Do We Need Improved Code Quality Metrics?

Tushar Sharma
Siemens Technology
Charlotte, USA
tusharsharma@ieee.org

Diomidis Spinellis
Athens University of Economics and Business
Athens, Greece
dds@aueb.gr

arXiv:2012.12324v1 [cs.SE] 22 Dec 2020

***Abstract*—The software development community has been using code quality metrics for the last five decades. Despite their wide adoption, code quality metrics have attracted a fair share of criticism. In this paper, first, we carry out a qualitative exploration by surveying software developers to gauge their opinions about current practices and potential gaps with the present set of metrics. We identify deficiencies including lack of *soundness*, *i.e.,* the ability of a metric to capture a notion accurately as promised by the metric, lack of support for assessing software architecture quality, and insufficient support for assessing software testing and infrastructure. In the second part of the paper, we focus on one specific code quality metric—LCOM as a case study to explore opportunities towards improved metrics. We evaluate existing LCOM algorithms qualitatively and quantitatively to observe how closely they represent the concept of cohesion. In this pursuit, we first create eight diverse cases that any LCOM algorithm must cover and obtain their cohesion levels by a set of experienced developers and consider them as a ground truth. We show that the present set of LCOM algorithms do poorly *w.r.t.* these cases. To bridge the identified gap, we propose a new approach to compute LCOM and evaluate the new approach with the ground truth. We also show, using a quantitative analysis using more than 90 thousand types belonging to 261 high-quality Java repositories, the present set of methods paint a very inaccurate and misleading picture of class cohesion. We conclude that the current code quality metrics in use suffer from various deficiencies, presenting ample opportunities for the research community to address the gaps.**

***Index Terms*—metrics, code quality, software metrics**

I. Background

A software metric quantifies a characteristic or an attribute of a software product, process, or project [1]–[3] that could be used later for assessment or prediction [1]. *Software product* metrics measure aspects such as size, complexity, and performance. *Process metrics* capture aspects of software development processes, such as the cost and effort required during the testing phase. Finally, *project metrics* concern project-level management and resource allocation, such as the number of active developers or the schedule overrun. In this paper, we focus on code quality metrics, a sub-category of product metrics. We take a critical look at the existing metrics and explore whether the software development and research community require an improved set of metrics.

Code quality metrics have a long history, which can be divided into three eras of research and practice. In the first era, researchers mainly concentrated on size and complexity metrics, such as the lines of code (LOC), cyclomatic complexity [4], and the Halstead metrics [5]. The second era brought a plethora of interest in object-oriented metrics. The most notable contributions of this period are the Chidamber and Kemerer (C&K) metrics suite [6], the MOOD metrics [7], and the design metrics by Martin [8]. We consider the year 2000 as the beginning of the third era of code quality metrics. Research then changed focus from proposing new metrics to utilizing code quality metrics in applications regarding bug prediction [9], maintainability prediction [10], and code smell detection [11]. Also, the era witnessed a wider adoption of metrics among practitioners to track their code quality.

Commonly used code quality metrics could be classified in the following categories—*size*, *complexity*, *coupling*, *cohesion*, and *inheritance* metrics. Size metrics compute the size of program entities or constructs. Particularly, LOC was one of the first metrics that the software developers started using in the late 1960s.

Halstead [5] proposed a metrics suite to measure program size (termed as volume and vocabulary), based on rather primitive operator and operand counts. Function points [12] compute size and complexity of a program considering number of user inputs, outputs, inquiries, and external interactions. Other commonly used size metrics are number of methods (NOM) in a class, number of parameters (NOP) in a method, and number of classes (NC) in an assembly, package, or a software system.

Cyclomatic complexity [4] is the most widely used complexity metric. McCabe computed the complexity using $v(G) = e - n + 2$ where e and n refer to number of edges and nodes in a control flow graph. C&K metric suite [6] introduced weighted methods per class (WMC) metric, which calculates the complexity of a class by summing up the cyclomatic complexities of each method of the class. Chidamber *et al.* [6] also proposed coupling between objects (CBO) metric, which indicates the number of classes coupled with a given class. The notion was further expanded to fan-in and fan-out metrics proposed by Henry and Kafura [13]. Fan-in and fan-out represent incoming and outgoing dependencies of a class respectively. Similarly, Martin [8] proposed afferent and efferent coupling metrics to quantify incoming and outgoing dependencies for a module. The Lack of Cohesion in Methods (LCOM) metric, proposed in C&K metrics suite, captures the cohesion characteristic among a class's methods based on their access to data members. The category of inheritance metrics include depth of inheritance (DIT) and number of children (NOC) [6]. Coleman *et al.* [14] proposed Maintainability Index

(MI) by combining Halstead volume, cyclomatic complexity, lines of code, and percentage of comments.

II. Deficiencies in Code Quality Metrics

Software metrics have always been on a roller-coaster ride. On one hand, researchers and practitioners have adopted metrics not only to reveal quality characteristics of their programs, but also to combine existing metrics into new ones, and use these to study more complex phenomena. On the other hand, metrics have drawn wide criticism. The majority of this criticism is related to *completeness* and *soundness*.

The first case covers the extent of completeness of the implementation details provided to compute the metrics independent from the programming language. For instance, implementation details of metrics in C&K suite [6] are missing, leaving it up to one's interpretation [2], [15]. Two examples of such deficiencies involve the lack of concrete details for the implementation of the *lack of cohesion in methods* (LCOM) metric [15], and the incomplete definition of the *coupling between objects* (CBO) metric. Particularly, the definition of the CBO metric does not clarify whether both incoming and outgoing dependencies or only outgoing dependencies should be used in the calculation.

Soundness of a metric refers to the ability of the metric to accurately capture the notion underlying its theoretical basis. Shepperd and Ince [16] presented their critical view on the Halstead, cyclomatic complexity, and information flow metrics. They wrote that the computing rules followed in the Halstead metrics are arbitrary and make use of magic numbers without providing sufficient theoretical foundation. For instance, programming time T is computed through the formula $T = E/(f \times S)$, where E is referred to as programming effort, f as seconds-to-minute ratio = 60, and S as Stround number $= 18$. Furthermore, these metrics use tokens and number of operators and operands, which are considered too simplistic and primitive, because they fail to capture control, data, and module structure [16]. Finally, Shepperd and Ince pointed out that Halstead metrics were developed in the era of batch processing where software systems were of a considerably smaller scale than today—often amounting to a few hundred lines of code. Similarly, other metrics such as cyclomatic complexity [15]–[17] and maintainability index [18]–[20] are also criticized for poor theoretical basis and validation.

A sound metric computing mechanism not only depends on the implementation, but also on the crisp definition of the metric [1]. From the implementation perspective, different metric tools often produce different results for the same source code [21]. More importantly, a metric is deprived from its usefulness if it does not measure what it intends to, especially in commonly occurring cases. SonarQube,[1], a widely used platform to measure and track code quality metrics and technical debt, identified that cyclomatic complexity fails to represent "complexity" of a snippet correctly (for instance, in the presence of switch-case statements). The platform has recently introduced a new metric, *i.e., cognitive complexity*, as an alternative to the cyclomatic complexity metric [22].

Similarly, despite the availability of many variants (LCOM1 [23], LCOM2 [6], LCOM3 [10], LCOM4 [24], LCOM5 [25]), the lack of cohesion in methods (LCOM) metric is another example where the intended characteristic, *i.e.,* cohesion [26], is not always correctly captured. In RQ2 of this study (Section V), we compare the five variants of the LCOM metric using different cases, and reveal their deficiencies in detail.

Apart from the issues of the existing metrics reported above, the following concerns are also identified regarding the commonly used code quality metrics. First, the present set of code quality metrics is mainly designed for object-oriented programs. Nowadays the software development community extensively uses programming languages, such as Python and JavaScript, that are not strictly object-oriented. Hence, the present set of metrics is not applicable in its original form on such languages. Second, the focus of code quality metrics has been limited to methods and classes. There is hardly any metric at the level of architecture granularity. For instance, akin to a complex method (inferred by using cyclomatic complexity metric), it is difficult to comment anything about the complexity of a component or the entire software system. Finally, recently, sub-domains of software, including tests and configuration code have become inseparable parts of the production code. The present set of metrics hardly support these sub-domains. Metrics such as code coverage for configuration and database code, as well as coupling between the test or infrastructure code and the production code could provide useful insight for the development community.

III. Goal and Methodology

The goal of this study is to first, understand the current practices in the software development community related to code quality metrics and identify potential gaps. Also, we aim to take one case study of a commonly used metric and explore the soundness of the metric.

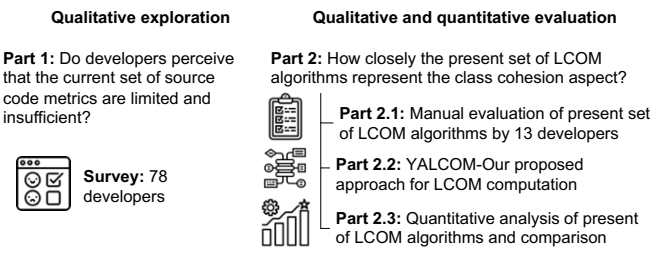

Fig. 1. Overview of the method

Towards the goals, we divide the study, as outlined in Figure 1, broadly into two parts. In the first part, we carry out a qualitative exploration by surveying software developers. The aim of the exploration is to gauge developers' opinions about current practices and potential gaps concerning the present set of metrics. For the case study, we considered cyclomatic complexity, Halstead metrics, and LCOM metrics. Halstead

[1]https://www.sonarqube.org/

metrics are not widely used; hence we dropped them from our consideration. Cyclomatic complexity has been studied and criticized by many authors; hence, the deficiencies of the metrics are well-known. The LCOM metric has been also studied widely; however, their deficiencies have not been explored and evaluated by the community. Therefore, we selected the metric for our case study. In the second part we evaluate, qualitatively and quantitatively, existing LCOM algorithms to observe their soundness *i.e.,* how closely they represent the concept of cohesion. In this pursuit, we first create eight diverse cases that any LCOM algorithm must be tested against. We first present these cases to experienced developers to obtain and establish the expected values of cohesion in each of these diverse cases. We compare the exiting set of LCOM algorithms and identify that the present set of LCOM algorithms do poorly *w.r.t.* these cases. To bridge the identified gap, we propose our approach of computing LCOM that we refer to as YALCOM (Yet Another Lack of Cohesion in Methods). Finally, we download 522 Java repositories, analyze them to compute the LCOM metric for all the classes in the analyzed repositories with existing methods as well as YALCOM. We compare the outcome of all the considered algorithms by computing *Euclidean distance* among them and deduce our observations based on the results of the experiments.

A. Research Questions

We explore two research questions in this study.

RQ1. *To what extent developers perceive the current set of code quality metrics comprehensive and sufficient?*

The first goal of this study is to explore the perception of software developers about code quality metrics. Specifically, we aim to gauge their opinions about sufficiency and completeness of the present set of commonly used metrics.

RQ2. *To what extent the existing set of* LCOM *algorithms capture the class cohesion aspect?*

As a case study, we focus on a specific commonly used metric—LCOM and explore to what extent the existing LCOM algorithms capture the class cohesion aspect in the produced metric values. We also aim to compare the existing algorithms with our method that we present in this study.

IV. RQ1. To what extent developers perceive the current set of code quality metrics comprehensive and sufficient?

We discuss the mechanism, experiment, and the obtained results corresponding to the first research question addressed in this study.

A. Approach

We designed a questionnaire to be used in an online survey to gather the opinions of software developers and researchers. The goal of the survey was to understand the practices followed by software developers and identify potential gaps related to code quality metrics.

The first question asked the participants their software development experience in the number of years. The second question (multiple-choice) attempted to gather the commonly used metrics by the participants. The next set of questions asked the participants whether the current set of metrics provides sufficient insights about specific software engineering concerns (*i.e.,* design, architecture, test, and infrastructure). We offered five-point Likert scale-based options (*i.e.,* strongly agree, agree, neither agree nor disagree, disagree, and strongly disagree) and asked them to choose one of the options. We also asked additional code quality metrics that they would like to have (open text). The next two questions were targeted towards the extent to which the promised aspect is represented and implemented correctly by the metrics (five-point Likert scale). Finally, we asked the participants about their suggestions and feedback (open text).

We initially ran an internal pilot study, which helped us refine the survey questionnaire according to the received suggestions, and then publicly distributed the final survey questionnaire through various social media platforms. We kept the survey anonymous. Table I presents the final set of questions in the survey.

B. Results

We received 78 complete responses. The participants belonged to various experience groups—no experience (5%), 1–2 years (12%), 3–5 years (17%), 6–10 years (21%), 10–20 years (32%), and more than 20 years of experience (13%).

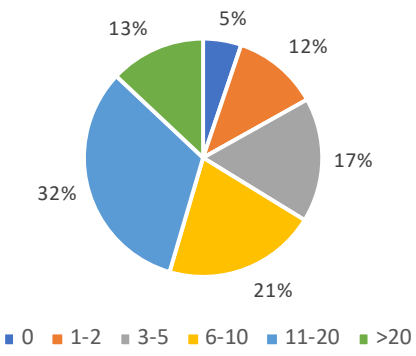

Fig. 2. Software developer experience of the survey participants

In the second question, respondents were asked to choose from a list of commonly used code quality metrics, which ones they use and monitor on a regular basis. As expected, lines of code (LOC) and cyclomatic complexity (CC) are the most commonly used metrics; 54% and 52% of the participants selected them, respectively. Almost none of the participants use Halstead metrics. Apart from the provided options, participants further mentioned that they rely on test coverage and clone percentage. Still, 26% of the participants affirmed that they do not use any metrics.

Next, respondents were asked whether the current set of metrics provides sufficient insight regarding software design. The majority (46% agree and 6% strongly agree) of the participants acknowledged that they get insight about software design using the current set of metrics. We asked a similar

Q#	Questions
1.	Please mention your software development experience in years.
2.	Which code quality metrics [1] do you use/monitor for your source code. Check all that apply.
	Lines of Code (LOC), Halstead volume, Average metrics (such as average method size and average number of methods per class), Cyclomatic complexity (CC), Weighted methods per class (WMC), Halstead difficulty, Halstead programming effort, Lack of cohesion in methods (LCOM), Coupling between objects (CBO), Fan-in/Fan-out, Afferent and efferent coupling (Ca/Ce), Depth of inheritance tree (DIT), Number of children (NOC), Response for a class, Comment percent, Others (specify), I/We do not use metrics,
	1. Stephen H. Kan. 2002. Metrics and Models in Software Quality Engineering (2nd ed.). Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
3.	Do existing metrics provide the insights you need regarding the software's design?
4.	Do existing metrics provide the insights you need regarding the software's architecture?
5.	Do existing metrics provide the insights you need regarding the testing sub-domain of software development?
6.	Do existing metrics provide the insights you need regarding the infrastructure (operations, site reliability engineering, production engineering) sub-domain of software development?
7.	Which additional code quality metrics would you like to have?
8.	A metric's accuracy represents the extent to which the metric captures and represents the promised aspect. How often do you observe that some of the metrics are inaccurate in certain cases? For instance, value of a metric is reported 0 while you expected 0.5 due to the wrong implementation in the tool that you used.
9.	A metric's accuracy represents the extent to which the metric captures and represents the promised aspect. How often do you observe that some code quality metrics are inaccurate in certain cases? For instance, value of a metric is reported 0 while you expected 0.5 due to the specification of the algorithm used to calculate the metric (rather than a software implementation fault).
10.	Do you have any comments, suggestions, reservations, feelings, or objections regarding the code quality metrics used commonly? Feedback to improve the survey is also welcome. Kindly provide your email address if you would like to receive compiled results of the survey (optional).

TABLE I
QUESTIONS IN THE DEVELOPERS' SURVEY

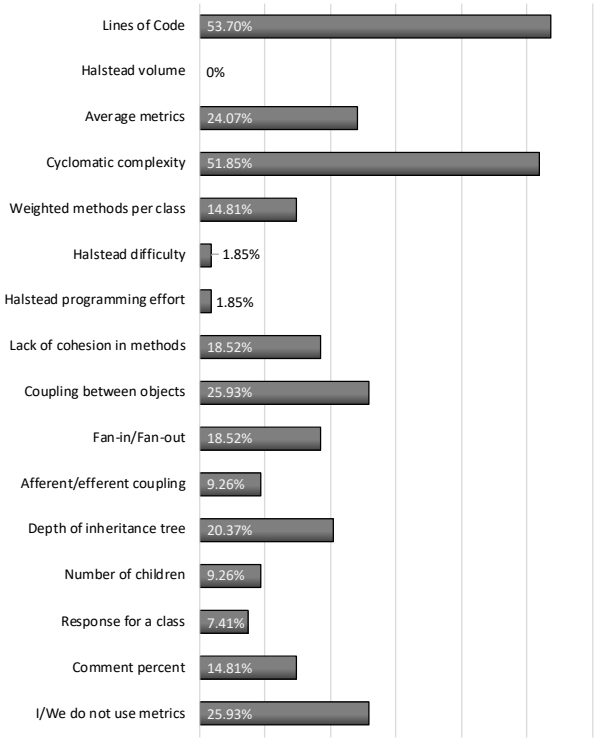

Fig. 3. Commonly used metrics by the survey participants

question about software architecture, and the majority (31% disagree and 15% strongly disagree) of the participants stated that the current set of code quality metrics is inadequate to measure quality of software architecture. In the same context, participants claimed that it is not easy to quantify architectural aspects. As a respondent stated, *"you cannot easily distill architectural insight into numbers"*. Other participants mentioned that the current set of metrics is only suited to measure implementation and design aspects, but not architectural characteristics.

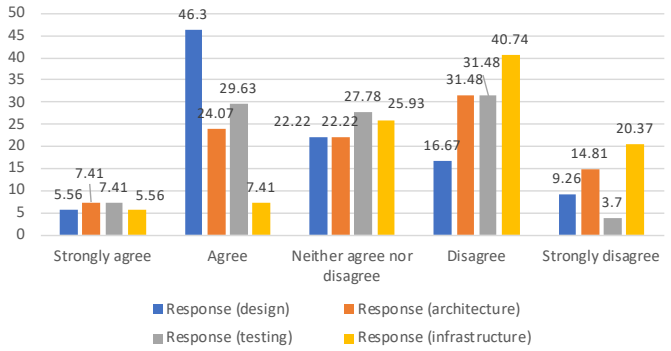

Fig. 4. Survey responses toward assessing sufficiency of the commonly used metrics

The next question inquired whether code quality metrics provide sufficient insight for the sub-domains of software development testing and infrastructure. Regarding the testing sub-domain, participants seemed divided. The majority (31% disagree and 4% strongly disagree) of the participants stated that code quality metrics fall short with respect to testing. In the words of a respondent, *"test coverage is not enough"*. On the other hand, another group of participants (30% agree and 7% strongly agree) claimed that the current set of testing-related metrics is sufficient. The negative opinion was amplified for the infrastructure sub-domain. A clear consensus emerged among 61% (41% disagree and 20% strongly disagree) of the participants who supported that code quality metrics are insufficient to provide insight regarding infrastructure. Undoubtedly, a lot of progress is expected toward measuring infrastructure code quality.

We got interesting and insightful answers in the response of the question that asked for additional metrics that they would like to use. The desired set of metrics mentioned, include architecture metrics including specialized metrics for new architecture styles (e.g. microservices), more context-sensitive metrics, 'should I refactor this [method/class]' metric, code change frequency, accidental complexity metrics, number of tests to cover all paths for different scopes, the degree of adherence to a paradigm, and metrics to show the extensibility of a module. This set of desirable metrics coming directly from the practitioners may lead new efforts towards novel code quality metrics.

A metric's accuracy represents the extent to which the metric captures and represents the promised aspect. Two additional questions were included regarding metrics showing inaccurate values due to wrong implementation and specification of the applied algorithm. 32% of the participants agreed that they often see inaccurate metric values due to wrong implementation of the metric tools. Another 39% of the participants did not notice any incorrect metric values. Concerning algorithmic inaccuracies related to wrong specification, half of the participants expressed their ignorance (by choosing *'I don't know'* option).

We summarize the results of the exploration below.

- The participants use a variety of code quality metrics, with LOC and CC being the most prevalent.
- The participants agreed that the present set of metrics provides insight for software design; however, they recognized that the currently used metrics are insufficient to measure architecture quality or to assess the testing and infrastructure sub-domains.
- The participants would like to see new metrics covering architectural aspects, code churn, and module extensibility.
- Finally, participants indicated that they often find inaccurate metric values due to wrong implementation.

V. RQ2. To what extent the existing set of LCOM algorithms capture the class cohesion aspect?

This research question takes a concrete case of LCOM metric and attempts to gauge the soundness of the existing set of methods to compute the metric.

A. RQ2.1: Manual Evaluation of LCOM Algorithms

1) Approach: In order to establish whether the commonly used set of LCOM metrics sufficiently capture the cohesion aspect of abstractions, we handcrafted eight classes/interfaces representing different cases. They are designed to cover various common cases involving interplay of method calls, fields—their type (a class or an interface) and their accesses, and inheritance that impact class cohesion and may potentially reveal the deficiencies of the existing algorithms to compute LCOM. Table II summarizes these cases; their corresponding source code can be found in our online replication package.[2]

Case	Class member relationships	Description
Case1	$m_1 \rightarrow \{a_1, a_2, a_3\}$, $m_2 \rightarrow \{a_2\}$, $m_3 \rightarrow \{a_3\}$	
Case2	$m_1 \rightarrow \{a_1, a_2, a_3\}$, $m_2 \rightarrow \{a_2\}$, $m_3 \rightarrow \{a_3\}$	a_3 is static
Case3	$m_1 \rightarrow \{a_1, a_2\}$, $m_2 \rightarrow \{a_1, a_2\}$, $m_3 \rightarrow \{a_3\}$	
Case4	$m_1 \rightarrow \{a_1, a_2, a_3\}$, $m_2 \rightarrow \{a_1\}$, $m_3 \rightarrow \{a_2, a_3\}$	a_1 is in super class as a protected member
Case5	$m_1 \rightarrow \{a_1\}$, $m_2 \rightarrow \{a_2\}$, $m_3 \rightarrow \{a_3\}$	
Case6	$m_1 \rightarrow \{a_1, a_2, a_3\}$, $m_2 \rightarrow \{a_2\}$, $m_2 \Rightarrow m_3$	
Case7	m_1, m_2, m_3	Type is an interface
Case8	$m_1, m_2, m_3, m_2 \Rightarrow m_3$	There are no fields (utility class)

TABLE II
CASES COVERING VARIOUS SCENARIO CONCERNING CLASS COHESION; M_i REFERS TO A METHOD, A_j REFERS TO AN ATTRIBUTE (FIELD) OF A CLASS, AND RELATIONS '$\rightarrow$' AND '$\Rightarrow$' REFER TO *access* AND *method invocation*, RESPECTIVELY.

To establish a ground truth, we put all of these cases in the form of a survey and sent out the survey to experienced software developers in our network. The survey had nine questions, one question per case showing the source code and one question about their programming experience. We gave four options to the participants for each case—the class/interface is cohesive, the class/interface is partially cohesive, the class/interface is incohesive, cohesiveness could not be determined.

We consider five commonly used variants of the metric, LCOM1 [23], LCOM2 [6], LCOM3 [10], LCOM4 [24], and LCOM5 [25] for the experiment and implemented these algorithms. We computed LCOM for all the eight cases using these methods and compared them.

2) Results: We received 13 complete responses from our survey to establish ground truth. Figure 5 shows that the participants were highly experienced with software development. They had 3–5 years of experience at the minimum (23%) while the majority of the participants belonged to the 6–10 years of experience category (46%).

Figure 6 summarizes the responses from the participants. It is evident from the survey that *Case1–2*, *Case4*, and *Case6* are considered *cohesive* by the majority of the participants. Similarly, *Case3* is tagged as *partially cohesive* and *Case5* is perceived as *incohesive*. Participants marked *Case7–8* as cases where cohesiveness could not be determined because the available information is not sufficient. We consider the observations obtained from the survey as our ground truth.

Then, we applied the existing LCOM algorithms and computed the metric value for each of the eight cases. Table III

[2]https://anonymous.4open.science/r/5f0fd91a-8884-4977-8edb-47950f7a2f13/

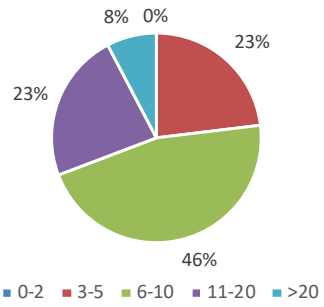

Fig. 5. Software development experience in years of the survey participants

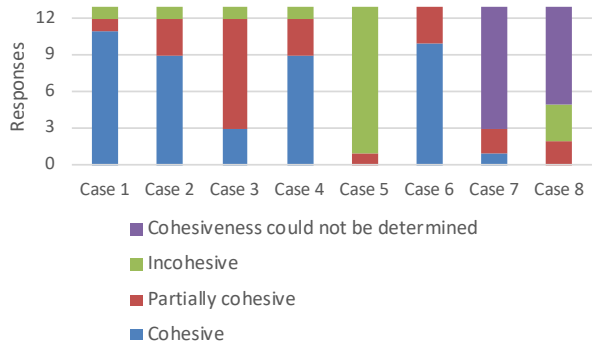

Fig. 6. Responses to the survey for each case considered for establishing ground truth

presents a comparison of the computed values using existing methods with their corresponding expected ground truth.

Case	L1	L2	L3	L4	L5	L	Ground truth
1	1	0	1	1	0.67	0.00	Cohesive
2	2	1	2	2	0.67	0.00	Cohesive
3	2	1	2	2	0.67	0.67	Partially cohesive
4	2	1	2	2	0.50	0.00	Cohesive
5	3	3	3	3	1.00	1.00	Incohesive
6	2	1	2	1	0.83	0.00	Cohesive
7	3	0	3	3	0.00	-1.00	Could not be determined
8	3	0	3	2	0.00	-1.00	Could not be determined

TABLE III
COMPARISON OF THE LCOM COMPUTED BY EXISTING METHODS WITH THE GROUND TRUTH. HERE, L1–5 REFER TO THE COMMONLY USED LCOM VARIANTS AND L REFERS TO THE LCOM METRIC PROPOSED IN THIS STUDY

We derive the following observations from the experiment.

- LCOM1–4 take into account only the instance attributes of a class, ignoring any static attributes. Although the dynamic properties of a static attribute differ from the instance attributes (for instance, static attributes can be accessed without creating an object), these attributes are part of the class they reside in. In the context of a metric that measures the similarity among class members, their dynamic property is irrelevant. Therefore, ignoring static attributes while assessing cohesion is inappropriate.
- The existing LCOM algorithms fail to distinguish the cases where the metric cannot be measured, from the perfectly cohesive cases, by always emitting the lowest metric value in the former cases. For instance, LCOM2 reports 0 not only when the type under measurement is completely cohesive (Case1), but also when the type is an interface (Case7) and a utility class, *i.e.,* a class with no attributes (Case8). Such an approach produces the illusion to the user that all cases with a metric value of zero are cohesive, while in reality, the algorithm was not provided with enough information to measure the metric and the algorithm fails to communicate this to the user.
- Method invocations within a class show that methods are working together to achieve a goal and thus must be considered while computing LCOM. However, LCOM1–3 and LCOM5 do not consider method invocations to compute the metric.
- Furthermore, the existing LCOM implementations focus on the common attribute access among methods within a class; however, they ignore common attribute access where the attribute is defined in a superclass. Classes are extensions of their superclasses, and it is very common to elevate data and method members to superclasses to avoid duplication among siblings. Hence, two methods that share attribute access or method invocation that is defined in a superclass contribute to cohesion and thus must be considered while computing the metric.
- Lastly, Fenton and Pfleeger [1] stated that a metric may follow a suitable measurement scale (such as nominal, ratio, and absolute) depending on the aspect being measured. LCOM1–4 measure cohesion on an absolute scale that may emit an arbitrary large number as the metric value making it almost impossible for the user to gain any insight from it. For instance, given that m is the number of methods of a class, the maximum value that LCOM2 may produce is $(m \times (m-1))/2$, which could be a considerable number for large classes. To facilitate metric interpretation and comparison, bounded concept such as cohesion must be better represented by a normalized value.

We summarize the results of the experiment below.

- Existing methods do not consider relevant source code semantics such as static fields and method invocations to compute cohesiveness. Lack of such information leads to produce a metric value that is far from the real.
- The existing LCOM algorithms fail to distinguish the cases where the metric cannot be measured, from the perfectly cohesive cases. Such an approach produces an illusion to the user that all cases with a minimum metric value are cohesive.
- The majority of the existing methods produce metric value on an absolute scale giving no clue to the user to interpret the value and giving no indication about what could be a good value to target for.

B. RQ2.2: Our Proposed Approach for LCOM Computation

1) Approach: To address the above-discussed deficiencies, we propose a new method referred to as YALCOM (Yet Another Lack of Cohesion in Methods) to compute the LCOM metric. Algorithm 1 presents the mechanism to compute the metric.

```
Input: Type t
Output: LCOM value
G = initialize a graph
if isMetricComputable(t) then
    G.addVertex(t.methods())
    G.addVertex(t.attributes())
    G.addVertex(t.supertype().attributes())
    for m : t.methods() do
        G.addEdge(m, m.attributesAccessed())
        G.addEdge(m, m.methodInvocations())
    end
    d = G.disconnectedGraphs()
    if d > 1 then
        LCOM = d/t.methods().size()
    else
        LCOM = 0
    end
else
    LCOM = -1
end

isMetricComputable()
Input: t
Output: True/False
begin
    if t.methods().Count = 0 Or t.isInterface() then
        return False
    end
    return True
end
```

Algorithm 1: YALCOM—The proposed LCOM metric

The algorithm takes a type *i.e.,* a class or an interface as an input. The algorithm returns -1 when the algorithm finds that the metric is not computable otherwise it returns a LCOM metric value $[0, 1]$. The metric is not computable when the number of methods is zero, or when the analyzed type is an interface. The algorithm creates a graph where the methods and attributes of the class are treated as vertices. Here, attributes from superclasses that are accessible from the class are also included. Relationships, *i.e.,* field accesses and method invocations, among the methods and attributes, make the edges. For example, if a method $\mathtt{m}_1$ accesses attributes $\mathtt{a}_1$ and $\mathtt{a}_2$ as well as calls method $\mathtt{m}_2$, then the node corresponding to method $\mathtt{m}_1$ will have edges to nodes representing attributes $\mathtt{a}_1$ and $\mathtt{a}_2$ as well as to method $\mathtt{m}_2$. Once the graph is constructed for the input class, the algorithm finds the disconnected subgraphs of methods. If the number of disconnected subgraphs is one, then all the attributes and methods are connected to each other and hence the class is perfectly cohesive (and thus assigned as 0 as the metric value). If the number of disconnected subgraphs is more than one, it implies that there are many *islands* of functionality within the class and hence the class is not cohesive. Here, the higher number of such subgraphs implies poorer cohesion. We compute the metric by dividing the number of disconnected subgraphs by the number of methods in the class. Since the number of disconnected subgraphs cannot be more than the number of methods (when none of the methods is associated with rest of the methods in the class), the maximum value that the algorithm can produce is 1.

2) Results: We evaluate the proposed algorithm on the established ground truth (refer to Section V-A). Table III shows the LCOM values (column L) produced by the algorithm along with the expected values shown by the ground truth column for all cases. The table clearly shows that the proposed algorithm computes the metric accurately for all the cases.

In addition to the ground truth cases, we manually checked the metric values computed by the proposed method on a few Java repositories. We present two examples one each from AWS Dynamodb Encryption[3] and Apache Metamodel[4] repositories. Figure 7 presents the relationships of methods and attributes of class `Builder` in Dynamodb Encryption repository. Methods are shown by square-ended rectangles whereas attributes are shown in rounded rectangles. Arrows show *access* relationship from a method to an attribute and *method invocation* from a method to another method. In this example, there are total of five methods and the method nodes make four disconnected subgraphs and thus the computed LCOM would be $4/5 = 0.8$ which is considered an incohesive class. Another example is from Apache Metamodel repository. Figure 8 shows methods and fields along with their relationships for class `ColumnTypeResolver`. There are seven methods; they are well-connected with the rest of the nodes in the graph. Hence the number of disconnected graphs is one and that implies that the class is perfectly cohesive with LCOM $= 0$.

> The ground truth test, as well as manual verification, provide sufficient indications that the proposed LCOM algorithm produces metric values as expected by the ground truth and performs superior compared to existing LCOM computation methods.

C. RQ2.3: Quantitative Analysis and Comparison

This part of the exploration carries out a quantitative analysis to compare the existing algorithms and their deviation from the values of the metric perceived by the developers.

1) Approach: To carry out a quantitative analysis, we downloaded high-quality Java repositories from GitHub, compiled, and analyzed them to measure LCOM metric values for all the types in all the analyzed repositories.

Download subject systems:
We used the following protocol to identify our subject systems.

[3]https://github.com/aws/aws-dynamodb-encryption-java
[4]https://github.com/apache/metamodel

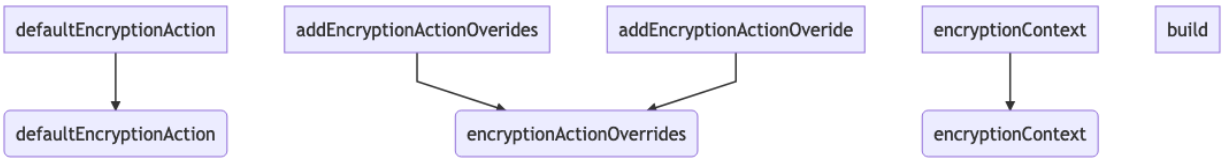

Fig. 7. Class methods, fields, and relationship among them for class `Builder` belonging to AWS Dynamodb Encryption repository

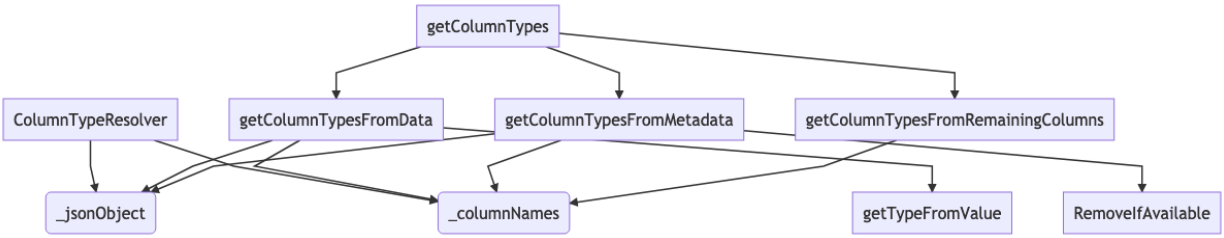

Fig. 8. Class methods, fields, and relationship among them for class `ColumnTypeResolver` belonging to Apache Metamodel repository

- We use RepoReapers [27] to filter out low-quality and too small repositories among the abundant repositories present on Github. RepoReapers analyzed a huge number of GitHub repositories and evaluated each of the repositories on eight dimensions providing a fair idea about their quality characteristics. These dimensions are architecture (as evidence of code organization), continuous integration and unit testing (as evidence of quality), community and documentation (as evidence of collaboration), history, issues (as evidence of sustained evolution), and license (as evidence of accountability).
- We select all the repositories containing Java code where all the eight RepoReapers' dimensions have suitable scores. We consider a score suitable if it has a value greater than zero.
- Next, we remove the repositories that have less than 10 stars as well as contain less than $1,000$ lines of code.
- Following these criteria, we selected and downloaded 522 repositories.

A complete list of the selected Java repositories along with their analyzed results can be found in our replication package.

Analyze subject systems:
Our implementation of LCOM algorithms uses Eclipse JDT[5] to prepare Abstract Syntax Tree and resolve symbols. It is necessary for the JDT libraries to have compiled `.class` files along with source code files to resolve the symbols and identify the various kinds of relationships correctly. Therefore, we first compile the subject systems. To carry out the compilation automatically, we checked for the usage of one of the two commonly used build systems for Java *i.e., Maven* and *Gradle* and trigger the corresponding command to compile the projects automatically. This approach could compile a total of 261 repositories; we discarded rest of the repositories from further analysis.

[5]https://www.eclipse.org/jdt/

We implement the existing methods to compute LCOM as well as our proposed algorithm. We analyzed all the 261 repositories and computed LCOM metric values for all the types contained in the repositories. We consolidated the obtained results, analyzed them, and documented our observations.

Euclidean distance is a common method to find the straight-line difference in Euclidean space [28] between two points. Equation 1 shows the mechanism to compute collective Euclidean distance for a series of points in Euclidean n-space [28]. The computed distance shows the extent to which two series are similar or different.

$$d(\vec{u}, \vec{v}) = \sqrt{\sum_{n=1}^{n} (u_i - v_i)^2} \tag{1}$$

We compute the Euclidean distance between our proposed method and all the existing LCOM methods. Since four out of five existing methods for LCOM produce metric values that are not normalized, we computed their normalized values based on the min-max normalization technique [29]—we use the minimum and maximum values of LCOM values for each specific repository individually for each considered method to produce their corresponding normalized values. The normalization process brought all the metrics on the same scale to have an appropriate comparison. We computed Euclidean distance for the normalized values also.

2) Results: We analyzed all the $90,029$ types belonging to 261 repositories and computed LCOM metric value using each of the existing LCOM algorithms along with our proposed method. Then, we computed Euclidean distance for all the types between LCOM1–5 and YALCOM. Before computing the distance, we first identified and separated the types where the number of methods is zero or the type is an interface. We found that there are $15,356$ such types. It implies that for approximately 17% of the types, the current LCOM methods were not able to compute the metric value because either the

type has no methods or the type is an interface and hence there are no relationships between methods and attributes. Despite this, the existing methods show a perfect 0 (or 1 for LCOM3–4) indicating they are perfectly cohesive types giving a false perception to the user of the metric about the class cohesion.

We computed the Euclidean distance for the rest of the types. Figure 9 shows the computed absolute Euclidean distance between each of the existing LCOM method and YALCOM. It is evident from the figure that the metric values produced by the existing methods are hugely different than the proposed method. Relatively, LCOM1–2 values are massively different from the proposed approach because their LCOM values are on an absolute scale and hence, often, the produced values are arbitrary very large. For example, the maximum metric values produced by LCOM1 and LCOM2 are $22,221,184$ and $20,689,090$ for class `PlanProto` belonging to Apache Tajo[6] repository. We figured out that it is the case because the `PlanProto` class has a total of $6,893$ methods belonging to its 207 nested types where the huge set of methods do not interact directly or indirectly with other methods to form subgraphs that leads to the huge LCOM value. For the same class, LCOM3 and LCOM4 report $4,081$ and $4,076$ respectively while LCOM5 gives 0.99. Such arbitrary large metric values not only confuse the users since it is not known whether a specific value is good or bad but also fails to provide some actionable insight. Table IV summarizes some key characteristics of the analyzed LCOM values.

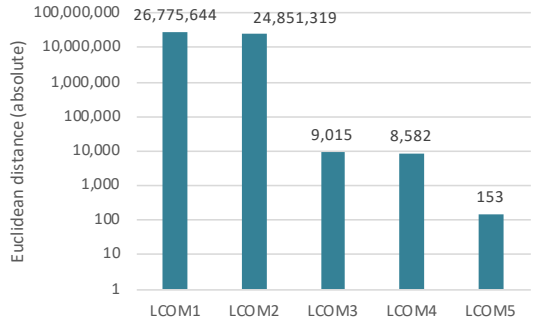

Fig. 9. Euclidean distance (absolute) between LCOM values generated by each of the LCOM1–5 and the proposed YALCOM

Algorithm	Maximum	Minimum	Median	Average
LCOM1	22,221,184	0	4	960.77
LCOM2	20,689,090	0	0	817.23
LCOM3	4,081	1	2	6.46
LCOM4	4,076	1	2	5.26
LCOM5	2	0	0.5	0.44
YALCOM	1	0	0.22	0.36

TABLE IV
CHARACTERISTICS OF THE LCOM VALUES GENERATED BY ALL CONSIDERED LCOM METHODS DERIVED FROM OUR QUANTITATIVE ANALYSIS

One may argue that it is not fair to compare the absolute values of the LCOM, though generated correctly by definition, from existing methods where the values are not normalized. To observe the difference and compare them on the same scale, we normalized the metric values in the range of [0, 1]. Figure 10 presents the comparison of all the considered LCOM methods after the normalization. Though the scale is different, the figure reflects a similar pattern as we see in Figure 9. LCOM5 computes the closest values to YALCOM compared to the rest of the existing methods. However, in summary, all the existing methods produce LCOM values very different from the expected values.

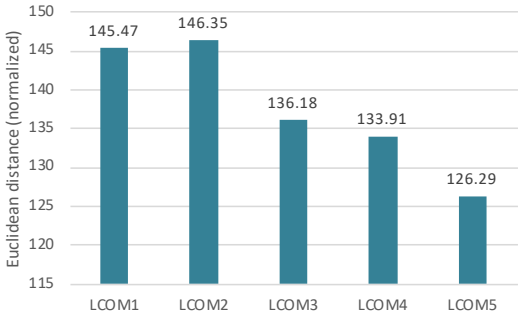

Fig. 10. Euclidean distance (normalized) between LCOM values generated by each of the LCOM1–5 and the proposed YALCOM

We summarize the findings of the quantitative analysis as follows.

- The existing LCOM algorithms report a perfect score indicating the type is perfectly cohesive even when there are not enough details available to deduce the conclusion. We found 17% of the types belonging to this category. The deficiency gives a false perception to the user about the cohesiveness of the class.
- The LCOM values computed from existing methods produce values that are far from the values produced by the validated method. Hence, the existing methods do not capture the notion of the metric accurately.

VI. Discussion and Implications

We identify the following implications for the software engineering community.

- For software engineering researchers, the exploration presented in the study reveals ample opportunities to fill the identified gaps by, for instance, proposing relevant and effective code quality metrics. Furthermore, the study paves the way to look for alternatives to deterministic and traditional code quality metrics. Additionally, the presented study invites researchers to gauge the effectiveness and soundness of existing metrics by presenting a case study of the LCOM metric. Similar studies can be explored for other commonly used metrics.

[6]https://github.com/apache/tajo

- Software tool developers and vendors may take a cue from the study and innovate novel mechanisms to perceive the specific quality aspects of a software system. Also, they could refine the existing metric implementations that better represent the measured aspect.
- The take away for a developer from this study is that the metrics that are being used today do not necessarily represent the aspect that the metric claim to and hence the over-reliance on strict quality processes based only on such metrics may not be needed.

Naturally, the discussed deficiencies and criticism can be addressed by proposing new mechanisms and metrics. Alternatively, advances in machine learning technologies and their practical feasibility have introduced an additional twist regarding code quality metrics. Traditionally, code quality metrics based on intuition, theory, or empirical evidence, offered a simple shortcut to assess source code quality. Essentially, researchers find rules and heuristics, and justify them with their theory and empirical evidence. On the contrary, machine learning-based methods look at the data and derive their heuristics from them. This implies that the community may come up with advanced machine learning-based quality estimation models, that are not necessarily derived from existing metrics but are trained on a gamut of software engineering development data. Such an approach offers advantages, issues, and challenges. On the one hand, machine learning approaches may offer increased accuracy by flagging quality problems that are not discernible through simple metrics. On the other hand, machine learning results might be less actionable, since they may point to questionable code without providing an explanation on why the specific code is problematic. And yet, a positive aspect of the machine learning approach may be that developers will be less likely to game the system by manipulating the metrics and their thresholds, rather than actually improving the software's quality. A widely known challenge concerning machine learning-based methods is the availability of well-curated labeled data, as the success of these methods depends heavily on them. Our take is that the software engineering community should consider machine learning approaches regarding the judgment of code quality as another type of metric, one that, in common with existing metrics, should stand or fall based on its usefulness backed by empirical evidence.

VII. Replication Package

We provide the following in our anonymized online replication package[7].

- Source code of all the cases used to establish the ground truth
- Implementation of existing LCOM algorithms as well as the proposed algorithm
- A list of Java repositories used for quantitative analysis
- Results generated from the LCOM implementations for all the considered repositories
- Consolidated results containing LCOM absolute as well as normalized values for all the types

VIII. Threats to Validity

Construct validity concerns the appropriateness of observations and inferences made on the basis of measurements taken during the study. We propose a new algorithm to compute the LCOM metric and derive observations based on the algorithm and the corresponding implementation. We validated the proposed algorithm against a ground truth that itself is prepared from the inputs provided by the experienced developers. In addition, we carried out random manual testing to ensure that the implementation is working as intended.

We derive a set of cases to cover different scenarios of a class configuration with methods, attributes, and their relationships. Though we tried to cover all common cases but it is possible that we have missed other relevant cases. We encourage researchers to extend the presented cases so that a benchmark to measure cohesion can be set up.

External validity concerns the generalizability and repeatability of the produced results. We carried out our experiments on a large number of repositories to keep the results generic. We have made all the scripts as well as data produced in the experiments available to promote reproducibility.

Internal validity refers to the validity of the research findings. It is primarily concerned with controlling the extraneous variables and external influences that may impact the outcome. Our online survey carried out by distributing it over the Internet via social media. Given the anonymous nature, we do not know whether the participants were indeed software developers and actually represent a common perception of a typical software developer. However, given a relatively large number of participants (78), we rely on the overall trend that we observe from the obtained choices.

IX. Conclusions

We provide an overview of the common code quality metrics in use since the 1960s, present criticism derived from the literature, and outline deficiencies in the currently used code quality metrics. We also present a developers' perspective, which clearly shows that the current set of metrics is not sufficient for a variety of reasons. The exploration reveals deficiencies in the current set of code quality metrics such as poor support to assess software architecture quality as well as to assess testing and infrastructure aspects of software systems. As a case study, we provide a detailed qualitative and quantitative analysis of presently used LCOM metric algorithms and point out specific deficiencies. We proposed a new algorithm to compute LCOM and show that it produces metric values as expected from the established ground truth. We also show that the presently used LCOM algorithms are not sound.

In the future, we would like to propose traditional code quality metrics for sub-domains, such as infrastructure, where sufficient metrics are not yet invented. We are also interested to explore machine learning approaches to spot source code entities where code quality is not optimal.

[7]https://anonymous.4open.science/r/5f0fd91a-8884-4977-8edb-47950f7a2f13/

REFERENCES

[1] N. Fenton and S. L. Pfleeger, *Software Metrics (2Nd Ed.): A Rigorous and Practical Approach*. Boston, MA, USA: PWS Publishing Co., 1997.

[2] A. L. Timóteo, A. Álvaro, E. S. De Almeida, and S. R. de Lemos Meira, "Software metrics: a survey," 2008, online: http://www.cin.ufpe.br/ alt/mestrado/qsic2008.pdf.

[3] S. H. Kan, *Metrics and Models in Software Quality Engineering*, 2nd ed. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2002.

[4] T. J. McCabe, "A complexity measure," *IEEE Transactions on Software Engineering*, vol. SE-2, no. 4, pp. 308–320, Dec 1976.

[5] M. H. Halstead, *Elements of Software Science (Operating and Programming Systems Series)*. New York, NY, USA: Elsevier Science Inc., 1977.

[6] S. R. Chidamber and C. F. Kemerer, "A metrics suite for object oriented design," *IEEE Transaction of Software Engineering*, vol. 20, no. 6, pp. 476–493, Jun. 1994.

[7] F. B. Abreu and R. Carapuça, "Object-oriented software engineering: Measuring and controlling the development process," in *Proceedings of the 4th international conference on software quality*, vol. 186, 1994, pp. 1–8.

[8] R. C. Martin, *Designing Object-oriented C++ Applications: Using the Booch Method*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1995.

[9] M. D'Ambros, M. Lanza, and R. Robbes, "An extensive comparison of bug prediction approaches," in *2010 7th IEEE Working Conference on Mining Software Repositories (MSR 2010)*. IEEE, 2010, pp. 31–41.

[10] W. Li and S. Henry, "Object-oriented metrics that predict maintainability," *Journal of Systems and Software*, vol. 23, no. 2, pp. 111 – 122, 1993, object-Oriented Software.

[11] M. Lanza and R. Marinescu, *Object-oriented metrics in practice: using software metrics to characterize, evaluate, and improve the design of object-oriented systems*. Springer Science & Business Media, 2007.

[12] C. R. Symons, "Function point analysis: Difficulties and improvements," *IEEE Trans. Softw. Eng.*, vol. 14, no. 1, pp. 2–11, Jan. 1988.

[13] S. Henry and D. Kafura, "Software structure metrics based on information flow," *IEEE Transactions on Software Engineering*, vol. SE-7, no. 5, pp. 510–518, Sep. 1981.

[14] D. Coleman, D. Ash, B. Lowther, and P. Oman, "Using metrics to evaluate software system maintainability," *Computer*, vol. 27, no. 8, pp. 44–49, 1994.

[15] T. Mayer and T. Hall, "A critical analysis of current oo design metrics," *Software Quality Journal*, vol. 8, no. 2, pp. 97–110, Oct 1999.

[16] M. Shepperd and D. Ince, "A critique of three metrics," *Journal of Systems and Software*, vol. 26, no. 3, pp. 197 – 210, 1994.

[17] N. E. Fenton and M. Neil, "Software metrics: success, failures and new directions," *Journal of Systems and Software*, vol. 47, no. 2-3, pp. 149–157, Jul. 1999.

[18] D. I. Sjøberg, B. Anda, and A. Mockus, "Questioning software maintenance metrics: A comparative case study," in *Proceedings of the ACM-IEEE International Symposium on Empirical Software Engineering and Measurement*, ser. ESEM '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 107–110. [Online]. Available: https://doi.org/10.1145/2372251.2372269

[19] A. V. Deursen, "Think twice before using the maintainability index," 2014, https://avandeursen.com/2014/08/29/think-twice-before-using-the-maintainability-index/.

[20] T. Kuipers and J. Visser, "Maintainability index revisited–position paper," in *Special session on system quality and maintainability (SQM 2007) of the 11th European conference on software maintenance and reengineering (CSMR 2007)*. Citeseer, 2007.

[21] T. Sharma and D. Spinellis, "A survey on software smells," *Journal of Systems and Software*, vol. 138, pp. 158 – 173, 2018.

[22] G. A. Campbell, "Cognitive Complexity, Because Testability != Understandability," 2016, https://blog.sonarsource.com/cognitive-complexity-because-testability-understandability.

[23] S. R. Chidamber and C. F. Kemerer, "Towards a metrics suite for object oriented design," in *Conference Proceedings on Object-oriented Programming Systems, Languages, and Applications*, ser. OOPSLA '91. New York, NY, USA: ACM, 1991, pp. 197–211.

[24] M. Hitz and B. Montazeri, "Chidamber and kemerer's metrics suite: a measurement theory perspective," *IEEE Transactions on Software Engineering*, vol. 22, no. 4, pp. 267–271, April 1996.

[25] B. Henderson-Sellers, L. L. Constantine, and I. M. Graham, "Coupling and cohesion (towards a valid metrics suite for object-oriented analysis and design)," *Object Oriented Systems*, vol. 3, pp. 143–158, 1996.

[26] L. H. Etzkorn, S. E. Gholston, J. L. Fortune, C. E. Stein, D. Utley, P. A. Farrington, and G. W. Cox, "A comparison of cohesion metrics for object-oriented systems," *Information and Software Technology*, vol. 46, no. 10, pp. 677 – 687, 2004.

[27] N. Munaiah, S. Kroh, C. Cabrey, and M. Nagappan, "Curating github for engineered software projects," *Empirical Software Engineering*, vol. 22, no. 6, pp. 3219–3253, Dec 2017. [Online]. Available: https://doi.org/10.1007/s10664-017-9512-6

[28] B. O'Neill, "Chapter 1 - calculus on euclidean space," in *Elementary Differential Geometry (Second Edition)*, second edition ed., B. O'Neill, Ed. Boston: Academic Press, 2006, pp. 3 – 42. [Online]. Available: http://www.sciencedirect.com/science/article/pii/B9780120887354500055

[29] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal, "Chapter 8 - data transformations," in *Data Mining (Fourth Edition)*, fourth edition ed., I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal, Eds. Morgan Kaufmann, 2017, pp. 285 – 334. [Online]. Available: http://www.sciencedirect.com/science/article/pii/B9780128042915000088