

What to Fix? Distinguishing between design and non-design rules in automated tools

Neil A. Ernst, Stephany Bellomo, Ipek Ozkaya, Robert L. Nord
Carnegie Mellon University Software Engineering Institute
Pittsburgh, PA

Email: {nernst,sbellomo,ozkaya,rn}@sei.cmu.edu

Abstract—Technical debt—design shortcuts taken to optimize for delivery speed—is a critical part of long-term software costs. Consequently, automatically detecting technical debt is a high priority for software practitioners. Software quality tool vendors have responded to this need by positioning their tools to detect and manage technical debt. While these tools bundle a number of rules, it is hard for users to understand which rules identify design issues, as opposed to syntactic quality. This is important, since previous studies have revealed the most significant technical debt is related to design issues. Other research has focused on comparing these tools on open source projects, but these comparisons have not looked at whether the rules were relevant to design. We conducted an empirical study using a structured categorization approach, and manually classify 466 software quality rules from three industry tools—CAST, SonarQube, and NDepend. We found that most of these rules were easily labeled as either not design (55%) or design (19%). The remainder (26%) resulted in disagreements among the labelers. Our results are a first step in formalizing a definition of a design rule, in order to support automatic detection.

Index Terms—Software quality, software design, software cost

I. INTRODUCTION

Static analysis tools evaluate software quality using rules that cover many languages, quality characteristics, and paradigms [1]. For example, the SonarQube tool¹ can handle C, C++, Java, Javascript, and many others, using language-dependent rule sets (called quality profiles). Software quality rules have accumulated as practitioners gradually recognize code ‘smells’ and poor practices. The first tool to automate rule-checking was the C language tool `lint` in 1979 [2]. Today most languages have linters which extend the standard syntax and type checking efforts of a compiler/interpreter to warnings and smell detection. Static analysis tools have traditionally focused on what we might call the *code-level*, rather than design. For instance, Johnson’s initial description of `lint` mentions type rules, portability restrictions, and “a number of wasteful, or error prone, constructions which nevertheless are, strictly speaking, legal” [2].

Increasingly, however, quality rules are targeting design problems, such as paradigm violations or architecture pattern violations (e.g., the work of Aniche et al. on MVC frameworks [3]). This is because design problems are often more significant than coding errors for long-term software maintenance

costs. This view is supported by the results of recent survey and interview [4] and issue tracker analysis [5] studies we performed, which found that syntax and coding problems are rarely important sources of these long-term costs, which we call technical debt; instead, sources of debt are created by poor design choices.

The challenge for users of static analysis tools is making sense of the results [1], [6]. Static analysis tools often generate many false positives, leading developers to ignore the results [7]. One potential improvement to this problem is to separate *design rules* from other rules. We examine the software quality rules of three typical tools (CAST, SonarQube, and NDepend) to understand the extent to which their quality rules are design-related. This raises the question of what we mean by *design*, a thorny question in software engineering research [8, p.14]. We use an extensional definition [9, ch 1.1] of design by creating a design rule classification rubric over several iterations, using rater agreement on classification labels as our metric. Design is clearly more than what the (imperfect) rubric suggests. Most importantly, it is limited to the rules we used as input, and each rater’s understanding of design. We expand on this notion of design in Section II, below.

Our contributions and findings in this work include:

- A conceptual model for thinking about design problem detection.
- A classification rubric for evaluating design rules.
- Existing tools do have rules that can check for design quality. 19% of the rules we examined were design-related.
- The rules we analyzed included examples of complex design concepts, such as design pattern conformance and run-time quality improvements. 68% of the rules that were labelled as design rules were examples of such rules.

II. BACKGROUND

A. Separating Design from Non-Design Rules

Our goal is to be able to label a given quality rule as either design or non-design. This requires a rubric, a schema that can guide labeling (in the quantitative context, this would be a decision tree). To create this rubric, we began by bootstrapping several concepts of design into an initial taxonomy, shown in Fig. 1. This we based on our experience from conducting numerous software architecture assessments, as well as related

¹sonarqube.com

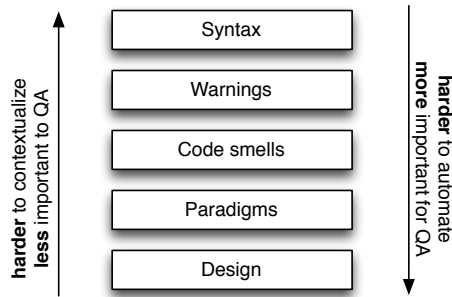


Fig. 1. Visual depiction of levels of quality analysis.

literature. This gave a rough classification of design as the upper levels of that taxonomy, i.e., levels “Paradigm” and “Design”. We defined each level as follows:

- 1) **Syntax**—the interpreter or compiler’s components (parsers, lexers, intermediate forms) assess language/configuration file syntax correctness. (Note that for some specialized cases, levels below this exist: to optimize compiled code with assembler, for example.)
- 2) **Warning**—compiler warnings (e.g., `-wall`. These warnings tend to be close to syntax in their complexity. For example, technically a fall through switch statement is syntactically correct in Java, but there is the `-Xlint:fallthrough` tag to catch this and force a compile failure. Often IDEs such as Eclipse will flag these automatically with warning icons. Linters traditionally operate at this level of analysis.
- 3) **Code smells**—Code conforms to commonly accepted best practices for that language (e.g., for Java, visibility modifiers are suitable, in C, no buffer overflows, memory is released appropriately). Some cross-language practices apply: documentation, tested, and so on.
- 4) **Paradigm**—Would someone writing object-oriented, functional, embedded, etc. code consider this reasonable? Includes principles like SOLID, functional side effects, memory management, etc. Distributed code demonstrates awareness of fundamentals of distributed computing. Also includes proper use of language idioms (e.g., proper use of Javascript callbacks, Ruby blocks).
- 5) **Design quality**—Given the knowledge available, the code is architecturally appropriate for the applicable quality attribute requirement (QAR), for example, modular, performant, and secure. The key here is understanding the relevant QARs.

Table I shows representative rules from all five levels of quality analysis: Java language specification (for level 1: syntax/compile correctness), Java warnings in Eclipse (level 2), and SonarQube rules detecting code smells, paradigm suitability, and design quality (levels 3/4/5). We use this taxonomy to bootstrap our initial rubric from Section IV-C.

We used a rubric to separate design from non-design rules.

To do so, we used inter-rater agreement on the design/not-design labels to assess rubric performance (i.e., how accurately it captured the rater understanding of the two classes). We then reconciled our labels where there was disagreement, and evolved the rubric to account for the lessons from the reconciliation. We then evaluated the rubric a second time, using different raters than the first iteration. To a first approximation, then, *design* is whatever eventual rubric most improves rater agreement (Cohen’s kappa [10]) for classifying rules, using appropriate safeguards against bias.

B. State of the Practice

We give a brief overview of the state of the practice with respect to design rule detection, derived from the static analysis tool literature.

While `lint` and its relatives (`jslint`, `pylint`, `FindBugs`, etc.) focus on code quality that compilers do not catch, increasingly vendors are offering much more than what linters traditionally supported. These *software quality tools* differ in several ways.

- 1) They take a wider perspective on software quality than code faults, incorporating also aspects of design quality;
- 2) They are used in managerial roles to check historical trends, as much as for individual developers to check specific faults. Most typically, software quality tools include time as a significant dimension to track trends. Linters tend to flag individual rule violations for follow-up, but provide little support for trends.
- 3) They focus on higher-order analysis. Sadowski et al. [7] call these ‘more sophisticated static analysis results’ and choose not to focus on them, seeing a higher payoff with simpler rules with clear fixes. However, they are focused on developer interaction where time is critical.
- 4) Many times they consist of a portfolio/ensemble approach to software metrics and rules, integrating diverse reports on bugs, rule violations, and observations on correlations between different rules.

Thus a software quality tool has a time-series focus, and pays more attention to code-smells, idioms, and design-related issues. These tools fall into the software analytics space [11], an emerging research area that seeks to leverage the large volumes of data collected as part of the development process (e.g., lines of code, bugs closed, team structure and activity). No “tools” beside IDEs provide compiler level assistance (Eclipse, for example, flags unknown Java types). Very few tools provide integration with compiler flags (see Java’s compiler flags²).

Software quality monitoring goals can be classified by looking at tool *coverage*. Coverage measures the extent to which quality aspects of software can be measured. That is, for given design questions, what answers can be garnered from tools? This might be called coverage or cognitive support—how well the tools support the analysis goals users have. The inverse of this question is to see what answers the tools provide

²<https://docs.oracle.com/javase/7/docs/technotes/tools/windows/javac.html#xlintwarnings>

TABLE I
EXAMPLES OF QUALITY RULES AND THEIR SCOPE

Message	Level	QAR	Classification
< X > cannot be resolved to a type	Syntax/compiler	-	ND
Return type missing	Syntax/compiler	-	ND
Empty statement after 'if/else'	Warning/XLint	-	ND
Deprecated	XLint	Maintainability	ND
Overrides	XLint	Maintainability	ND
Catch clauses should do more than rethrow	Good code	Maintainability	ND
Public methods should throw at most one checked exception	Good code	Maintainability	ND
Threads should not be used where "Runnables" are expected	Paradigm	Concurrency	DR
Abstract classes without fields should be converted to interfaces	Paradigm	Modularity	DR
Component Balance within 1 S.D. of industry average	Design	Modularity	DR
Classes should not be coupled to too many other classes	Design	Modularity	DR

that are simply not useful, or do not map to any known existing questions.

Tools that incorporate aspects of design analysis ideally provide reliable, automated, repeatable results to address the following goals.

- Find poor architectural decisions and shortcuts and identify refactoring opportunities [4], [12].
- Understand when payoff is economically justified [13].
- Find increased numbers of {defects bugs churn} above baseline (hotspots) [14].
- Understand the trends and rate of change in key indicators (such as lines of code, test coverage, or rule violations) [15].
- Provide traceability across architectural tiers, frameworks, and languages [3].

III. STUDY SCOPE

Context is important to the question of how to use automated tools for finding software quality problems. This is because there are contextual factors that make some of what a tool might do superfluous or even irritating for a given project [1].

Fig. 2 shows a conceptual model to organize this space and better articulate our research approach and contributions.

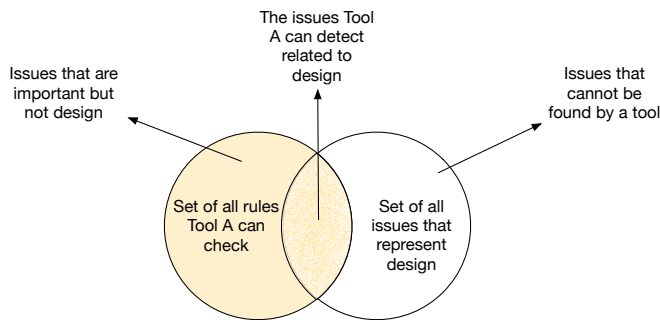


Fig. 2. What tools can detect, and what projects want detected. Design issues and rules are both context-dependent on the project properties.

Figure 3 captures our overall research agenda. Modulo the persistent challenge of diffusion of innovations, the goal is to provide insight into moving the two circles closer

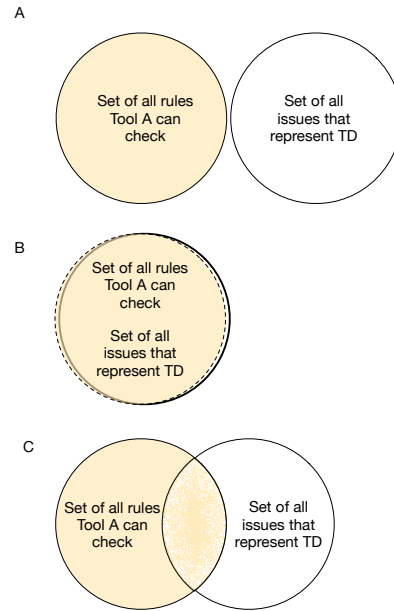


Fig. 3. **A.** The worst case; tools have no capabilities needed. Result: tool goes unused. **B.** Ideal case. Near-perfect overlap, and the tool is used constantly. **C.** Reality. The tool has some useful capabilities, but there is a large unmet need.

together. This paper focuses on understanding and improving the automation of design rules. That may be with new tool capabilities (such as usability improvements, new rules, better integration). The other aspect is process change in the organization or project (e.g., moving to continuous integration with automated quality checks).

IV. METHODOLOGY

Our research focuses on how some current tools detect design-related technical debt. This approaches the issue of design problem detection from the left side of Fig. 2; rather than asking what support a given project needs, we look at what support a tool provides to understand design-related rules in the commercial tooling landscape. In particular, we identify two questions:

- What is a *design-related* rule?

- What kinds of design-related rules do common tools provide?

We approached these two questions by evaluating how well three commercial tools manage design rules. We begin with the taxonomy in Section II, and then expand that taxonomy into a rubric for classifying rules. We then validate our rubric with experts, and report the rubric performance.

A. Initial Taxonomy

We began with the taxonomy in Section II. The creation of that taxonomy was described previously.

B. Tool and Rule Selection

We selected the three commercial tools for our study opportunistically. We chose rule sets from software quality management tools that have stated they have capabilities to detect design and technical debt: NDepend³, SonarQube⁴, and CAST⁵, and that we have access to (primarily for reasons of licensing and installability). While the tools provide a broader set of capabilities for quality management, we narrowly focused on their quality measures and rules for static analysis of code. In particular, we were able to access all three rule explanations via the tool’s documentation. Then, we chose a subset of available code quality rules. We focused on Java and .Net rules and rules that the tool documentation stated applied generically to all code under analysis. Our analysis of the rules from these tools in no way implies endorsement or critique of the tool itself.

Each of these tools encapsulate well established code and design quality principles as part of their analysis features. NDepend provides quality analysis support for .Net applications. NDepend also incorporates some architectural visualization capabilities providing dependency graph and matrix views to interpret system dependencies, element relationships and the like. For NDepend, we analyzed all rules. In NDepend, the rules are grouped into categories, one of which is *Design*. We found that while most of these rules were design-related, several were more about basic principles or metrics, and many rules outside of the design category were still design-related.

SonarQube is an open source platform for code quality analysis that provides support for several languages such as Java, C/C++, Objective-C, Python among others. SonarQube organizes its reports under three categories: bugs, vulnerabilities and code smells and propagates these into a technical debt assessment based on number of days it would take to fix these. For SonarQube, we analyzed all Java rules with priority major or higher. SonarQube rules we analyzed did not explicitly list the categories, such as design or dead code as in NDepend or architecture and application specificity as in CAST.

CAST provides similar features, in addition to dashboards to allow different stakeholders to interpret the results. CAST operates with a service-based model, while it publishes the

quality checks its tools embodies through various documentations. The rules we analyzed from CAST were either JEE or “all”.

C. Rubric Creation and Refinement

Based on the taxonomy (see Fig. 1), empirical data collected in our previous studies [4], [5], and example rules extracted from the three tools, we created our initial design rules classification rubric. The rubric is also motivated by established design and architecture principles, such as assessing the scope of an issue as local, non-local and architectural [8]. Such principles are also accepted by tool vendors. For example, CAST’s Application Intelligence Platform categorizes its rules into unit level (the rule impacts local part of application code), technology level (the rule impacts several components) and system level (the rule impacts cross boundaries and architectural elements).⁶

We first created a simple definition of a design rule rubric, then iterated on it with examples from interviews and survey responses to ‘test’ how well it handled these cases. The input for the rubric is a single rule from one of the example rule sets. The labeler (person classifying the rule) considers the rule, then applies the decision criteria to the rule. In our classification guidelines we specified that labelers should look at each rule on its own, without considering long-term accumulated impact of multiple violations of the rule. For example, numerous ‘dead stores’ may indicate a bigger problem than a single instance would.

We refined the rubric into a second version, then conducted a final round of classification, rotating assignments so two new labelers approached the dataset. The results reported below apply to this final round.

D. Validation

We validated the results with arms-length experts in software quality analysis. Each expert applied the rubric to a system they were most recently working on and commented on the extent to which it would correctly distinguish a rule as design-related or not.

V. CLASSIFICATION RUBRIC

Following our annotations with the first version of the rubric, we reconciled our labeling results to understand how the rubric failed. For example, the first version used a quality attribute filter on the rules. We found that the quality attribute filter was not helpful in reducing the label space, since nearly every rule, even syntax rules, had some quality attribute impact. The initial version also had a ‘deprecated’ filter meant to discern design-related rules in the code structure, but was similarly unhelpful; deprecation is rather an orthogonal dimension of the codebase, and a quality rule on its own.

Fig. 4 shows our final version of the rubric. The first decision point is rule scope (DP1), with three potential branches

³<https://blog.ndepend.com/technical-debt-avoid-ndepend/>

⁴<https://blog.sonarsource.com/evaluate-your-technical-debt-with-sonar/>

⁵<http://www.castsoftware.com/research-labs/technical-debt>

⁶<https://goo.gl/6nGEs3>

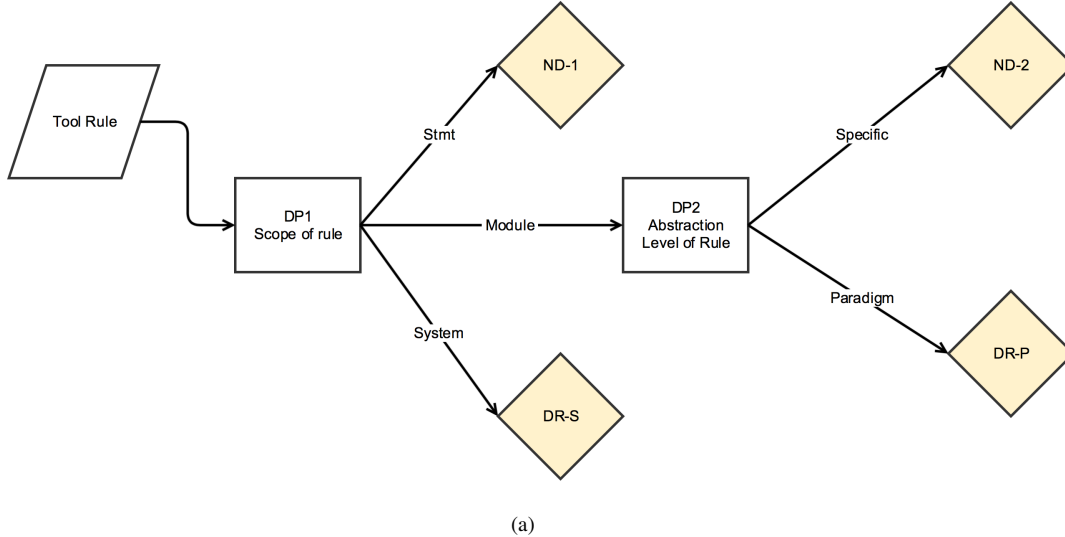


Fig. 4. Final version of Design Rule Classification Rubric

that we refer to as statement-level, module-level and system-level below. To explain how it works, we give examples of each branch of the tree below.

Statement-level: At the first filter, statement level, scope is limited to a single code statement and rules are typically syntax related. We categorized as ND-1 (not-design rule) if the rule scope is limited to single code statement (e.g., internal to method (switch, case, if/else, expression). Empty methods, dead stores, also fit here.

We give the following rule as an example here *for loop stop conditions should be invariant*. We labeled this rule ND-1 as the violation is likely to be found within a method, checking for how local variables are set before for loops, a basic coding construct.

Module-level: The next filter is the module-level. This includes groups of statements (that might be bundled into a file or package) or a language construct (method, class) that can be executed independently, reused, tested, and maintained or is a composition of other modules. This category includes module-level syntax specific violations, similar to statement level, as well as rules that check for design paradigms. This leads to a decision point for the abstraction level of the rule (DP2).

We discarded module-level syntax checking rules as non-design, ND-2 in our classification rubric. These rules typically cannot translate easily into another language or paradigm, encapsulate keywords or reserved words or concepts that can't be translated to generalizable concept (e.g., violating Spring naming conventions is a rule with no obvious commonalities in other frameworks). *Don't call GC.Collect() without calling GC.WaitForPendingFinalizers()* is an example for such rules.

Module-level rules that check for design paradigm, DR-P in our classification rubric, aim to identify rules that encapsulate known design paradigm principles. These include object-

oriented, functional, imperative programming, etc.; architectural styles, such as concurrent, model-view-control, pipe-filter, etc.; use of particular design patterns and paradigms, such as exception handling, singletons and factories, etc. *Action Classes should only call Business Classes* is a rule labeled as DR-P since it enforces an aspect of the MVC pattern.

System scope: The system scope filter includes problems detected cross system boundaries (e.g., between languages and/or architectural layers). This also includes rules where system-level metric thresholds are reported (e.g., complexity, dependency propagation). *Avoid having multiple artifacts deleting data on the same SQL table* is an example of such a rule labeled DR-S. This is a design rule because not only multiple system elements are involved, but their architectural responsibilities are critical, in this case enforcing the data model.

Our study artifacts that includes the design rules spreadsheet, categorization labeling results and categorization guidance document (with rubric) are available on a website.⁷

VI. RESULTS

A. Applying the Rubric to Software Quality Rule Sets

Out of the 466 rules we analyzed, 55% were easily labeled as not design (ND); 19% were labeled as clearly design, and the remaining had disagreements among the labelers and were hard to classify. Fig. 5 shows relative performance of the raters on each tool. In Table II we list some examples of design rules, non-design rules and hard-to-classify rules.

Table III summarizes the inter-rater agreement (using Cohen's Kappa) after applying this rubric on the data set. For simplicity of reporting, we collapse both design (DR-S and DR-P) and non-design (ND-1 and ND-2) labels together.

⁷http://www.sei.cmu.edu/architecture/research/arch_tech_debt/what-to-fix

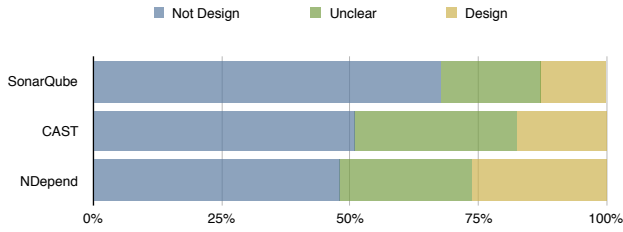


Fig. 5. Categories of rater agreement, normalized across tools. Non-design makes up the majority in all three tools.

TABLE II
EXAMPLES OF DESIGN, NON-DESIGN, AND HARD TO CLASSIFY RULES

<i>Design Rules</i>	
Action Classes should only call Business Classes	
Avoid high number of class methods invoked in response to a message	
Avoid Classes with a High Lack of Cohesion	
<i>Non-Design Rules</i>	
Try-catch blocks should not be nested	
Two branches in conditional structure should not have same implementation	
All script files should be in a specific directory	
<i>Hard to Classify</i>	
Avoid hiding attributes	
Avoid defining singleton or factory when using Spring	
Avoid declaring an exception and not throwing it	
lines of code covered [test coverage]	
Each method in an Action Class should have a small complexity	

We found it relatively easy to classify many rules on either end of the spectrum as either syntax (e.g., *methods should not return constants* or design (e.g., *action classes should only call business classes*, enforcing the implementation of the MVC pattern.) However, there were many rules that either had characteristics making it hard to classify or where our rubric could be clarified. The low Cohen’s κ values were due to a high level of disagreement over the hard-to-classify rules.

B. Validation Feedback

We validated our rubric with three senior architecture analysts from the Software Engineering Institute, not connected with our team. We had each person comment on the rubric itself, and then label a random sample ($n=74$) of only the hard

TABLE III
RATER AGREEMENT. DR=DESIGN RULE. ND=NOT DESIGN. GRAY CELLS SHOW AREAS OF RATER DISAGREEMENT.

	DR	ND		DR	ND
DR	39	33	DR	17	14
ND	5	71	ND	12	90
(a) NDepend. Cohen	$\kappa = 0.48. N = 148$		(b) SonarQube. Cohen	$\kappa = 0.44. N = 133$	
	DR	ND		DR	ND
DR	32	24	DR	32	24
ND	35	94	ND	35	94
(c) CAST. Cohen $\kappa =$	$0.28, N = 185$				

TABLE IV
VALIDATOR AGREEMENT. DR=DESIGN RULE. ND=NOT DESIGN. GRAY CELLS SHOW AREAS OF RATER DISAGREEMENT.

	DR	ND		DR	ND
DR	14	2	DR	23	6
ND	28	30	ND	19	26
(a) Validator #2. Co-	$\kappa = 0.25. N =$		(b) Validator #3. Co-	$\kappa = 0.34, N =$	
74			74		
	DR	ND		DR	ND
DR	21	5	DR	21	27
ND	21	27	ND	21	27
(c) Validator #1. Co-	$\kappa = 0.32. N =$				
74					

to classify rules from the three tools, using the final version of the rubric. We compared each label to a reconciled set of labels that the authors created (where reconciled means we discussed each disagreement to derive a final consensus label). We report Cohen’s Kappa and confusion matrices, like that in Table III, in Table IV.

It is important to note that the validation data set was composed primarily of rules where labelers initially disagreed and then discussed during reconciliation. Our rationale for this criteria was to gather input on effectiveness of classification rubric improvements made after the initial round of classification. While this approach did yield important feedback, the side effect is that focusing on only the challenging labels may have contributed, at least in part, to low agreement numbers. Following their labeling exercise, we asked the participants for their impressions of the rubric.

- 1) **Scope.** Validators 1 and 3 both expressed difficulty in how to interpret the ‘scope’ filter for its design relevance. Validator 3 explained using an example rule: *Constructors of abstract classes should be declared as protected or private*. Although the method visibility aspect of this rule appeared to have design implications, Validator 3 struggled with the scope aspect explaining, “I usually defaulted to ND-1 since the change seemed to be limited to a specific statement. However, I wondered should I have said the scope was ‘system’ because of the potential impact across the whole system”.
- 2) **DR-P/DR-S overlap.** The validators commented on the overlap between the decision branches leading to DR-P and DR-S. For example, a violation that occurs at one location makes DR-P applicable, however, system boundary implications make DR-S also applicable. Validator 1 said the decision was further complicated by trouble deciphering what it meant for problems to cross system boundaries, “Is the problem located at multiple points in the system? Does it affect multiple points...?” (we noted the latter comment is similar to the overarching scope comment given by Validator 3 above).

- 3) **Metric threshold rules.** Metric threshold is about rules such as *Module complexity over x limit*. Validator 1 said that scope frequently led him down the ND-2 module path because heuristics are frequently applied at the method or class level in the rules. However, metric thresholds show up as an example in DR-S, too.

VII. DISCUSSION

Research to date on the fitness of quality analysis tools for design analysis has taken the approach of running multiple tools on the same data sets and comparing the results with each other [16], [17], or with some measure of design ground truth [18]. Such approaches have limitations due to potential feature limitations of tools. Consequently, they assess the tools and report what cannot be done, but do not improve our understanding of design analysis based on core concepts. Our goal in analyzing the rule sets is to assess the characteristics of automatable rules that check for design problems.

A. The Nature of Design Rules

Design rules have one or more of the following properties:

- Design rules check for propagation of issues across system elements. All of the rules where there was agreement on as design rules had this property.
- Design rules encapsulate semantics of known design constructs, such as architectural patterns. While some rules were initially tossed out as syntactic, recognizing that their goal was to enforce design constructs or patterns helped clarify their assessment.
- Design rules break down semantics of quality attribute concerns. This is an area where most significant progress can be made. Certain quality attribute concerns can still be automatable and enforce design; exception handling, test coverage were such examples in the rules we analyzed. We see the biggest gap in existing rules in taking explicit advantage of a mapping between automatable analysis and quality attribute concerns and design tactics.
- Design rules assess system impact over time, which is tricky as this is where there is no magic number or metric and is observed over time in the context of the system. We did not include reporting rules in this set such as source lines of code, number of functions changed and the like.

Exploring the research question: *What kinds of design-related rules do common tools provide?* gave input into the intersection of what the rule sets of common tools can detect, and what design analysis is needed to detect technical debt. We found several problems with the software quality rules the tools provided:

- Rules evolved from metrics that are possible to calculate and automate, but those metrics are not necessarily the most useful for the information needs of tool users.
- Rules that the user chooses not to take action on are due to false positives, lack of context integration, and overload of data without actionable guidance on what to fix remain a large problem.

- Many rules focus on analysis at the ‘code smells’ level and below (taxonomy levels 1-3). Problems at the paradigm suitability and design quality levels (4,5) remain poorly supported. However, this was the source of the greatest amount of technical debt according to our earlier studies [4], [5].

B. Lessons Learned in Assessing Rules

- *Are syntactic rules checking for design conformance?* Our rubric led us to classify rules checking purely syntactic implementation as not design related. However, the goal of some syntactic rules is to enforce design conformance. Examples of such rules are “avoid declaring an exception and not throwing it” or “classes should not be empty” indicating dead code in some cases.
- *Are metric threshold rules indicative of design problems, and thus design rules?* While a number of software metrics violating a certain threshold are available as rules (e.g., cyclomatic complexity should not exceed 7, depth of inheritance should not exceed 5), there is also evidence that such heuristics have a wide range of false positives and disregard context. These thresholds make sense only when combined and correlated with other system observations. These rules are helpful in creating other complex rules, yet are not useful for design assessment solely by themselves.
- *Are reporting rules design rules?* Reporting rules measure source lines of code, class size, method size, line length, number of parameters used, and the like. We concluded that they are not design rules, but need to be treated as contextual parameters that allow improved analysis of systems with relevant subsets of the rules.

C. Threats to Validity

We identified the following threats to validity.

Manual inspection: Our research approach relies on manually classifying existing software quality rules as implemented by commercial tools for their ability to detect design issues. In order to minimize biased classification, we had multiple researchers label the rules, and revised the classification guidance based on the biases we identified. In addition, we also had experts external to the research team classify a random sampling of the issues. The team and external experts classifying the rules are experts in technical debt research and software design, but they are not all experts in the tools or languages studied.

External validity: At this stage of our study we do not claim generalizable results, i.e., we do not expect this rubric, and its associated definition of design, to necessarily work on other tools. We continue to address these open research questions by applying the outcome rule set to several customer and open source projects, which is ongoing work that is not in the scope of this paper.

Internal validity: We checked our inter-rater reliability by having two of us assess rules, and rotating raters after round one. We resolved disagreements and reflected the outcomes

of that in the rubric. We only focus on a subset of rules for Java and C# and those that these particular tools have defined. However, the rules focus on basic language constructs that are transferable and we included rules from several tools to minimize this threat.

Construct Validity: The initial taxonomy relies on our initial view of software design to characterize software quality rules. Other definitions of software design (as opposed to code quality) might result in a different taxonomy. Our rubric relies on our understanding of design, as well as the literature and tool’s internal documentation.

VIII. RELATED WORK

A. Finding Design Violations

A few approaches, embody architectural design constraints as first class constructs. Koziol surveyed architecture metrics and categorized a number of quantitative methods [19]. Some studies have looked at validating a collection of measures in a quality model (mostly for maintainability [20]). Others have looked in depth at one or more methods to assess the efficacy of the approach (for example [12]). The Consortium for IT Quality (CISQ) Automated Quality Characteristic Measures are used to identify violations of good coding and architectural practice in the source code of software [21]. CAST has rules to implement the measures, for example, it allows one to define constraints on relations between modules, and highlights violations [22]. Lattix is specifically targeted to architectural analysis, using static dependency clusters as the unit of analysis [23]. Reflexion models were the initial approach to defining this [24]; since extended into architecture patterns [25]. The other approach is architecture reconstruction by detecting architectural tactics and design patterns [26] or design rule spaces [27].

A number of studies have looked for relationships between software metrics as reported by software quality tools as those we used for our study and technical debt. This work has applied existing code smells, coupling and cohesion, and dependency analysis to identifying areas of technical debt [28], [16]. The outcomes of these studies correlate with our findings where some subsets of the rules do represent more significant issues in the code base. There is also work that investigates the analysis of technical debt problem as a design flaw issues, such as [29]. Similarly, Kazman et al. relates architectural modularity violations to number of bugs to detect technical debt [30]. The goal of our study is to create a framework by which we can extend and limit the scope of existing rules to repeatedly detect design quality. The design issues highlighted in these studies overlap with those rules that our study identified as design rules and represent the modifiability/maintainability space.

Software quality tool vendors are grappling with the same challenges. Recently SonarQube simplified its quality model to create a better encapsulation of code issues, maintainability issues and run-time aspects that are most critical such as

vulnerabilities⁸. CISQ recently put out a developer survey to understand how developers perceive the time to fix when they encounter certain violations that such rules tag to better assess technical debt⁹.

B. Use of Tools

Johnson et al. [1] write about interviews they did at Google with twenty developers asking how they used Lint, Findbugs etc. The main findings were it took them out of their workflow and there were too many false positives. A similar perspective was voiced in our survey in terms of using such tools for technical debt analysis [4].

Emanuelsson and Nilsson [31] describe how three static analysis tools are used at Ericsson to find bugs: Coverity, Polyspace, and Klocwork. After describing the different features of each tool, they turn to experiences at Ericsson with the tools. Among their conclusions were that the tools generally performed quite well, with manageable false-positive rates, if the tool was introduced early on. Here “false-positive” is used in the *irrelevant* sense to capture the developer’s perspective: any report from the tool a user chooses not to take action to resolve the report, rather than the strict sense, namely, a result that should never have occurred. Sadowski et al. [7] call this ‘effective false positives’. Falessi and Vgele discuss this distinction as well [16]. The Ericsson result also found that tools were complementary to one another, but desired better tuning and configurability options, e.g., “a mechanism to fine-tune the effort spent on deriving value ranges of variables and the effort spent on aliasing analysis”.

In Smith et al. [32] and Witschey et al. [33] the authors investigate the usefulness and patterns of work for security analysis tools. Smith et al. examine how developers use security analysis tools. This builds on earlier work that examines information needs for software developers, such as Ko’s work [34] that looked at questions asked in program understanding tasks. Smith et al. took ten developers and assigned them security analysis tool to walk through bug reports, categorizing the types of questions they asked while doing so. Card sorting produced over ten classes of common questions new tools should seek to address. Witschey et al. did a similar exercise looking at what intrinsic properties of those tools were important for adoption. The differences between the two papers—focused on information needs developers have, and also support tools provide—mirrors the framework we adopt for technical debt tools.

For security tools, Witschey et al. [33] created a regression model to predict whether a tool would be adopted. The factors in that model include that one can see what others do with the tool (Observability), frequent organizational training and mandates (Education and Policies), and a perception that security tools are valuable (Advantages). It seems reasonable to think that most of these factors would generalize to other quality attributes and software quality in general.

⁸<https://blog.sonarsource.com/bugs-and-vulnerabilities-are-1st-class-citizens-in-sonarqube-quality-model-along-with-code-smells/>

⁹<http://it-cisq.org/technical-debt-remediation-survey/>

Smith et al. [32] sought to understand the information developers need when using a static security analysis tool. The outcome was a table of important information needs, including “Preventing and Understanding Potential Attacks”, and understanding “Control and Call Information”, among several others. The results are mostly specific to security, but a class of needs about developer self-awareness and education are likely useful for technical debt tools as well.

There are also several tools to assist with refactoring. What is relevant to the design analysis and technical debt tool landscape is that despite an established and industry-championed set of practices (“Refactoring”, by Martin Fowler [35], is widely considered a standard book in the practice of software engineering), excellent IDE integration, many refactoring practices are poorly adopted. This puzzling state of affairs is investigated by Murphy-Hill and Black [36] and in more detail by Murphy-Hill, Parnin, and Black [37]. They note that close to 90 percent of refactorings are performed manually, despite good tool support. Rename was by far the most common.

In analyzing the tools themselves, Murphy-Hill et al. [37] notes that “Toolsmiths need to explore alternative interfaces and identify common refactoring workflows” to better integrate the refactoring with what developers need at that moment. They list three key factors in whether a developer will use the tool: *Awareness* of the tool, *Opportunity* to apply the tool at the right moment, and *Trust* that the tool will do what is needed. The authors also note two reasons developers gave for not using a tool: 1) unfamiliar touch points in the code, those places the refactoring would change; 2) disruption in normal routines by jumping into new UI elements and out of the problem at hand. Both the three factors and two obstacles seem relevant to technical debt tools as well.

There is a renewed interest in software analytics, and many corporations are adopting them as part of their practice. Sadowski et al. [7] describe how Google integrate static analysis and technical debt identification into Google’s development practices and environments, using a tool called Tricorder (now released as an external open source project called Shipshape). They begin by noting the problems with existing tools: Google’s giant codebase; engineers testing their own code and mandatory code reviews for each patch; high false-positive rates with tools in the past, such as Findbugs; and a wide variety of programming languages and contexts. The Tricorder tool addresses this last point by supporting domain-specific analysers using a contract approach, that mandates these plugins not have high false positive rates (that is, report results that developers label annoying or uninformative).

IX. CONCLUSION

Software quality tools mix design rules and code quality rules. Separating these is important since design problems often result in longer-term costs. This study labeled the software quality rules of three tools as *design* or *not-design*, using an iteratively defined classification rubric validated with experts.

Our study suggests that progress in automated design analysis can be achieved by addressing the following:

- 1) **Defining design rule scope.** Our classification results revealed design rules go beyond statement level quality checks. Tools that reported scope of impact (and not just time to fix, as many do currently), would aid in classification.
- 2) **Validating properties of design rules.** Based on the design rules we identified, we extracted initial properties of design rules. For example, we found that the quality attribute filter was not helpful since nearly every rule, even syntax rules, had some quality impact. Existing know-how on design tactics and expert observations can help validate and improve the properties of automatable design rules.
- 3) **Improving context sensitivity.** The hard to label rules in our data are an important outcome of our study. Often, the reason they were hard to correctly label was due to context-specificity. As reported by Microsoft researchers in [38], it is only when your analytics efforts work closely with the information needs of the stakeholders that there is real impact.

X. ACKNOWLEDGMENTS

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

[Distribution Statement A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

DM-0004376

REFERENCES

- [1] B. Johnson, Y. Song, E. Murphy-Hill, and R. Bowdidge, “Why don’t software developers use static analysis tools to find bugs?” in *Proceedings of the ACM/IEEE International Conference on Software Engineering*, San Francisco, CA, USA, 2013, pp. 672–681.
- [2] S. Johnson, “Lint, a C program checker,” Bell Labs, Tech. Rep. 65, 1979.
- [3] M. Aniche, G. Bavota, C. Treude, A. van Deursen, and M. A. Gerosa, “A validated set of smells in model-view-controller architecture,” in *Proceedings of the IEEE International Conference on Software Maintenance and Evolution*, 2016.
- [4] N. A. Ernst, S. Bellomo, I. Ozkaya, R. L. Nord, and I. Gorton, “Measure it? manage it? ignore it? software practitioners and technical debt,” in *Proceedings of the ACM SIGSOFT International Symposium on Foundations of Software Engineering*, 2015, pp. 50–60.
- [5] S. Bellomo, R. L. Nord, I. Ozkaya, and M. Popeck, “Got technical debt?: surfacing elusive technical debt in issue trackers,” in *Proceedings of the International Working Conference on Mining Software Repositories*, 2016, pp. 327–338.
- [6] A. Bessey, K. Block, B. Chelf, A. Chou, B. Fulton, S. Hallem, C. Henri-Gros, A. Kamsky, S. McPeak, and D. Engler, “A few billion lines of code later: Using static analysis to find bugs in the real world,” *Communications of the ACM*, vol. 53, no. 2, pp. 66–75, Feb. 2010.

- [7] C. Sadowski, J. van Gogh, C. Jaspan, E. Soederberg, and C. Winter, "Tricorder: Building a program analysis ecosystem," in *Proceedings of the ACM/IEEE International Conference on Software Engineering*, 2015.
- [8] R. Kazman and H. Cervantes, *Designing Software Architectures: A Practical Approach*, ser. SEI Series in Software Engineering. Addison-Wesley, 2016.
- [9] C. Menzel, "Possible worlds," in *The Stanford Encyclopedia of Philosophy*, winter 2016 ed., E. N. Zalta, Ed. Metaphysics Research Lab, Stanford University, 2016.
- [10] J. Cohen, "Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit." *Psychological Bulletin*, vol. 70, no. 4, p. 213, 1968.
- [11] C. Bird, T. Menzies, and T. Zimmermann, *The Art and Science of Analyzing Software Data*. Elsevier, 2015.
- [12] E. Bouwers, J. P. Correia, A. van Deursen, and J. Visser, "Quantifying the Analyzability of Software Architectures," in *Proceedings of the IEEE/IFIP Working Conference on Software Architecture*, Boulder, CO, Jun. 2011, pp. 83–92.
- [13] K. Sullivan, P. Chalasani, and S. Jha, "Software Design as an Investment Activity: A Real Options Perspective," K. Sullivan, P. Chalasani, and S. Jha, Eds. *Real Options and Business Strategy: Applications to Decision Making*, Dec. 1999, pp. 215–262.
- [14] L. Xiao, Y. Cai, and R. Kazman, "Titan: a toolset that connects software architecture with quality analysis," in *International Symposium on Foundations of Software Engineering*, 2014, pp. 763–766.
- [15] K. Power, "Understanding the impact of technical debt on the capacity and velocity of teams and organizations: viewing team and organization capacity as a portfolio of real options," in *International Workshop on Managing Technical Debt (MTD)*. IEEE, 2013, pp. 28–31.
- [16] D. Falessi and A. Voegelé, "Validating and prioritizing quality rules for managing technical debt: An industrial case study," in *International Workshop on Managing Technical Debt*, 2015.
- [17] I. Griffith, D. Reimann, C. Izurieta, Z. Codabux, A. Deo, and B. J. Williams, "The correspondence between software quality models and technical debt estimation approaches," in *Sixth International Workshop on Managing Technical Debt, MTD@ICSME 2014, Victoria, BC, Canada, September 30, 2014*, 2014, pp. 19–26.
- [18] T. Lutellier, D. Chollak, J. Garcia, L. Tan, D. Rayside, N. Medvidovic, and R. Kroeger, "Comparing software architecture recovery techniques using accurate dependencies," in *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 2*, 2015, pp. 69–78.
- [19] H. Koziolok, "Sustainability evaluation of software architectures: A systematic review," in *Proceedings of the Joint ACM SIGSOFT Conference – QoSA and ACM SIGSOFT Symposium – ISARCS on Quality of Software Architectures – QoSA and Architecting Critical Systems – ISARCS*. New York, NY, USA: ACM, 2011, pp. 3–12.
- [20] R. Plöesch, S. Schuerz, and C. Koerner, "On the validity of the it-cisq quality model for automatic measurement of maintainability," in *Annual Computer Software and Applications Conference*, vol. 2, July 2015, pp. 326–334.
- [21] Consortium for IT Software Quality, "CISQ specifications for automated quality characteristic measures," Consortium for IT Software Quality, Technical Report CISQ–TR–2012–01, 2012.
- [22] A. Szykarski, J. Sappidi, and B. Curtis, "Estimating the principal of an application's technical debt," *IEEE Software*, vol. 29, no. undefined, pp. 34–42, 2012.
- [23] N. Sangal, E. Jordan, V. Sinha, and D. Jackson, "Using dependency models to manage complex software architecture," in *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, 2005, pp. 167–176.
- [24] G. C. Murphy, D. Notkin, and K. Sullivan, "Software reflexion models: Bridging the gap between source and high-level models," in *International Symposium on Foundations of Software Engineering*, 1995, pp. 18–28.
- [25] M. Anvaari and O. Zimmermann, "Towards reusing architectural knowledge as design guides," in *Proceedings of the International Conference on Software Engineering and Knowledge Engineering*, 2013, pp. 181–186.
- [26] M. Mirakhorli, P. Mäder, and J. Cleland-Huang, "Variability points and design pattern usage in architectural tactics," in *International Symposium on the Foundations of Software Engineering*, Cary, NC, November 2012.
- [27] L. Xiao, Y. Cai, and R. Kazman, "Design rule spaces: A new form of architecture insight," in *Proceedings of the ACM/IEEE International Conference on Software Engineering*, 2014, pp. 967–977.
- [28] F. A. Fontana, V. Ferme, and S. Spinelli, "Investigating the impact of code smells debt on quality code evaluation," in *Proceedings of the Third International Workshop on Managing Technical Debt*, 2012, pp. 15–22.
- [29] R. Marinescu, "Assessing technical debt by identifying design flaws in software systems," *IBM Journal of Research and Development*, vol. 56, no. 5, p. 9, 2012.
- [30] R. Kazman, Y. Cai, R. Mo, Q. Feng, L. Xiao, S. Haziyeve, V. Fedak, and A. Shapochka, "A case study in locating the architectural roots of technical debt," in *Proceedings of the ACM/IEEE International Conference on Software Engineering*, 2015, pp. 179–188.
- [31] P. Emanuelsson and U. Nilsson, "A comparative study of industrial static analysis tools," *Electronic Notes in Theoretical Computer Science*, vol. 217, pp. 5 – 21, 2008, proceedings of the 3rd International Workshop on Systems Software Verification (SSV 2008).
- [32] J. Smith, B. Johnson, E. Murphy-Hill, B. Chu, and H. R. Lipford., "Questions developers ask while diagnosing potential security vulnerabilities with static analysis," in *International Symposium on Foundations of Software Engineering*, Bergamo, Italy, 2015.
- [33] J. Witschey, O. Zielinska, A. K. Welk, E. R. Murphy-Hill, C. B. Mayhorn, and T. Zimmermann, "Quantifying developers' adoption of security tools," in *International Symposium on Foundations of Software Engineering*, Bergamo, Italy, 2015, pp. 260–271.
- [34] A. J. Ko, H. H. Aung, and B. A. Myers, "Eliciting Design Requirements for Maintenance-Oriented IDEs: A Detailed Study of Corrective and Perfective Maintenance Tasks," in *Proceedings of the ACM/IEEE International Conference on Software Engineering*, St Louis, MS, 2005, pp. 126–135.
- [35] M. Fowler, K. Beck, J. Brant, W. Opdyke, and D. Roberts, *Refactoring: Improving the Design of Existing Code*. Addison Wesley, 1999.
- [36] E. Murphy-Hill and A. Black, "Refactoring tools: Fitness for purpose," *IEEE Software*, vol. 25, no. 5, pp. 38–44, Sept 2008.
- [37] E. Murphy-Hill, C. Parnin, and A. Black, "How we refactor, and how we know it," *IEEE Transactions on Software Engineering*, vol. 38, no. 1, pp. 5–18, Jan 2012.
- [38] D. Zhang, S. Han, Y. Dang, J.-G. Lou, H. Zhang, and T. Xie, "Software analytics in practice," *IEEE Software*, vol. 30, no. 5, pp. 30–37, 2013.