

Software Metrics: Calculation and Optimization of Thresholds

AUTHOR

Abhishek Kumar Maheswari
Faculty of Computer Science
Venkateshwara Group of Institutions,
Meerut.

Abstract

In this article, we present a algorithmic method for the calculation of thresholds (the starting point for a new state) for a software metric set. To this aim, machine learning and data mining techniques are utilized. We define a data-driven methodology that can be used for efficiency optimization of existing metric sets, for the simplification of complex classification models, and for the calculation of thresholds for a metric set in an environment where no metric set yet exists. The methodology is independent of the metric set and therefore also independent of any language, paradigm or abstraction level. In four case studies performed on large-scale open-source software metric sets for C functions, C++, C# methods and Java classes are optimized and the methodology is validated.

1. INTRODUCTION

Software has become part of the everyday life. Embedded software in modern cars controls the distance to the car in front of us. News portals on the Internet utilize sophisticated distributed software to report the news events as they occur. Users expect and need software to conform to a certain standard of quality. The *International Organization for Standardization (ISO)* defines quality as the “*degree to which a set of inherent characteristics fulfills requirements*” in the ISO 9000 standard. To uphold the required standard of quality, the assurance that software quality attributes are fulfilled is an important aspect of the execution of software projects. Quality attributes like maintainability and understandability are often assessed using *software metrics*. Software metrics provide means to put numbers on abstract attributes, such as complexity or size. Often, one metric is insufficient to effectively analyze a quality attribute. Instead, we use a set of metrics to determine whether a quality attribute is fulfilled or problematic. To determine if metric values are good or bad, clear indicators are required. Otherwise such metric sets are hard to interpret. For this purpose, we use thresholds for metric values: a quality attribute is said to be problematic, when at least one threshold for a metric is violated. For thresholds to be effective indicators, the quality of the threshold values themselves is of great importance. However, the thresholds often depend on the project environment, e.g., programming languages and tool support. Therefore, the definition of thresholds is often problematic and defined thresholds may not be valid in other environments.

During the last years, machine learning has been successfully applied and

has become a standard technique for data analysis in many different fields, such as gene analysis in biology, or data mining techniques companies use to optimize their marketing strategies. It has also been used in computer science, e.g., for defect prediction. In this article, we introduce an algorithmic approach for the optimization of the size software metric sets and threshold values used. To this aim, a machine learning algorithm is used to define an approach for the calculation of thresholds for a metric set. In a previous work, we used relatively simple brute-force approach for the calculation of threshold values for a metric set for the *Testing and Test Control Notation* (TTCN). However, such a brute force approach has scalability problems and is therefore infeasible for larger metric sets. This work presents a more sophisticated approach, which utilizes the learning of axis-aligned d -dimensional rectangles for the threshold calculation. The objective of this work is to reduce the complexity of metric-based classifiers for software quality to improve their understandability and interpretability, which will benefit both researchers and the industry as it allows pinpointing the source of deficits more effectively. To this end, we provide a versatile, data-driven means for both threshold calculation and the optimization of metric sets integrated into a single algorithm.

The contribution of this article is fourfold.

1. A machine learning based method for the computation of thresholds for metric sets.
2. A high-level methodology for the optimization of already existing metric sets with thresholds.
3. Using the same methodology to effectively replace existing classification methods, and thereby reducing their complexity.
4. An outline how a good metric set with thresholds can be determined in an environment where no thresholds exist yet.

For the first contribution, we show how the problem of rectangle-learning relates to sets of software metrics with thresholds and how rectangle learning can be utilized to compute thresholds. The second contribution defines a generic methodology for metric set efficiency optimization, which is in fact

not restricted to software metrics, but applicable to metric sets in general. For this, we assume that an effective metric set with thresholds already exists. We show how a smaller and effective set can be determined, which is due to its reduced size also efficient. The third contribution shows how this approach can be used to replace existing classification strategies that may not even be metric-based, with a threshold based classification. Such a replacement can be used to substitute hard-to-interpret or black-box classifiers with easy-to-interpret threshold classification. Finally, we show how our approach can be used to determine a good metric set in an environment where no means for the automated classification of software entities exists yet. In comparison to the first contribution, this includes not only the calculation of the threshold values, but also the selection of an appropriate subset of metrics from a possibly large set of candidate metrics.

All methodologies defined in course of this article are independent of the metric sets themselves and only depend on actually observed data. The methods are therefore independent of any specific programming language (e.g., C, Java) and level of abstraction (e.g., methods, classes). In four case studies, we validated that the approach works well for product metrics in large-scale open-source software projects. As part of the case studies, metric sets for C function, C++ and C# methods, and Java classes are analyzed.

The structure of this article is as follows. In Section 2, we introduce the concepts of software metrics and how they can be used in combination with thresholds for quality estimation. Afterwards, we briefly introduce machine learning and define the foundations of the learning approach used in this article in Section 3. In Section 4, we define the methodology for the optimization of software metric sets with thresholds and provide a description of how it can be applied to perform different tasks is.

2. SOFTWARE METRICS

According to Fenton and Pfleeger, “*Measurement is the process by which numbers or*

symbols are assigned to attributes of entities in the real world in such a way as to describe them according to clearly defined rules". A way to measure software is to use *software metrics*. The IEEE defines software metrics as "the quantitative measure of the degree to which a system, component, or process possess a given attribute". This means that software metric is a clearly defined rule that assigns values to software entities (e.g., components, classes, or methods) or attributes of development processes.

Fenton and Pfleeger divided software metrics into three categories: *process metrics* measure attributes of a development process itself; *product metrics* measure documents and software artifacts that were produced as part of a process; *resource metrics* measure the resources, which were utilized as part of a process. Furthermore, each metric measures either an *internal* or an *external* attribute. Internal attributes are those that can be measured by observing only the process, product or resource itself, without considering its behavior. External attributes on the other hand are attributes that are related to the behavior of software systems. In this work, the focus is on internal product metrics that measure source code. Some examples for internal attributes that relate to source code are size, reuse, modularity, algorithmic complexity, coupling, functionality, and control-flow structuredness. Further attributes are staticness, method complexity, or attributes that relate to object-oriented software, such as usage of inheritance.

2.1 METRIC SETS UNDER STUDY

The methods described in this article are general and may be used independent of a specific metric set. However, as part of this article, metric sets for the evaluation of the maintainability are studied exemplary. This is done with two different metric sets on different levels of abstraction: methods and classes. The maintainability describes non-functional aspects such as testability, understandability, or changeability of software. Because no single metric is able to cover all these aspects, we employ a set of metrics that covers

internal attributes like the structure, size, and complexity instead. We selected the metrics based on our experience and with the aim to cover the maintenance related aspects of the source code that can be measured automatically with internal product metrics.

For the analysis of methods and functions, we use four metrics listed in Table 1. With the control-flow structuredness measured by the *Cyclomatic Number* (VG) and *Nested Block Depth* (NBD), coupling measured by *Number of Function Calls* (NFC), and size measured by *Number of Statements* (NST), these metrics cover most of the maintainability-related attributes of methods, except the algorithmic complexity. Since algorithmic complexity is not really an attribute of the source code and cannot be measured automatically, it has been omitted. While both VG and NBD measure the control-flow structuredness, they measure different aspects of this attribute: NBD measures the maximum nesting of structural blocks, while (VG) measures the overall branching between blocks.

For the analysis of classes, the seven metrics listed in Table 1b are used. With these metrics, five internal attributes of classes are evaluated. The metric *Weighted Methods per Class* (WMC) measures the method complexity as the sum of the metric VG measured for all methods in a class. The metrics *Coupling Between Objects* (CBO) and *Response For a Class* (RFC) measure the coupling. For the measurement of the size of a class, the metrics *Number of Methods* (NOM) and *Lines of Code* (LOC) are utilized. The use of inheritance is measured by *Number of Overridden Methods* (NORM), the staticness of a class is measured by the metric *Number of Static Methods* (NSM). We included the attributes inheritance and staticness, as they greatly influence the maintainability of classes (Daly et al. 1996). Inheritance is often difficult to test and also decreases the understandability of the source code. Static methods and attributes can pose problems, as they are global for all instances of a class and can therefore introduce unwanted side effects.

Table-1
Metrics used in this Research Paper

Metric name	Internal attribute	Description
(a) Metrics for methods and functions		
<i>Cyclomatic Number</i> (VG)	Control-flow structuredness	Calculated based on the control flow graph $G = (V,E)$ and number of a method M as $VG(M) = E - V + p$, where p is the number of entries and exits.
<i>Nested Block Depth</i> (NBD)	Control-flow structuredness	Maximum number of nested blocks in a method.
<i>Number of Function Calls</i> (NFC)	Coupling	Number of functions called by a method
<i>Number of Statements</i> (NST)	Size	Number of statements of a method
(b) Metrics for classes		
<i>Weighted Methods per Class</i> (WMC)	Method complexity	Complexity of a class as the sum of the complexity of its methods. Here, VG is used as complexity measure.
<i>Coupling Between Objects</i> (CBO)	Coupling	Number of classes, to which a class is coupled.
<i>Response For a Class</i> (RFC)	Coupling	Size of the <i>response set</i> of a class, i.e. all methods that can be invoked directly or indirectly by calling a method of a class.
<i>Number of Overridden Methods</i> (NORM)	Inheritance	Number of methods defined by a parent that are overridden by a class
<i>Number of Methods</i> (NOM)	Size	Number of methods of a class
<i>Lines of Code</i> (LOC)	Size	Lines of code, excluding empty and comment-only lines.
<i>Number of Static Methods</i> (NSM)	Staticness	Number of static methods of a class

2.2 THRESHOLDS FOR SOFTWARE METRICS

In general, thresholds discriminate values. In case a threshold defines an upper bound, the values that are greater than a threshold value are considered to be problematic, the values lower are considered to be acceptable. Thus, by defining thresholds a simple analysis of measured values is possible. For the interpretation of software metrics thresholds are required. For example, consider a metric m that measures the size of an entity x . Then a threshold t can be used to determine if x is too large:

$$m(x) > t \Rightarrow x \text{ is too large}$$

While the above is an example of a threshold used as an upper bound, it might as well be a lower bound. For simplicity, we assume that thresholds are always upper bounds. However, this is no restriction

as lower bounds can be transformed into upper bounds. Let m a metric with threshold t that defines a lower bound, i.e., entities x are considered to be problematic if $m(x) < t$, which is equivalent to $1/m(x) > 1/t$ if $m(x)$ and t are non-negative, as metrics and thresholds usually are. By defining a new metric $m_2(x) = 1/m(x)$ and a new threshold $t_2 = 1/t$ a new metric with the opposite order is defined and with t_2 a threshold is obtained that defines an upper bound. However, by inverting the metric, its scale is hanged. Another way to transform a lower bound into an upper bound while keeping it to scale is to subtract the metric from a maximum value. Let m_{\max} the maximum value of metric m . Then

and a new metric $m_2(x) = m_{\max} - m(x)$ and a new threshold $t_2 = m_{\max} - t$ are obtained,

where t_2 is an upper bound for m_2 . However, this method has the disadvantage that a maximum value m_{\max} has to be known.

Thresholds are not without problems. The first is the generality of threshold values. In general, a threshold value is good in one setting must not necessarily be good every setting. Depending on the organization, the programming language, the tools used the qualification of the developers, among other factors that are project dependent, good threshold values may vary. This is a problem, as each organization, and maybe even each project, has to define thresholds that are chosen depending on its environment. This issue directly relates to a second issue, as good thresholds depend on so many factors, the definition of thresholds itself is a problem. Therefore, a methodology to determine environment specific thresholds is required.

To allow a more differentiated analysis more than one threshold value for one metric can be defined. In this article, we assume source code to be either problematic or un-problematic. However, further shades of gray exist in between. For example, there may be two thresholds, a low one for *weak infractions* and a higher one for *critical infractions*.

3. FOUNDATIONS OF MACHINE LEARNING

In this section, we introduce the concepts of machine learning essential for this work. After a brief description of machine learning in general, we define the learning framework used in this work in Section 3.1. Finally, we discuss an algorithm to learn axis-aligned d -dimensional rectangles in Section 3.2. The approach for the optimization of metric sets is based on this algorithm.

In general, machine learning is a way to analyze data. Learning theory assumes that observed data can be described by an underlying process. The type of the process varies and depends on the type of learning. For example, it could be an automaton, but also a stochastic process. The aim of machine learning is to identify this process. Often, this is not accurately possible. However, in most cases it is still

possible to detect patterns within the data. Assuming that the underlying stochastic process does not change, it is possible to predict properties of unseen data using the detected patterns. A more detailed introduction to machine learning in general can be found in the literature.

3.1 CONCEPT LEARNING IN THE PRESENCE OF NOISE

In this work, we use *concept learning*. A concept defines how to divide vectors from the \mathbb{R}^d into positive and negative examples. The task of a learning algorithm is to infer a *target concept* g out of a *concept class* \mathcal{C} .

The target concept can also be interpreted as the *Bayesian classifier* of the concept. A concept can also be understood as a map $g: \mathbb{X}^d \rightarrow \{0,1\}$, where $\mathbb{X}^d \subset \mathbb{R}^d$ denotes the *input space*. A learning sample is of the form (X, Y) where the *input element* X is randomly distributed according to the *sample distribution* D defined over the input space \mathbb{X} , Y is the *random label* or *output element* associated with X . In a noise free setting, the value of Y depends only on the random vector X and the target concept g and $Y = g(X)$. To obtain samples U , the concept of an *oracle* is used. On request, an oracle $EX(D, g)$ randomly draws an input element X according to the distribution D , classifies X using g and returns a sample $U = (X, g(X))$. In practical applications, the oracle can be seen as a training sample that contains classified entities to be used for the learning.

Real-life applications are usually not noise-free, i.e., the property $Y = g(X)$ is not always fulfilled. Most algorithms designed to work on noise-free data often perform poorly or do not work at all in the presence of noise. Therefore, noise modeling and algorithms that use these models for learning in the presence of noise are important. One way to introduce noise into a learning model is the *classification noise* model. In the classification noise model, the label of the output variable Y is changed with a fixed probability and $Y = g(X) \oplus S$, where \oplus denotes the symmetric difference. The *random noise* $S \in \{0,1\}$ is 1 with probability ζ , i.e., $\mathbb{P}(S=1) = \zeta$, where ζ denotes the *noise rate*. In the In the

classification noise model, S is independent of the input element X . It follows directly that $P(Y = g(X)) = \zeta$. In combination with oracles, noise can be seen as an attacker that corrupts the output element of a sample generated by an oracle.

In the SQM proposed by Kearns *query functions* of the form $x: x^d \times \{0,1\} \rightarrow [a,b]$ are used to infer information about the data. For this purpose, a statistical oracle is introduced that returns the expected result of the queries within a specified degree of precision. The estimation is based on noise models.

Attacker
 $P(S=1)=\zeta$

Attacker
 $P(S=1)=\zeta(x)$

Oracle $E(X, g(X)) \rightarrow Y = g(x) \oplus S - (X, Y) \rightarrow$ Algorithm

We use a generalization of the classification noise model, where the random noise rate is orthogonal to the target concept. The restriction that S is independent of the input element X is dropped. Instead, we introduce a *random noise rate* $\zeta(X)$ that depends on the input element, as shown in Fig. 1b. Hence, $\zeta(x) = P(S=1|X=x)$ and thus the random noise depends on the input. For a given $x \in X^d$, the noise rate is $\zeta(x) = P(Y = g(X)|X=x)$ and for $y_0 \in \{0,1\}$ the *conditional expected noise rate* given $g(X) = y_0$ is $n_{\zeta} := E(n(X) | g(X) = y_0)$

Using the conditional expected noise rates ζ_0, ζ_1 , and the *expected noise rate* can be calculated as $n_{\zeta} := E(n(X)) = \zeta_0 P(g(X) = 0) + \zeta_1 P(g(X) = 1)$

Furthermore, we assume that query functions are *admissible*. A query function \div is admissible, if it is not correlated to the noise rate $\zeta(X)$ conditioned on the concept $g(X)$. The geometrical uncorrelation is orthogonality, hence it is said that the noise is orthogonal to the target concept. For the learning, this means that it is not possible to infer the value of \div by simply considering the noise rate $\zeta(X)$. This is a reasonable assumption, as usually no information about the result of a query is obtained by simply considering the noise rate.

Based on the introduced concepts and definitions, we can state the central theorem of the learning framework. This theorem describes how the expected value of an admissible query can be calculated if the conditional expected noise rates ζ_0 and ζ_1 are known.

3.2 A RECTANGLE LEARNING ALGORITHM

In this work, we adapted the algorithm for learning axis-aligned d -dimensional rectangles

proposed to the noise model described above. The main adaptations are that the conditional expected noise rates ζ_0 and ζ_1 have both to be sampled, instead of only the expected noise rate ζ . Furthermore, the statistical oracle used by the algorithm is changed from the SQM to the random noise model by calculating the expected results of statistical queries based on Theorem 1. The algorithm has two phases. In the first phase, the training data is partitioned according to its distribution. In the second phase, the rectangle is computed based on this partition. Both phases are described in the following.

The aim of the first phase of the algorithm is to find a partition of the d -dimensional real-space, such that $P(X_i \in I_{i,p}) = P(X_i \in I_{i,q}) = \frac{1}{|I_{i,\epsilon}|} \approx \epsilon$ for each dimension $i=1, \dots, d$ and $p, q=1, \dots, \#I_{i,\epsilon}$ for $X \in \mathbb{R}^d$ randomly distributed according to \mathbf{D} , where X_i denotes the i -th component of X and ϵ an *error bound* that the calculate hypothesis should abide. This means that it is equally likely that the i -th component of the randomly drawn vector X falls into any of the intervals $I_{i,\epsilon}$. In the implementation of the algorithm, a sorting algorithm is utilized to obtain these intervals according to the empirical distribution

of a discrete training sample. After sorting the values for the i -th dimension, the intervals $I_{i,p}$ can be defined by assigning the first $\frac{p}{\lfloor 1/\epsilon \rfloor}$ vectors to $I_{i,1}$, the next $\frac{p}{\lfloor 1/\epsilon \rfloor}$ to $I_{i,2}$ and so on. These intervals fulfill the property defined by (3.6). If there are n samples in a training set, the complexity of the first phase is $O(dn \log n)$, as for each dimension the samples have to be sorted and efficient sorting algorithms are $O(n \log n)$.

In the second phase, the boundaries of the target rectangle are calculated. For each dimension separately, the probability $PI_{i,p} = P(X_i \in I_{i,p} / g(X) = 1)$, i.e., the probability that the target rectangle intersects an interval $I_{i,p}$ is calculated. This probability is calculated using admissible queries and (3.5). If the target rectangle intersects an interval, the probability $PI_{i,p}$ should be significantly larger than 0. Thus, for each dimension i , the probabilities $PI_{i,p}$ are calculated from the left, i.e., $p = 1, 2, \dots$. The first interval, for which is significant defines the left, i.e., lower boundary l_i of the rectangle in the i -th dimension. The same is done from the right, i.e., $p = \#1/\epsilon, \dots, \#1/\epsilon$ to determine the right, i.e., upper boundary u_i . Using this procedure for each dimension, boundaries (l_i, u_i) are calculated.

In the second phase, for each dimension, the probability $PI_{i,p}$ is calculated for at most $\#1/\epsilon$ intervals from the left and analogously from the right. The estimation of this probability is $O(n)$. Thus the complexity of the second phase is $O(dn \frac{1}{\epsilon})$ and the overall complexity of the algorithm is $O(dn \log n + dn \frac{1}{\epsilon})$.

4. OPTIMIZATION OF METRIC SETS AND THRESHOLDS

In this section, we introduce our machine learning based approach to optimize metric sets with thresholds for the detection of problematic entities. First, we describe in Section 4.1 how the rectangle learning algorithm is utilized to calculate thresholds. Based on that, we define a threshold optimization algorithm for the calculation of an optimized metric set with thresholds. Then we show three applications for this threshold optimization algorithm:

- 1) Optimization of an existing metric set with thresholds to obtain an effective and efficient subset;

- 2) Reduction of the complexity of the used classification method;
- 3) Determination of environment specific thresholds.

4.1 CALCULATION OF THRESHOLDS USING RECTANGLE LEARNING

The analysis approach is based on a given metric set $M = \{m_1, \dots, m_d\}$ and a set of software entities X with known classifications Y . The aim is to obtain thresholds $T = \{t_1, \dots, t_d\}$ for the metrics in M such that the metric set can be used to discriminate software entities in the same way, as it is done by the pair (X, Y) . By measuring the software entities X with M , we transform the set of software entities X into a set of vectors in the d -dimensional real space, such $M(X) := \{(m_1(x), \dots, m_d(x)) : x \in X\} \subset \mathbb{R}^d$

The pair $(M(X), Y)$ is the input for the axis-aligned rectangle learning algorithm, introduced in Section 3.2. As result, the algorithm yields pairs of upper and lower bounds (l_i, u_i) for each dimension $i = 1, \dots, d$. As the i -th dimension represents the values the software entities calculated using the metric m_i and under the assumption that high metric values are bad, we interpret the upper bound of the rectangle in the i -th dimension as the threshold for the metric m_i . Therefore, with $t_i = u_i$ a set of thresholds $T = \{t_1, \dots, t_d\}$ for the metric set M is obtained. For an entity x , the classification of the

Algorithm 1: Algorithm for the calculation of thresholds
 Input : Set of software entities X with classifications Y , metric set $M = \{m_1, \dots, m_d\}$
 Output : thresholds $T = \{t_1, \dots, t_d\}$
 $M(X) \leftarrow \{(m_1(x), \dots, m_d(x)) : x \in X\}$;
 Apply the rectangle learning algorithm to $(M(X), Y)$ and obtain $(u_i, l_i), i = 1, \dots, d$;
 $t_i \leftarrow u_i$, for all $i = 1, \dots, d$;
 $T \leftarrow \{t_1, \dots, t_d\}$;
 return T

metric set M and the thresholds T is defined as

$$fo(x, M, T) = \begin{cases} 1 & \text{if } \{i \in \{1, \dots, d\} : m_i(x) > t_i\} = \emptyset \\ 0 & \text{if } \{i \in \{1, \dots, d\} : m_i(x) > t_i\} \neq \emptyset \end{cases}$$

i.e., $f_0(x, M, T)$ is zero when at least one metric m_i exceeds its threshold t_i , and is one when none of the metrics exceeds its threshold.

Under the assumption that metric values are positive, this classification describes a rectangle with upper bounds t_i and 0 as the lower bound.

The *classification error* is defined as the probability that a randomly drawn sample (X,Y) is classified wrongly $\epsilon = P(f_o(X, M, T) \neq Y)$, Consequently, the *empirical classification error* on a given training sample (X,Y) is defined as

4.2 THRESHOLD AND METRIC SET OPTIMIZATION ALGORITHM

Next, we define a threshold optimization algorithm that computes an optimized metric set based on the calculation of thresholds for a metric set. This means a metric set that is not only effective with respect to the classification it yields, but also efficient in terms of its size. To achieve this, we reduce the dimension of the metric set and recalculate the threshold values for the reduced sets. Recalculating the thresholds allows the algorithm to, e.g., enforce a stronger classification using one metric while dropping another from the set.

The algorithm uses an existing method f for the classification of software entities X . By applying f to the entities $x \in X$, the classification Y can be calculated as $Y = \{f(x) : x \in X\}$. The resulting pair (X, Y) is the basis for the calculation of thresholds.

Let M be a metric set to be used as basis for the determination of an optimized, i.e., effective and efficient metric set with thresholds. A metric set is called effective if its classification error is close to 0, i.e., less than or equal to a threshold for the error ϵ . A metric set is called efficient if it is the smallest

set to do so. Therefore, we need to calculate a subset M_2 with thresholds $T_2 = \{t_{2_1}, \dots, t_{2_{d_2}}\}$ that yields classification error smaller than $\bar{\epsilon}$. To this aim, we determine thresholds based on the training set (X, Y) for all subsets of M . In other words, all sets that are element of the power set of M : 2^M . Then, for each subset M_2 the empirical classification error $\hat{\epsilon}_{X,Y}$ is calculated. The smallest set M_2 that has a classification error $\hat{\epsilon}_{X,Y} \leq \bar{\epsilon}$ is an effective and efficient subset of M . Algorithm 2 describes the whole threshold optimization algorithm in a step-wise fashion. We discuss the run time and scalability of the algorithm.

The value $\bar{\epsilon}$ can be used as a steering parameter, depending on the accuracy expected of the optimized set and the available data. The higher the accuracy shall be and the more data is available, the smaller $\bar{\epsilon}$ should be. It is possible that no M_2 , T_2 satisfies the condition that its error is below $\bar{\epsilon}$. In this case, there are three options to proceed: 1) choose a larger value for $\bar{\epsilon}$; 2) use a different metric set M ; 3) abort the optimization efforts and conclude that a metric set M_2 with threshold T_2 is insufficient to describe the classification. As a practical matter, $\bar{\epsilon}$ can be sampled, e.g., by starting with $\bar{\epsilon} = 0.01$ and increasing it in 0.01 steps till a metric set found.

4.3 OPTIMIZATION OF THE EFFICIENCY OF METRIC SETS WITH THRESHOLDS

Given an existing effective metric set, the threshold optimization algorithm can determine an effective and efficient subset. Let M be a metric set

Algorithm 2: The threshold optimization for metric set optimization

Input : Effective classification method f , set of software entities X , metric set M , error threshold δ

Output : Effective and efficient metric set M^* with thresholds T^*

$Y \leftarrow \{f(x) : x \in X\}$;

foreach $M^1 \in P(M)$ do

 Calculate thresholds T_1 for M^1 , X , Y with Algorithm 1

end

$M^* \leftarrow \min\{M^1 \in P(M) : \epsilon_{X,Y}(M^1, T^1) \leq \delta\}$;

return M^* and the corresponding T^* ;

with thresholds T and X a set of software entities. The function $f_0(x, M, T)$ defines a classification method for X . Then, f_0 , X , M , and an appropriate value for \tilde{a} are the input for the threshold optimization algorithm which will compute an optimized subset M^* with thresholds T^* .¹ As an example, the classification obtained by two metrics is approximated by only one of the two metrics. The dashed lines visualize the thresholds of the two metrics, used to classify the samples for the training.

4.4 REDUCTION OF THE CLASSIFICATION COMPLEXITY

Another way to utilize the threshold optimization algorithm is to reduce the complexity of the classification scheme. With thresholds, a simple kind of classification is described: if a threshold is violated, an entity is problematic. This makes it clear why an entity is problematic and also provides an indicator what the problem might be. A slightly more complex approach is to allow a number of infractions \tilde{e} , i.e., \tilde{e} thresholds may be violated. The function

$$f\lambda(z, M, T) = \begin{cases} 1 & \text{if } |\{i \in \{1, \dots, n\} : m_i(x) > t_i\}| \leq \lambda \\ 0 & \text{if } |\{i \in \{1, \dots, n\} : m_i(x) > t_i\}| > \lambda, \end{cases}$$

describes the classification defined by a metric set M with thresholds T . With $\tilde{e} = 0$, this $f_{\tilde{e}}$ is equal to f_0 .^{4.1}, hence $f_{\tilde{e}}$ is a generalization of f_0 .

One reason to use such a rule is to grant the developers more freedom, e.g., allowing short methods with a high structural complexity or long methods with a low structural complexity. But methods that are both long and structurally complex are forbidden. The classification with \tilde{e} allowed infractions introduces an additional complexity to understand *why* a problematic entity was classified as such and which counter measures can be taken. Complex approaches that may yield a very good classification may be difficult to impossible to interpret, e.g., SVM based techniques. Other techniques, e.g., classification trees show directly why an entity was classified as problematic, but it not clear how to fix as the tree may hide other reasons why the entity is problematic. In general, the classification could be performed by an arbitrary complex function f . A metric set that yields the same

classification, with no infractions whatsoever allowed is preferable, because as Occam's razor suggests, the simplest solution is preferable.

4.5 LEARNING OF ENVIRONMENT SPECIFIC THRESHOLDS

An important aspect of thresholds for metrics is that they are often dependent on the properties of the project environment such as the requirements, the developer qualification or the programming language. Therefore, the best results are achieved with thresholds tailored to the specific environment. In the previous two sections, we have only shown how the threshold optimization algorithm can optimize already existing classification methods. However, the algorithm is also able to determine thresholds where currently no method of classification exists. For this, an expert has to select a set of software entities X that are typical for the project environment. Afterwards, the expert manually classifies them into good and bad based on his or her expertise. As basis for this, the expert may use intricate knowledge, but also information about the software, e.g., the fault history to identify which sections are probably problematic. This is the traditional approach to determine the quality of software, without metric sets and thresholds. Using the thus obtained knowledge, we can determine a metric set with environment specific thresholds that mimics the expert's knowledge. To conform to our nomenclature, the expert can be seen as a function f that classifies software. Then, given a metric set M , the threshold optimization algorithm is able to determine an effective, efficient, and environment specific metric set M^* with thresholds T^* that emulates the expert's knowledge.

5. CASE STUDIES

For the validation of the approach for the optimization of metric sets, we performed four case studies consisting of a total of nine experiments. After we describe the general methodology used to perform the case studies, we present the results of the case studies. The case studies were designed to answer the following research questions:

R1: Is the method to optimize the efficiency of metric sets effective?

- R2: Is the method to reduce classification complexity effective?
- R3: Are the methods applicable and effective to different levels of abstraction (e.g., methods, classes, packages) and programming languages?
- R4: Is threshold recalculation with the rectangle learning algorithm necessary or is it sufficient to reuse known thresholds?
- R5: Is the exponential nature of the approach a threat to its scalability?

We answer these questions with respect to the case study results in Section 6.1.

5.1 METHODOLOGY

The case studies are based on metric data mined from archives of large scale open source software projects. By measuring code checked out from source code repositories, we obtained sets of software entities X with metric values $M(X)$. To guarantee the validity of the results, the measured data is randomly split into three disjunctive sets: a *training set* (X_{train}, Y_{train}) that contains 50% of the data; a *selection set* (X_{sel}, Y_{sel}) that contains 25% of the data; an *evaluation set* (X_{eval}, Y_{eval}) that contains 25% of the data. Each of the three sets is used at a different stage of our learning approach. The training set is used to calculate a set of hypotheses $h_{p,q}$ for sampled noise rates $\zeta_{0,p}, \zeta_{1,q}$ using the rectangle learning algorithm. The selection set is used to select the best of these hypotheses, i.e., an optimal hypothesis h^* with respect to the empirical classification error $\epsilon_{X_{cel}, Y_{cel}}$. The evaluation set is used to calculate the empirical classification error $\epsilon_{X_{eval}, Y_{eval}}$ of h^* on data that has not been part of the learning process. The error threshold \bar{a} for the threshold optimization algorithm is gradually increased in steps of 0.005 until a set is found that abides the threshold.

To further evaluate the case study results, we employ two additional measures for the quality of a hypothesis. The first is the MCC, a measure for the quality of binary classifications often used in machine learning (Matthews 1975). It is based on the so called confusion matrix. In the confusion matrix, a hypothesis is compared to the actual values

separately for positive and negative samples by counting *true positive* (tp), *true negative* (tn), *false positive* (fp), and *false negative* (fn) classified samples. In Fig. 4, the structure of the confusion matrix is visualized. The MCC is defined as

$$MCC = \frac{tp \cdot tn - fp \cdot fn}{\sqrt{(tp + fp)(tp + fn)(tn + fp)(tn + fn)}}$$

Its value is distributed between “-1 and 1, whereas 1 represents a perfect prediction, “-1 an inverse prediction, and 0 a random prediction. In contrast to the classification error, the MCC provides a balance between fp and fn predictions. Thus, MCC is more sensitive if the hypothesis has a bias towards rather falsely classifying positive samples than negative ones and vice versa.

		Actual classification	
		Positive	Negative
Hypothesis	Negative	true positive (tp)	false positive (fp)
	Positive	false negative (fn)	true negative (tn)

Fig. Confusion matrix

The second is the *F-Score*, another measure for the quality of hypothesis based on the confusion matrix. It is based on the *precision* and *recall* of a hypothesis. The precision measures how many of the positive predicted values of a hypothesis are actually positives. The recall is a measure for how many of the actual positive values are predicted correctly.

They are defined as follows:

$$\begin{aligned} precision &= \frac{tp}{tp + fp} \\ recall &= \frac{tp}{tp + fn} \end{aligned} \quad (5.2)$$

The F-score is then defined as harmonic mean between prediction and recall:

$$F - score = \frac{precision \cdot recall}{precision + recall} \quad (5.3)$$

Its value is distributed between 0 and 1, with 1 being a perfect score and zero being the worst.

5.2 CASE STUDY 1:

Optimization of Metric Sets for Methods : In the first case study, we analyzed the methodology for the optimization of metric sets for methods and functions. For this purpose, we measured software from various domains implemented in the languages C, C++ , and C#. Hereafter, we use the terms method and function interchangeably.

For C, we measured the *Apache HTTP Server*, an open source HTTP server for Unix/Linux and Windows systems developed and maintained by the Apache Foundation. We measured C++ methods for two of the main components of the *K Desktop Environment* (KDE) for Linux, the *kdebase* and the *kdelibs* components. The *kdebase* component contains most of the core applications of KDE, e.g., the window manager, an X terminal emulator, and the file manager Dolphin. The *kdelibs* provide a library of important core functions that are used by KDE, e.g., for networking, printing, and multi-threading. For C#, we measured three projects. The first C# project measured is *AspectDNG*,⁵ an aspect weaver that enables *Aspect Oriented Programming* (Kiczales et al. 2002) in C#. The second is the *NetTopologySuite*,⁶ a *Geographic Information System* (GIS) solution for the .NET platform. Finally, we measured *SharpDevelop*,⁷ an *Integrated Development Environment* (IDE) for C#, VB.NET, and Boo. Following table gives further information about the

analyzed versions, as well as the size of the projects.

5.3 CASE STUDY 2:

Optimization of Metric Sets for Classes:

In the second case study, we analyzed the optimization of metric sets for Java classes. The basis for this case study are two large-scale open source projects, both run by the Eclipse Foundation:⁸ the Eclipse Platform⁹ and the Eclipse *Java Development Tools* (JDT).¹⁰ The Eclipse Platform Project defines the main components of the Eclipse Platform, like the handling of resources, the workbench, and the editor framework. For the analysis, we excluded the test code and the *Standard Widget Toolkit* (SWT), a framework for user-interface programming. The rational being, that test code is inherently different from product code and thus test classes should not be compared to other classes. For example, test-cases can be highly repetitive as lists of values have to be compared to expected values, leading to a large size of test classes. On the other hand, the structure of test code should be less complex to prevent errors in the test code itself. The thresholds of the related metrics, like LOC and WMC should therefore be different than for normal code. As for the SWT, while it is formally a part of the Eclipse Platform Project, it is mainly independent. The Eclipse JDT implements an IDE for Java development on top of the Eclipse Platform. Again, we excluded the test code from the analysis.

Table-2
Statistical information about the measured projects

Project name	Version	Language	Number of methods	
			Total	Problematic
(a) Projects used for method-level analysis				
Apache Webserver	2.2.10	C	6718	1995
Kdebase	12/05/2008	C++	21404	4161
Kdelibs	12/05/2008	C++	37444	4921
AspectDNG	1.0.3	C#	2759	232
NetTopologieSuite	1.7.1.RC1	C#	3059	317
SharpDevelop	2.2.1.2648	C#	15700	1950
Project name	Version	Language	Number of classes	
			Total	Problematic
(b) Projects used for class-level analysis				
Eclipse java development tools	3.2	Java	4833	3349
Eclipse platform project	3.2	Java	5399	3517

The metric set under study was $M = \{WMC, CBO, RFC, NORM, LOC, NOM, NSM\}$ with thresholds as defined in Table 2b in the same manner as in case study 1. The metrics DIT and NOC were initially also part of this set, but we had to exclude them beforehand due to their poor distribution. As for DIT, ~98% of the classes had an inheritance depth of 0 or 1. With the metric NORM another inheritance related measure is still part of the metric set, thus DIT can be excluded

Table-3

Statistical information about the measured Java classes

Metric	Median	Arithmetic mean	Max value	Threshold
WMC	12	27.48	2138	100
CBO	8	13.40	212	5
RFC	20	35.21	675	100
NORM	0	0.96	166	3
LOC	24	82.95	6619	500
NOM	6	10.79	418	20
NSM	0	0.81	128	4

without reducing the internal attributes measured. The distribution of NORM is not ideal either, with only ~83% of all values greater than or equal to 2. However, this is still better than the distribution of DIT. The same argument is used to exclude NOC, where ~91% are 0 or 1.

The optimization yielded the set $M^* = \{CBO, NORM, NSM\}$, with thresholds $t_{CBO} = 5$, $t_{NOM} = 3$ and $t_{NSM} = 4$. Similar to case study 1, the calculated threshold values are the same as the ones used for the classification. The empirical error of this set is 0.27%. The MCC and F-score reveal no weaknesses either, both have values above 0.99. Therefore, by using the set M^* of size $|M^*| = 3$ instead of M of size $|M| = 7$, the size of the metric set is reduced by 57% without loss of generality. The following Table depicts detailed results and statistical information about the metrics.

5.4 CASE STUDY 3:

Reduction of the Classification Complexity for Methods : We performed this case study on the same data as case study 1 (see Section 5.2). The case study is designed to test the capability of the threshold optimization algorithm to reduce the classification complexity. To this aim, we calculated

the classification Y for the training using the metric set $M = \{VG, NBD, NFC, NST\}$, thresholds T as defined in Table 2a, and $f_1(\cdot, M, T)$ (see (4.4)) to calculate the classification. Thus, an entity is only considered problematic if the threshold of more than one metric is violated.

In contrast to case study 1, the result is different for the various languages. In case of C, the metric NST with a threshold of $t_{NST, C} = 5$ yields the best result with an empirical error of 0.84%. For C++ and C#, the metric VG with thresholds $t_{VG, C++} = 10$ and $t_{VG, C\#} = 9$ performs best with an empirical error of 0.87%, respectively 1.36%. The calculated threshold value in the C# experiment is different to the one used in the initial classification, while remains the same in the C and C++ experiments. The MCC revealed no weakness for the C and C++ experiments. However, in the C# experiment, the MCC dropped to 0.7598. While this is still a very good value, it indicates a possible weakness of this result. The F-score revealed no such weakness and was above 0.9930 for all three languages. Thus, we were able to use a simpler classification methodology, while also reducing the size of the metric set by 75% for all three languages. Table 6c summarizes the results of this case study.

5.5 CASE STUDY 4:

Reduction of the Classification Complexity for Classes : We performed the fourth case study on the same data as case study 2 (see Section 5.3). Like case study 3, it is designed to test the capability to reduce classification complexity. The methodology is similar to the one used in case study 3. Again, we use f_e instead of f_0 for the classification of software entities. Here, we use $\ddot{e} = 1, 2$, i.e., we perform two experiments:

- 1) one threshold violation allowed;
- 2) two threshold violations allowed.

Allowing more infractions would render the metric set ineffective, as more than half of the thresholds would have to be violated to even classify a class as problematic.

In both experiments, we determined effective and efficient metric sets. In the first experiment, with

one violation allowed, the metric set $M^* = \{RFC, NORM, NOM, NSM\}$ with thresholds $t_{RFC,1} = 98$, $t_{NORM,1} = 3$, $t_{NOM,1} = 20$, and $t_{NSM,1} = 4$ performed best with an empirical error of 1.71%. In the second experiment, the metric set $\{WMC, RFC\}$ with thresholds $t_{WMC,2} = 99$ and $t_{RFC,2} = 97$ was effective and efficient with a classification error of 2.21%. Half of the threshold values calculated in this case study deviated from the ones used for the classification. While the empirical error of the experiment with $\ddot{e} = 1$ was higher than with $\ddot{e} = 2$, the MCC performed the other way around. While the MCC of the experiment with $\ddot{e} = 1$ is unproblematic with 0.9449, it drops slightly for $\ddot{e} = 2$ to 0.8494. This suggest, that the hypothesis in the second experiment has a slight bias towards positive samples, as the F-score revealed no such weakness. It is above 0.98 for both experiments. The results show that a simpler classification can be used in both cases and, furthermore, the metric set sizes can be reduced by 42% and 71%, respectively. Table 6d summarizes the results of this case study.

6. DISCUSSION

In this section, we discuss the research questions R1–R5 with respect to the case study results. Afterwards, we discuss other methods for metric set optimization and compare them to our methodology.

6.1 DISCUSSION OF RESEARCH QUESTIONS

R1: Is the method to optimize the efficiency of metric sets effective?

The results of the three experiments of case study 1 and the experiment performed in case study 2 show that the methodology is capable of decreasing the size of metric sets between 57% and 75% without a significant loss of classification precision. Based on these four successful experiments, each of them performed in a different environment, the answer to this research question is yes.

R2: Is the method to reduce classification complexity effective?

In case studies 3 and 4 we classified the data with a method more complex than the simple threshold classification. A total of five experiments were performed, in all of which simple thresholds were sufficient to reproduce the original classification. Furthermore, the resulting metric sets were also 42% to 85% smaller than the ones used for the classification. Thus, the answer to this research question is yes.

R3: Are the methods applicable and effective to different levels of abstraction (e.g., methods, classes, packages) and programming languages?

In the case studies 1 and 3, we analyzed methods and functions, while classes were under consideration in case studies 2 and 4. Thus, the approach does not depend on the level of abstraction. Furthermore, in the case studies, we used projects written in four different programming languages: C, C++, C# and Java. These four languages cover the procedural and the object-oriented paradigm. Moreover, C is a low-level and close to the system programming language, whereas Java and C# are relatively high level. Therefore, the results indicate that the programming language has no impact on the capabilities of the methodology and the answer to this question is yes.

R4: Is threshold recalculation with the rectangle learning algorithm necessary or is it sufficient to reuse known thresholds?

On one hand, the results of case studies 1 and 2 suggest that recalculation of threshold values is not required when optimizing a metric set. In all experiments conducted, the calculated threshold values were the same as the original ones. On the other hand, the results of case study 3 and 4 suggest that when the classification method is changed, recalculation of threshold values is beneficial even if the formerly used method is based on thresholds. In addition to the problems analyzed in the case studies, there are possible applications where no thresholds are available, e.g., if a non-threshold based classification method is to be optimized. In such cases threshold calculation is integral and may not be omitted. In conclusion, whether the recalculation of

threshold adds value to the proposed method depends on the application of the method.

R5: Is the exponential nature of the approach a threat to its scalability?

The execution of all nine experiments performed as part of the four case studies took 139 seconds in total on a normal desktop workstation running on an Intel Core2 Duo E8400 processor. For these experiments, the rectangle learning algorithm was executed a total of 480 times, therefore, a single execution took about 0.29 seconds in average. As there are 2^{20} different subsets of a metric set of size 20, the execution would take $2^{20} \cdot 0.29 \text{ H} \approx 304.000$ seconds, thus, approximately 3.5 days. While this is a pretty long time, it has to be taken into account that such an optimization must only be performed once and does not need to run regularly. Furthermore, run time can be reduced by using multiple parallel threads of execution. Of course, with even greater metric sets, this does not resolve the problem. In conclusion, it can be said that the approach is able to handle metric sets with a size of about 20 in an acceptable amount of time. For larger metric sets, a heuristic for the selection of subsets to be analyzed needs to be employed.

6.2 LIMITATIONS

We only analyzed open source software in the case studies; non-open-source software has not been analyzed. However, the work by Werner showed that a similar approach worked with TTCN-3 test suites, i.e., software written in a *Domain Specific Language* (DSL) in a non-open-source environment.

The metric sets we analyzed only consist of internal product metrics on the method and class level. Metric sets on higher levels of abstraction, as well as metric sets including process or resource metrics have not been analyzed. Furthermore, the chosen threshold values may have been inadequate to begin with, leading to misclassified training data.

The proposed methodology produces a binary classification and can therefore only differentiate between “good” and “bad”, further shades of grey are not possible.

7. RELATED WORK

Research on how environment specific metric sets can be obtained was performed by Basili and Selby. In contrast to this work, the authors use a *Goal/Question/Metric* (GQM) approach to determine a metric set and condense it using factor analysis. A statistical method to obtain threshold values was introduced by French who used it to derive thresholds for object-oriented and procedural software.

An approach to determine classification trees to identify quality critical modules was proposed by Porter and Selby. The tree makes its decisions based on intervals of metric values, which is similar to using thresholds.

A methodology to determine metric sets to predict quality critical modules using Boolean Discriminant Functions (BDFs) has been introduced by Schneidewind. The BDFs consist of boolean disjunctions of threshold violations to identify critical modules, which is just another formalization of the classification model used in this work. They determine the thresholds using Kolmogorov–Smirnov tests. This model is extended to Generalized BDFs by introducing conjunctions into the Boolean functions.

Use environment specific thresholds to determine whether metric values are low, average, or high, based on the arithmetic mean and the standard deviation of observed metric data. These thresholds are then used in an *overview pyramid* to provide an overview of object-oriented software based. The metrics are divided into three aspects: inheritance; size and complexity; coupling. Using the thresholds, a coloring scheme is defined that visualizes the software properties. In comparison to this work, the authors do not assume thresholds to define metric values as problematic, but rather use them to discriminate metric values into low, average, and high values.

An instantiation of the maintainability characteristic of the ISO 9126 quality model (ISO/IEC 2001–2004) is described by Heitlager. They use both internal and external product metrics to define ratings for the source code properties

volume, complexity per unit, duplication, unit size, and unit testing. Based on the property ratings, the sub-characteristics of maintainability are rated from which maintainability is inferred. The ratings are based on intervals, which are similar to using thresholds. In comparison to our work, they have five rating classes instead of a binary classification. Furthermore, very good ratings for one property allow bad ratings for another, which is different to the strict threshold classification we apply.

A paper similar to this work, but using a less sophisticated approach for the optimization of metric sets for TTCN-3 is presented by Werner. However, the machine learning methodology used in this work is more mature and the case studies analyze it in a wider setting, i.e. various programming languages and levels of abstraction.

8. CONCLUSION

We defined a novel high-level approach for the calculation of thresholds for software metrics to evaluate quality attributes. The method is purely data driven and utilizes machine learning techniques. Based on this, we defined a methodology to determine optimized metric sets that replicate a given

classification of a quality attribute. We outlined how the methodology can be applied to improve the efficiency of existing metric sets with thresholds, reduce the complexity of a used classifier and how a new metric set can be introduced using the methodology. In two case studies, we showed that the methodology is able to greatly improve the efficiency of existing metric sets. In two further case studies, we reproduced complex classifications successfully with simple thresholds.

Future projects may include more case studies, on how well the approach works in other environments, e.g., domain specific languages or how well it handles sparse data. Moreover, it may be investigated how learning of *Disjunctive Normal Forms* (DNFs) of thresholds instead of conjunctions affects the hypothesis quality, the metric set reduction, and the interpretability of the resulting classifiers. Furthermore, a detailed comparison with black-box classification techniques like *Artificial Neural Networks* (ANNs) or SVMs is an interesting topic for the future. Another research direction is to determine metric sets and thresholds that can be used to steer software project decisions. To this aim, the approach needs to be adapted for process data.

REFERENCES

1. Angluin D, Laird P (1988) Learning from noisy examples. *Mach Learn* 2(4):343–370. doi:10.1023/A:1022873112823
2. Basili V, Rombach H (1988) The TAME project: towards improvement-oriented software environments. *IEEE Trans Softw Eng* 14(6):758–773
3. Basili V, Weiss D (1984) A methodology for collecting valid software engineering data. *IEEE Trans Softw Eng* 10(6):728–738
4. Basili VR, Selby RW Jr (1985) Calculation and use of an environment's characteristic software metric set. In: *ICSE '85: proceedings of the 8th international conference on Software engineering*. IEEE Computer Society Press, Los Alamitos, CA, USA, pp 386–391
5. Basili VR, Briand LC, Melo WL (1996) A validation of object-oriented design metrics as quality indicators. *IEEE Trans Softw Eng* 22(10):751–761. doi:10.1109/32.544352
6. Benlarbi S, Emam KE, Goel N, Rai S (2000) Thresholds for object-oriented measures. In: *ISSRE '00: proceedings of the 11th international symposium on software reliability engineering*. IEEE Computer Society, Washington, DC, USA, p 24
7. Brodag T, Herbold S, Waack S (2010) A generalized model of pac learning and its applicability. *Mach Learn* (manuscript in revision)
8. Chidamber SR, Kemerer CF (1994) A metrics suite for object oriented design. *IEEE Trans Softw Eng* 20(6):476–493. doi:10.1109/32.295895
9. Daly J, Brooks A, Miller J, Roper M, Wood M (1996) Evaluating inheritance depth on the maintainability of object-oriented software. *Empir Softw Eng* 1(2):109–132
10. Devroye L, Györfi L, Lugosi G (1997) A probabilistic theory of pattern recognition. Springer, New York
11. Duda R, Hart P (1973) *Pattern classification and scene analysis* ETSI (2007) ETSI Standard (ES) 201 873-1 V3.2.1 (2007-02): the testing and test control notation version 3; part 1: TTCN-3 core language. European Telecommunications Standards Institute (ETSI), Sophia-Antipolis, France, also published as ITU-T Recommendation Z.140

13. Fenton N, Pfleeger S (1997) *Software metrics: a rigorous and practical approach*. PWS Publishing Co. Boston, MA, USA.
14. French V (1999) *Establishing software metric thresholds*. In: *International workshop on software measurement (IWSM99)*
15. Grabowski J, Hogrefe D, Réthy G, Schieferdecker I, Wiles A, Willcock C (2003) *An introduction to the testing and test control notation (ttcn-3)*. *Comput Netw* 42(3):375–403. doi:10.1016/S1389-1286(03)00249-4
16. Heitlager I, Kuipers T, Visser J (2007) *A practical model for measuring maintainability*. In: *6th international Conference on the Quality of information and communications technology, 2007. QUATIC 2007*, pp 30–39. doi:10.1109/QUATIC.2007.8
17. IEEE (1990) *Ieee glossary of software engineering terminology. ieee standard 610.12*. Tech. rep., IEEE
17. ISO/IEC (2001–2004) *ISO/IEC standard no. 9126: software engineering—product quality; parts 1–4*. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), Geneva, Switzerland
17. ISO/IEC (2005) *ISO/IEC Standard No. 9000*. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), Geneva, Switzerland
18. Kearns M (1998) *Efficient noise-tolerant learning from statistical queries*. *J ACM* 45(6):983–1006. doi:10.1145/293347.293351
19. Khoshgoftaar TM (2002) *Improving usefulness of software quality classification models based on boolean discriminant functions*. In: *ISSRE '02: proceedings of the 13th international symposium on software reliability engineering*. IEEE Computer Society, Washington, DC, USA, p 221
20. Kiczales G, Lamping J, Lopes C, Hugunin J, Hilsdale E, Boyapati C (2002) *Aspect-oriented programming*. US Patent 6,467,086
21. Lanza M, Marinescu R, Ducasse S (2005) *Object-oriented metrics in practice*. Springer-Verlag New York, Inc., Secaucus, NJ, USA
22. Lilliefors HW (1967) *On the Kolmogorov–Smirnov test for normality with mean and variance unknown*. *J Am Stat Assoc* 62(318):399–402. <http://www.jstor.org/stable/2283970>
23. Lorenz M, Kidd J (1994) *Object-oriented software metrics: a practical guide*. Prentice Hall PTR
24. MacKay DJ (2003) *Information theory, inference, and learning algorithms*. Cambridge University Press
25. Mammen E, Tsybakov AB (1999) *Smooth discrimination analysis*. *Ann Stat* 27(6):1808–1829
26. Matthews BW (1975) *Comparison of the predicted and observed secondary structure of t4 phage lysozyme*. *Biochim Biophys Acta, Protein Struct* 405(2):442–451. doi:10.1016/0005-2795(75)90109-9. URL:<http://www.sciencedirect.com/science/article/B73GJ-47T22GD-132/2/b5b0dbd824d44e6edeefb7b8d2613775>
27. Nagappan N, Ball T, Zeller A (2006) *Mining metrics to predict component failures*. In: *ICSE '06: proceedings of the 28th international conference on software engineering*. ACM, New York, NY, USA, pp 452–461. doi:10.1145/1134285.1134349
28. Porter AA, Selby RW (1990) *Empirically guided software development using metric-based classification trees*. *IEEE Softw* 7(2):46–54. doi:10.1109/52.50773
29. Quinlan JR (1986) *Induction of decision trees*. *Mach Learn* 1:81–106. doi:10.1007/BF00116251
30. Rosqvist T, Koskela M, Harju H (2003) *Software quality evaluation based on expert judgement*. *Softw Qual J* 11:39–55. doi:10.1023/A:1023741528816
31. Schneidewind NF (1997) *Software metrics model for integrating quality control and prediction*. In: *ISSRE '97: proceedings of the eighth international symposium on software reliability engineering*. IEEE Computer Society, Washington, DC, USA, p 402
32. Schneidewind NF (2000) *Software quality control and prediction model for maintenance*. *Ann Softw Eng* 9(1–4):79–101. doi:10.1023/A:1018920623712
33. Schölkopf B, Smola AJ (2002) *Learning with kernels*. MIT Press
34. Selby RW, Porter AA, Schmidt DC, Berney J (1991) *Metric-driven analysis and feedback systems for enabling empirically guided software development*. In: *ICSE '91: proceedings of the 13th international conference on software engineering*. IEEE Computer Society Press, Los Alamitos, CA, USA, pp 288–298
35. Shawe-Taylor J, Cristianini N (2004) *Kernel methods for pattern analysis*. Cambridge University Press
36. Tsybakov AB (2004) *Optimal aggregation of classifiers in statistical learning*. *Ann Stat* 32(1):135–166
37. Werner E, Grabowski J, Neukirchen H, Rottger N, Waack S, Zeiss B (2007) *TTCN-3 quality engineering: using learning techniques to evaluate metric sets*. *Lect Notes Comput Sci* 4745:54