after release*.
- **complexity metrics:** We computed for each case several complexity metrics (see Table 1). Metrics that are computed for classes or methods are aggregate by using average (avg), maximum (max), and accumulation (sum) to file and package level.
- **structure of abstract syntax tree(s):** For each case, we list the size (=number of nodes) of the abstract syntax tree(s) of the file or package, respectively. Abstract syntax trees also consist of different types of nodes (see Figure 2). In addition to size, we also list the frequency of each of these nodes. These counts allow constructing new metrics without any additional processing of the source code.

Table 2. Number of cases

Release	Number of	
	Files	Packages
2.0	6740	376
2.1	7900	433
3.0	6614	429

<i>AnnotationTypeDeclaration</i>	<i>MethodInvocation</i>
<i>AnnotationTypeMemberDeclaration</i>	<i>MethodRef</i>
<i>AnonymousClassDeclaration</i>	<i>MethodRefParameter</i>
<i>ArrayAccess</i>	<i>Modifier</i>
<i>ArrayCreation</i>	<i>NormalAnnotation</i>
<i>ArrayInitializer</i>	<i>NullLiteral</i>
<i>ArrayType</i>	<i>NumberLiteral</i>
<i>AssertStatement</i>	<i>PackageDeclaration</i>
<i>Assignment</i>	<i>ParameterizedType</i>
<i>Block</i>	<i>ParenthesizedExpression</i>
<i>BlockComment</i>	<i>PostfixExpression</i>
<i>BooleanLiteral</i>	<i>PrefixExpression</i>
<i>BreakStatement</i>	<i>PrimitiveType</i>
<i>CastExpression</i>	<i>QualifiedName</i>
<i>CatchClause</i>	<i>QualifiedType</i>
<i>CharacterLiteral</i>	<i>ReturnStatement</i>
<i>ClassInstanceCreation</i>	<i>SimpleName</i>
<i>CompilationUnit</i>	<i>SimpleType</i>
<i>ConditionalExpression</i>	<i>SingleMemberAnnotation</i>
<i>ConstructorInvocation</i>	<i>SingleVariableDeclaration</i>
<i>ContinueStatement</i>	<i>StringLiteral</i>
<i>DoStatement</i>	<i>SuperConstructorInvocation</i>
<i>EmptyStatement</i>	<i>SuperFieldAccess</i>
<i>EnhancedForStatement</i>	<i>SuperMethodInvocation</i>
<i>EnumConstantDeclaration</i>	<i>SwitchCase</i>
<i>EnumDeclaration</i>	<i>SwitchStatement</i>
<i>ExpressionStatement</i>	<i>SynchronizedStatement</i>
<i>FieldAccess</i>	<i>TagElement</i>
<i>FieldDeclaration</i>	<i>TextElement</i>
<i>ForStatement</i>	<i>ThisExpression</i>
<i>IfStatement</i>	<i>ThrowStatement</i>
<i>ImportDeclaration</i>	<i>TryStatement</i>
<i>InfixExpression</i>	<i>TypeDeclaration</i>
<i>Initializer</i>	<i>TypeDeclarationStatement</i>
<i>InstanceofExpression</i>	<i>TypeLiteral</i>
<i>Javadoc</i>	<i>TypeParameter</i>
<i>LabeledStatement</i>	<i>VariableDeclarationExpression</i>
<i>LineComment</i>	<i>VariableDeclarationFragment</i>
<i>MarkerAnnotation</i>	<i>VariableDeclarationStatement</i>
<i>MemberRef</i>	<i>WhileStatement</i>
<i>MemberValuePair</i>	<i>WildcardType</i>
<i>MethodDeclaration</i>	

Figure 2. Abstract syntax tree nodes.

3.3. Data relevance

The Eclipse bug data set can be used to build and assess models for defect prediction. Figure 3 shows a histogram of the number of defects for packages. Most packages (almost 50%) have no observed defects; the remaining packages have up to 103 defects reported.

This distribution calls for two interesting research questions: Which files/packages have defects (a *classification* problem)? And which are the files/packages with the most defects (a *ranking* problem)? Having reliable predictions for both supports allocation of resources for quality assurance (such as testing) to parts of a system that are most defect-prone.

3.3.1. Classification

Classification tries to predict whether a file or package will have at least one defect reported. When applied to all files/packages, the outcome is a classification table such as the following:

		Defects are observed.		
		True	False	
Model predicts defects.	Positive	True Positive (TP)	False Positive (FP)	→ Precision
	Negative	False Negative (FN)	True Negative (TN)	↓ Accuracy
		↓ Recall		

For assessing the quality of a classification model we recommend to use precision, recall, and accuracy:

- **Precision.** The precision relates the number of true positives (predicted and observed as defect-prone) to the number of files/packages predicted as defect-prone.

$$\text{precision} = \frac{TP}{TP + FP}$$

A value close to one is desirable and would mean that every file/package that was predicted to have defects actually had defects.

- **Recall.** The recall relates the number of true positives (predicted and observed as defect-prone) to the number of files/packages that actually had defects.

$$\text{recall} = \frac{TP}{TP + FN}$$

A value close to one is best and would mean that every file/package that had defects observed was predicted to have defects.

- **Accuracy.** The accuracy relates the number of correct classifications (true positives and true negatives) to the total number of files/packages.

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

A value of one is best and would mean that the model classified perfectly, i.e., made not a single mistake.

In order to interpret these measures correctly, one additionally needs to know the percentage of files (or packages) that have defects. Assume that 80% of all files have defects and a model classifies every file as defect-prone. In this case, the model has a precision of 80%, recall of 100%, and accuracy of 80%. Still such a model is not helpful for classification purposes.

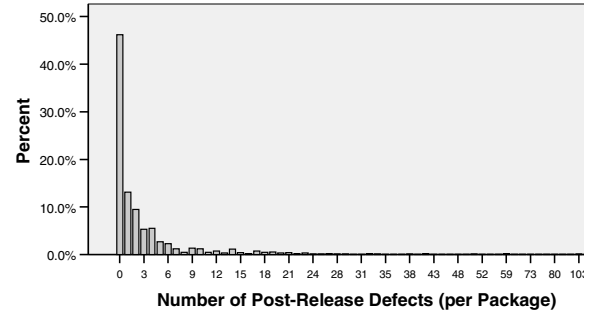


Figure 3. Histogram of post-release defects.

3.3.2. Ranking

Ranking tries to predict an order of files/packages where the files/packages with most defects come first. To measure the quality of such a ranking, we recommend using the Spearman correlation.

- **Spearman correlation.** The Spearman correlation coefficient measures the correlation between a predicted and observed ranking. High correlations and thus a high quality of the predicted ranking are indicated by values close to 1 and -1: values of 1 indicate an identical ranking and values of -1 indicate an opposite ranking. Values close to 0 indicate no correlation.

4. Experiments

In this section, we present a few experiments using the *Eclipse bug data set*. We do not attempt to give definite answers for defect prediction, but merely highlight the potential of bug data when it comes to address this problem.

4.1. Correlations

At first, we computed the Spearman correlation between the number of **pre-release and post-release defects** and the complexity metrics in the data set. Table 3 lists the correlation values for both, file and package level (release 3.0). Correlations significant at the 0.05 level are marked with (*), the ones significant at 0.01 are marked with (**).

The high correlation values (0.907 for files, 0.921 for packages) between the number of pre-release and post-release defects indicate that the files/packages having the most pre-release defects are likely to also have the most post-release defects and vice versa.

Most correlations with metrics are positive and significant—the more complex a file/package the more defects it will have. However, on **file level** only the

accumulated McCabe complexity (VG_sum), the total lines of code (TLOC), and the closely related sum of method lines of code (MLOC_sum) have correlation values above 0.400. This indicates that the size of files (or of control flow graphs for McCabe) seems to be a good indicator for defect-prone files.

On **package level**, more correlations are above the 0.400 threshold: fan out as measured by the number of calls in a method (FOUT), nested block depth (NBD), number of non-static fields (NOF) and methods (NOM), and the number of files in a package (NOCU). Most accumulated metrics (_sum) have high correlations values (except NSF and NSM, the number of static fields and methods, respectively). However, in many cases the maximum metrics (_max) show correlation values comparable to the accumulated ones (except for PAR_max).

Finding a single indicator or predictor for the number of defects is extremely unlikely. In the next subsections, we will combine input features by building regression models (linear for ranking, and logistic for classification).

4.2. Classification

The previous section showed that individual code metrics correlate with the number of defects. But how can we use metrics to predict whether a file/package will have defects? One solution is to build statistical models that classify files/packages as defect-prone (*has_defects=1*) or not (*has_defects=0*) based on the values of the code metrics.

We built logistic regression models for the *Eclipse bug data set* to predict whether files/packages have **post-release defects**. Logistic regression models typically predict likelihoods (between 0 and 1); when the predicted likelihood was above 0.5, we classified a file/package as defect-prone, otherwise as defect-free. In total we built six models for two levels of granularity (files, packages) and three releases (2.0, 2.1, 3.0). We tested the models across releases of Eclipse, but always on the level they were built from.

Table 4 lists the precision, recall, and accuracy values for the **file level**. The recall values are low for all tests, meaning that only few of the defect-prone files were correctly identified as defect-prone. However, the precision values are above 0.600 in most cases, suggesting that there are only few false positives, i.e., when a file is classified as defect-prone, this decision is most likely to be correct. Most of the precision, recall, and accuracy values remain comparable across releases. This means that a model learned from one release, say 2.0, can be applied to a later release, say 3.0, without losing too much predictive power (only a decrease of 0.084 in accuracy).

Table 3. Spearman correlation between pre- and post-release defects and metrics.

Release 3.0	File level		Package level	
	pre	post	pre	post
pre	1.000	.907(**)	1.000	.921(**)
post	.907(**)	1.000	.921(**)	1.000
FOUT_avg	.298(**)	.301(**)	.190(**)	.173(**)
FOUT_max	.356(**)	.365(**)	.414(**)	.368(**)
FOUT_sum	.377(**)	.387(**)	.434(**)	.371(**)
MLOC_avg	.324(**)	.326(**)	.279(**)	.307(**)
MLOC_max	.383(**)	.388(**)	.442(**)	.433(**)
MLOC_sum	.394(**)	.405(**)	.475(**)	.428(**)
NBD_avg	.323(**)	.325(**)	.261(**)	.269(**)
NBD_max	.377(**)	.383(**)	.483(**)	.467(**)
NBD_sum	.378(**)	.392(**)	.464(**)	.404(**)
NOF_avg	.241(**)	.251(**)	.264(**)	.264(**)
NOF_max	.248(**)	.256(**)	.431(**)	.396(**)
NOF_sum	.251(**)	.257(**)	.445(**)	.404(**)
NOM_avg	.289(**)	.313(**)	.273(**)	.232(**)
NOM_max	.300(**)	.321(**)	.415(**)	.348(**)
NOM_sum	.303(**)	.321(**)	.447(**)	.376(**)
NSF_avg	.166(**)	.196(**)	.246(**)	.226(**)
NSF_max	.174(**)	.202(**)	.352(**)	.314(**)
NSF_sum	.174(**)	.203(**)	.387(**)	.341(**)
NSM_avg	.112(**)	.131(**)	.149(**)	.134(**)
NSM_max	.116(**)	.134(**)	.298(**)	.259(**)
NSM_sum	.116(**)	.134(**)	.303(**)	.264(**)
PAR_avg	.114(**)	.111(**)	.057	.082
PAR_max	.256(**)	.263(**)	.298(**)	.276(**)
PAR_sum	.339(**)	.353(**)	.447(**)	.387(**)
VG_avg	.322(**)	.331(**)	.300(**)	.334(**)
VG_max	.373(**)	.385(**)	.451(**)	.451(**)
VG_sum	.385(**)	.401(**)	.475(**)	.431(**)
NOCU			.406(**)	.344(**)
ACD	.191(**)	.164(**)		
ACD_avg			.251(**)	.164(**)
ACD_max			.327(**)	.235(**)
ACD_sum			.338(**)	.247(**)
NOI	-.170(**)	-.179(**)		
NOI_avg			.016	.006
NOI_max			.134(**)	.121(*)
NOI_sum			.138(**)	.107(*)
NOT	.171(**)	.180(**)		
NOT_avg			.000	.010
NOT_max			.198(**)	.205(**)
NOT_sum			.385(**)	.336(**)
TLOC	.407(**)	.420(**)		
TLOC_avg			.324(**)	.309(**)
TLOC_max			.461(**)	.419(**)
TLOC_sum			.487(**)	.433(**)

Correlation significant at the (**) 0.01 level, (*) 0.05

Table 5 presents the results for the **package level**. The results are almost the same as for file level (precision slightly increases, accuracy slightly decreases), however, the recall increases substantially. Intuitively, it is easier to predict defect-prone packages, as already one defect-prone file makes a package defect-prone.

4.3. Ranking

In order to predict the files/packages that have most **post-release defects** we used linear regression models. Using these models we predicted for each file/package the number of expected post-release defects and compared the resulting ranking to the observed ranking using Spearman correlation. Again, we built models for each level (file, package) and release (2.0, 2.1, 3.0) and tested them across releases.

Table 6 shows the results for **file level**. In addition to Spearman, we also list the R^2 value of the trained model and the Pearson correlation. The R^2 value measures the variability in a data set that is accounted for by a statistical model. In our case only up to 41% of variance are explained by complexity metrics. Pearson correlation assumes a linear relation between the correlated variables; we report Pearson values only for completeness. The Spearman correlation values ρ are low, reaching 0.416 at most, but positive. Therefore, files having a higher predicted rank are more likely to have a high observed rank, too.

The results for **package level** are listed in Table 7. The R^2 values substantially improve; for release 3.0 81.2% of variability is explained by the linear model. The Spearman correlation values slightly increase to up to 0.449. Again the increase in correlation and R^2 values demonstrates that it is easier to make predictions for packages than for files. Also, models learned from earlier releases can be used to predict for future releases; for instance the model trained from release 2.1 showed a correlation of 0.449 on release 3.0.

5. Conclusions

Where do bugs come from? By mapping failures to components, our *Eclipse bug data set* offers the opportunity to research this question. The experiments in this paper showed that the combination of complexity metrics can predict defects, suggesting that the more complex code it, the more defects it has.

However, our predictions are far from being perfect. They therefore raise follow-up questions: Are there better indicators for defects than complexity metrics? How applicable are models across projects and over time? How do we integrate prediction models into the development process?

**Table 4. Classification of files.
(Precision P , Recall R , and Accuracy Acc).**

Training	Testing	Defects	P	R	Acc
2.0	2.0	0.265	0.657	0.242	0.766
	2.1	0.234	0.578	0.259	0.783
	3.0	0.355	0.638	0.244	0.682
2.1	2.0	0.265	0.673	0.185	0.760
	2.1	0.234	0.645	0.219	0.789
	3.0	0.355	0.687	0.195	0.682
3.0	2.0	0.265	0.476	0.330	0.726
	2.1	0.234	0.453	0.369	0.748
	3.0	0.355	0.663	0.379	0.711

**Table 5. Classification of packages.
(Precision P , Recall R , and Accuracy Acc).**

Training	Testing	Defects	P	R	Acc
2.0	2.0	0.536	0.742	0.735	0.721
	2.1	0.510	0.667	0.732	0.677
	3.0	0.570	0.641	0.724	0.612
2.1	2.0	0.536	0.733	0.617	0.675
	2.1	0.510	0.720	0.700	0.708
	3.0	0.570	0.713	0.664	0.656
3.0	2.0	0.536	0.653	0.719	0.645
	2.1	0.510	0.632	0.751	0.651
	3.0	0.570	0.785	0.789	0.757

**Table 6. Ranking files with linear regression.
(Pearson r and Spearman ρ correlation)**

Files	Testing		
	2.0	2.1	3.0
2.0	$r=0.491$	$r=0.492$	$r=0.470$
	$R^2=0.241$	$\rho=0.309$	$\rho=0.347$
2.1	$r=0.456$	$r=0.536$	$r=0.553$
	$R^2=0.287$	$\rho=0.307$	$\rho=0.380$
3.0	$r=0.331$	$r=0.422$	$r=0.640$
	$R^2=0.410$	$\rho=0.225$	$\rho=0.416$

All correlations are significant at the 0.01 level.

**Table 7. Ranking packages with linear regression.
(Pearson r and Spearman ρ correlation)**

Packages	Testing		
	2.0	2.1	3.0
2.0	$r=0.614$	$r=0.490$	$r=0.368$
	$R^2=0.378$	$\rho=0.356$	$\rho=0.319$
2.1	$r=0.391$	$r=0.761$	$r=0.722$
	$R^2=0.580$	$\rho=0.387$	$\rho=0.449$
3.0	$r=0.120$	$r=0.369$	$r=0.901$
	$R^2=0.812$	$\rho=0.238$	$\rho=0.330$

All correlations are significant at the 0.01 level.

The above questions indicate the potential of future empirical research based on bug data. To support this very research, we are happy to make the *Eclipse bug data* set publicly available.

Overall, we would like this dataset to become both a challenge and a benchmark: Which factors in programs and processes are predictors of future bugs, and which approach gives the best prediction results? The more we learn about past mistakes, the better are our chances to avoid them in the future—and build better software at lower cost.

For access to the Eclipse bug data set, as well as for ongoing information on the project, see

<http://www.st.cs.uni-sb.de/softevo/>

Acknowledgments. Our work on mining software repositories is funded by Deutsche Forschungsgemeinschaft, grant Ze 509/1-1. Thomas Zimmermann is additionally funded by the DFG-Graduiertenkolleg-Leistungsgarantien für Rechnersysteme. Thanks to Adrian Schröter for computing the initial version of the *Eclipse bug data set*.

6. References

- [1] V. R. Basili, L. C. Briand, and W. L. Melo, "A validation of object-oriented design metrics as quality indicators." *IEEE Transactions on Software Engineering* vol. 22, pp. 751-761, 1996.
- [2] A. B. Binkley and S. R. Schach, "Validation of the coupling dependency metric as a predictor of run-time failures and maintenance measures." in *Proceedings of the International Conference on Software Engineering*, 1998, pp. 452-455.
- [3] D. Cubranic and G. C. Murphy, "Hipikat: Recommending pertinent software development artifacts." in *25th International Conference on Software Engineering (ICSE)*, Portland, Oregon, 2003, pp. 408-418.
- [4] G. Denaro, S. Morasca, and M. Pezzè, "Deriving models of software fault-proneness." in *Proceedings of the 14th International Conference on Software Engineering and Knowledge Engineering* Ischia, Italy, 2002 pp. 361 - 368.
- [5] G. Denaro and M. Pezzè, "An empirical evaluation of fault-proneness models." in *Proceedings of the International Conference on Software Engineering (ICSE 2002)*, Orlando, Florida, USA, 2002, pp. 241-251.
- [6] M. Fischer, M. Pinzger, and H. Gall, "Populating a release history database from version control and bug tracking systems." in *Proc. International Conference on Software Maintenance (ICSM 2003)*, Amsterdam, Netherlands, 2003.
- [7] T. L. Graves, A. F. Karr, J. S. Marron, and H. Siy, "Predicting fault incidence using software change history." *IEEE Transactions on Software Engineering*, vol. 26, 2000.
- [8] J. P. Hudepohl, S. J. Aud, T. M. Khoshgoftaar, E. B. Allen, and J. Mayrand, "Emerald: Software metrics and models on the desktop." *IEEE Software*, vol. 13, pp. 56-60, September 1996.
- [9] T. M. Khoshgoftaar, E. B. Allen, N. Goel, A. Nandi, and J. McMullan, "Detection of software modules with high debug code churn in a very large legacy system." in *ISSRE '96: Proceedings of the The Seventh International Symposium on Software Reliability Engineering (ISSRE '96)*, Washington, DC, USA, 1996, p. 364.
- [10] N. Nagappan and T. Ball, "Explaining failures using software dependences and churn metrics," Microsoft Research, Redmond, WA 2006.
- [11] N. Nagappan and T. Ball, "Use of relative code churn measures to predict system defect density." in *Proceedings of the International Conference on Software Engineering (ICSE 2005)*, St. Louis, Missouri, USA, 2005, pp. 284-292.
- [12] N. Nagappan, T. Ball, and A. Zeller, "Mining metrics to predict component failures." in *Proceedings of the International Conference on Software Engineering (ICSE 2006)*, Shanghai, China, 2006.
- [13] N. Ohlsson and H. Alberg, "Predicting fault-prone software modules in telephone switches." *IEEE Trans. Software Eng.*, vol. 22, pp. 886-894, 1996.
- [14] T. J. Ostrand, E. J. Weyuker, and R. M. Bell, "Predicting the location and number of faults in large software systems." *IEEE Trans. Software Eng.*, vol. 31, pp. 340-355, 2005.
- [15] A. Schröter, T. Zimmermann, R. Premraj, and A. Zeller, "If your bug database could talk..." in *Proceedings of the 5th International Symposium on Empirical Software Engineering. Volume II: Short Papers and Posters*, 2006, pp. 18-20.
- [16] A. Schröter, T. Zimmermann, and A. Zeller, "Predicting failure-prone components at design time." in *Proceedings of the 5th International Symposium on Empirical Software Engineering (ISESE 2006)*, Rio de Janeiro, Brazil, 2006.
- [17] J. Śliwerski, T. Zimmermann, and A. Zeller, "When do changes induce fixes? On Fridays." in *Proc. International Workshop on Mining Software Repositories (MSR)*, St. Louis, Missouri, U.S., 2005.
- [18] R. Subramanyam and M. S. Krishnan, "Empirical analysis of ck metrics for object-oriented design complexity: Implications for software defects." *IEEE Trans. Software Eng.*, vol. 29, pp. 297-310, 2003.