



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *Fourth International Workshop on Artificial Intelligence for Requirements Engineering (AIRE'17), Lisboa.*

Citation for the original published paper:

Femmer, H., Unterkalmsteiner, M., Gorschek, T. (2017)
Which requirements artifact quality defects are automatically detectable?: A case study
In: *Proceedings - 2017 IEEE 25th International Requirements Engineering Conference Workshops, REW 2017*, 8054884 (pp. 400-406). IEEE
<https://doi.org/10.1109/REW.2017.18>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:bth-15453>

Which requirements artifact quality defects are automatically detectable? A case study

Henning Femmer
Institut für Informatik
Technische Universität München, Germany
femmer@in.tum.de

Michael Unterkalmsteiner, Tony Gorschek
Software Engineering Research Lab,
Blekinge Institute of Technology, Sweden
{mun,tgo}@bth.se

Abstract—[Context:] The quality of requirements engineering artifacts, e.g. requirements specifications, is acknowledged to be an important success factor for projects. Therefore, many companies spend significant amounts of money to control the quality of their RE artifacts. To reduce spending and improve the RE artifact quality, methods were proposed that combine manual quality control, i.e. reviews, with automated approaches. [Problem:] So far, we have seen various approaches to automatically detect certain aspects in RE artifacts. However, we still lack an overview what can and cannot be automatically detected. [Approach:] Starting from an industry guideline for RE artifacts, we classify 166 existing rules for RE artifacts along various categories to discuss the share and the characteristics of those rules that can be automated. For those rules, that cannot be automated, we discuss the main reasons. [Contribution:] We estimate that 53% of the 166 rules can be checked automatically either perfectly or with a good heuristic. Most rules need only simple techniques for checking. The main reason why some rules resist automation is due to imprecise definition. [Impact:] By giving first estimates and analyses of automatically detectable and not automatically detectable rule violations, we aim to provide an overview of the potential of automated methods in requirements quality control.

Index Terms—Requirement Engineering, Artifact Quality, Automated Methods

I. INTRODUCTION

Requirements Engineering (RE) artifacts play a central role in many systems and software engineering projects. Due to that central role, the quality of RE artifacts is widely considered a success factor, both in academia, e.g. by Boehm [1] or Lawrence [2], and also by practitioners [3].

As a result, companies invest heavily into quality control of RE artifacts. Since RE artifacts are written mostly in natural language [4], quality control is usually applied manually, e.g. in the form of manual reviews. However, besides all of its advantages, manual quality control is slow, expensive and inconsistent, heavily dependent on the competence of the reviewer. One obvious approach to address this is combining manual reviews with automated approaches. The goal of a so-called *phased inspection* [5], [6] is to reduce the effort in manual reviews and to improve the review results by starting into the review with a better (e.g. readable) artifact.

Therefore, various authors have focused on automatically detecting quality defects, such as ambiguous language (i.a. [7], [8], [9], [10]) or cloning [11]. However, it is still an open

question to what degree quality defects can be detected automatically or require human expertise (i.e. manual work). In previous work [10], we took a bottom-up perspective by qualitatively analyzing which of the quality review results could be automatically detected.

Research Goal: In this work, we take a top-down perspective by focusing on requirements writing guidelines from a large company. Furthermore, we systematically classify and quantify which proportion of the rules can be automated.

II. RELATED WORK

Researchers and practitioners have been working on supporting quality assurance with automated methods (at least since the end of the 1990's [7]). We want to give only a brief, non-exhaustive summary here. Please refer to our previous work [10] for a more detailed analysis.

Defect types: Most works in this area focus on the detection of various forms of ambiguity, e.g. [8], [12], [13], [14]. Other works try to detect violations of syntactic [11] or even semantic duplications [15]. Other works focus on correct classifications [16] or on the question whether an instance follows given structural guidelines, e.g. for user stories [9] or for use cases [17].

Criteria: The aforementioned works used different sets of criteria. Most prominently are definitions of ambiguity [18], previously summarized lists of criteria [19], or requirements standards [10], [20].

Techniques: So far, various techniques have been applied, including machine learning [16], [21] and ontologies [22]. However, Arendse and Lucassen [23] hypothesize that we might not need sophisticated methods for most aspects of quality. In this paper, we provide data regarding this hypothesis. All in all, few works have tried to take a different viewpoint and understand what *cannot* be automatically checked. In previous work [10], we approached this question in a qualitative manner, by looking not at definitions, but at instances of defects. We did not quantify the portion of automatically discoverable defects, since this depends heavily on the requirements at hand (which defects does an author introduce and a reviewer find?).

Research Gap: Various authors have shown how to automatically detect individual quality defects. In previous work [10] we qualitatively analyzed which requirements quality defects can be detected. In this work, we provide first

evidence, based on requirements writing rules used in a large organization, on the proportion between automatically / not automatically detectable requirements quality issues.

III. STUDY DESIGN

We conducted this study in a research collaboration with the Swedish Transport Administration (STA), the government agency responsible for planning, implementing and maintaining long-term rail, road, shipping and aviation infrastructure in Sweden. In particular, we studied their requirements guidelines that were developed by editors who review and quality assure specifications. A total of 129 rules were analyzed in this paper. While our long-term goal in this research collaboration, is described in more detail elsewhere [24], the specific research goal of this paper is to *characterize requirements writing rules with respect to their potential to be automatically checked from the viewpoint of a requirements quality researcher in the context of an industrial requirements quality control process*. From this goal definition we derive our research questions:

RQ1: How many rules for natural language requirements specifications can be automated?

RQ2: To what degree can rules be categorized into groups and to what degree can these groups be eligible for automation?

RQ3: What information is required to automatically detect rule violations?

RQ4: Which rules resist automation and why?

A. Rule classification

A lack of classification schema for requirements writing rules prompted us to formulate the following schema (see Tbl. I).

1) *Rule type*: We distinguish between the lexical, grammatical, structural and semantic rule type (see rules 160, 56, 78 and 81 in Tbl. I). A lexical rule refers to constraints on the use of certain terms or expressions that may induce ambiguity, reduce understandability or readability. Similarly, a grammatical rule refers to constraints on sentence composition. A structural rule refers to the form in which information is presented and formatted. Finally, a semantic rule refers to constraints on the text content and meaning.

2) *Rule context*: We introduced this dimension to characterize in which context of the requirements specification the rule is relevant. An appropriate automated check flags only violations that occur in the correct context, e.g. in requirements (if they are separated from informative text), figures, tables, references, headings, enumerations, comments.

3) *Information scope*: This dimension describes the scope that needs to be considered in order to decide whether the rule is violated or not. We defined five levels: word/phrase, sentence, section, document and global. For example, to check rule 56 in Tbl. I, it is enough to inspect a sentence. However, rule 24 requires access to information that is not in the requirements specification, hence we classified it as global information scope. This characterization provides indication that can be used to estimate the relative required effort to implement the automated check of the rule.

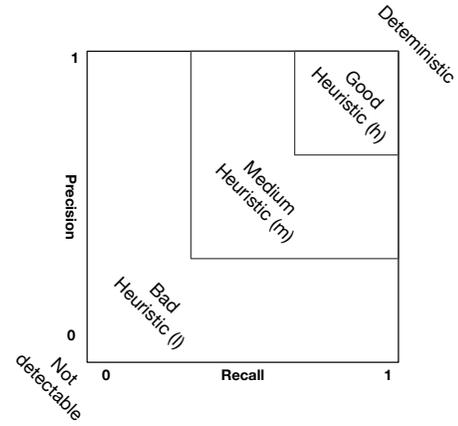


Fig. 1. The categories of detection accuracy as used in this study

4) *Necessary information*: This dimension describes NLP-based and domain-specific information needed to detect rule violations. NLP-based information refers to language and document structure, such as Part-of-Speech (POS) tags, lemmas and word stems, morphological tags, parse trees and meta-data on formatting. Domain-specific information is only available in the specific domain in which the rules apply, e.g. lists of referenced documents or a domain model / ontology. For example, rule 50 in Tbl. I can be decided with POS tags while rule 56 requires a parse tree that indicates where the subject is positioned in the sentence.

5) *Detection accuracy*: This dimension provides a rough estimate, based on the experiences of previous work [20], on the expected accuracy for detecting rule violations. We have defined a five-level scale, illustrated in Fig. 1, spanning from deterministic, i.e. 100% detectable, to not detectable at all. Good heuristics feature both high recall and precision, while bad heuristics always trade-off between precision and recall. For example, while assigning POS tags is a probabilistic algorithm, we classified rule 50 in Tbl. I as a good heuristic since this particular problem has been solved before, with demonstrably high precision and recall. We classified rule 81, on the other hand, as bad heuristic since, while conceptually feasible, we lack an accurate solution, i.e. a technique to extract a domain model and use that to determine whether a requirement statement contains supplemental information. Then, there are also rules that we do not expect to be automatically detectable at all (rule 54), because they turn out to be challenging, even in manual reviews. We classified these not automatically detectable rules along main reasons (categories resulted from previous work [10], see Tbl. III).

B. Data Collection, Classification and Analysis

We received a total of 192 writing rules from STA, of which we filtered unapproved rule ideas (63), resulting in 129 original rules. In case a rule contained discernible sub-rules, we split them up to facilitate the classification, resulting in 166 classified rules. We then developed an initial version of the classification schema illustrated in Section III-A. While all

TABLE I
CLASSIFICATION SCHEMA WITH RULE EXAMPLES

ID	Rule	Type	Context	Scope	Necessary information	Detection accuracy
160	The term “function” shall be used instead of the term “functionality”.	Lexical	Anywhere	Word/Phrase	Lemma / Dictionary	Deterministic
56	Requirements shall start with the subject.	Grammatical	Requirement	Sentence	Parse tree	Heuristic (h)
78	Text consisting of a definition shall be preceded with the identifier “Definition:”.	Structural	Requirement	Section	Lemma / Dictionary	Heuristic (m)
81	If a functional requirement is supplemented with additional information to clarify how the requirement can be met, the additional information must be formulated as a separate requirement.	Semantic	Requirement	Section	Domain model	Heuristic (l)
24	References to other documents in the specification are done by reference to the document title.	Structural	Anywhere	Global	Regular expressions, Document list	Deterministic
50	Requirements must be understandable independently, i.e. the subject must be indicated in the respective requirements (the subject must not be only defined in the section title).	Semantic	Requirement	Sentence	POS tags	Heuristic (h)
54	The introductory section of the specification shall not contain any requirements.	-	-	-	-	Not detectable

dimensions and the categories for type and detection accuracy were defined a-priori, the categories for context, scope and necessary information were identified during the classification process. During this first workshop we classified 39 rules, stabilizing the schema and fostering our shared understanding. Then, the second author proceeded to classify the remaining 127 rules alone. The first author sampled 20 rules from this set, independently classified them and calculated the inter-rater agreement ($\kappa = 0.79$) which is considered substantial [25]. The first author then reviewed all 127 rules, marked those where he disagreed, and finally consolidated all classifications with the second author in a second workshop.

We then used the classifications of accuracy for RQ1, the type, context and scope for RQ2, the necessary information for RQ3, and the reasons for RQ4.

IV. RESULTS

RQ1: How many rules for natural language requirements specifications can be automated?

In Fig. 2, we show the results from classifying the estimated detection accuracy of the rules. We estimate that 41% of the rules can be deterministically checked, meaning that an algorithm finds each violation. 34% of the rules are heuristic, with 12% of high accuracy, and 11% of medium and low accuracy. We estimate that the remaining 25% cannot be checked at the current state of art and at the current state of the rule definitions.

Discussion: Whether rules can be automatically detected is not a binary question. In fact, it depends on the context. However, most rules we can put into a certain category, indicating their potential to be automatically checked. We were surprised by the large number of rules that can be automated. This indicates the potential for automation, as we will discuss in future work.

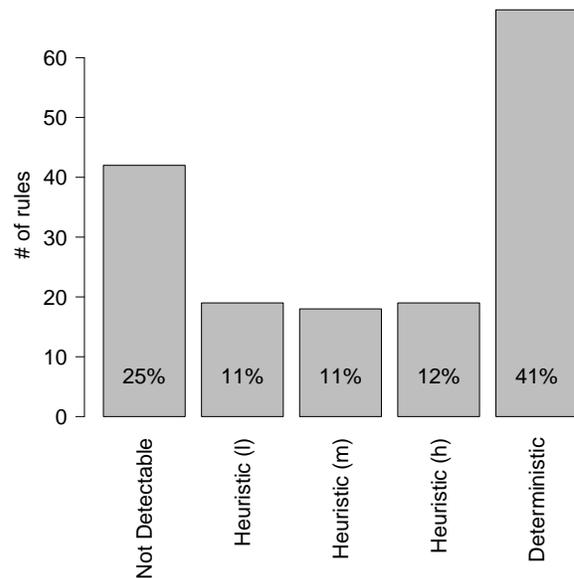


Fig. 2. Frequency of rules falling into one of the detection accuracy categories.

RQ2: To what degree can rules be categorized into groups and to what degree can these groups be eligible for automation?

In Fig. 3, we show the results from classifying the automatically detectable rules by their type and estimated detection accuracy. The results indicate an estimated high detection accuracy for structural and lexical rules, medium accuracy for grammatical rules, and medium to low accuracy for semantic rules. Fig. 4 shows that most rules are at the level of words or phrasing or at the level of sentences. Lastly, Fig. 5 shows that most rules hold anywhere or specifically concern the requirements of the RE artifact.

Discussion: The further a rule goes into semantic aspects, the harder it is to detect violations. For structural rules, e.g. where

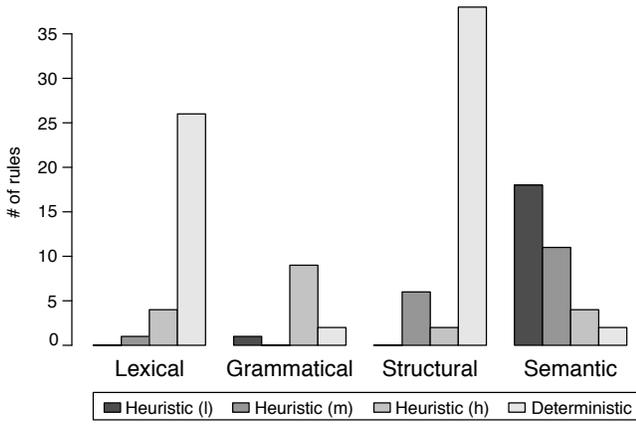


Fig. 3. Estimated detection accuracy for each category.

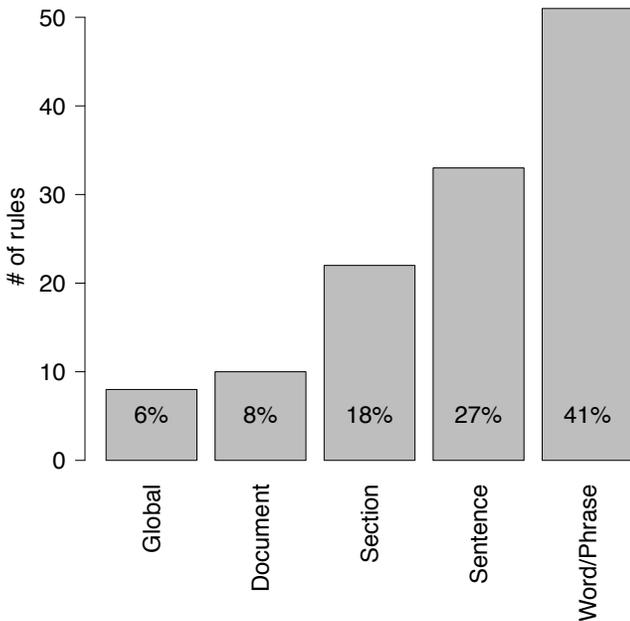


Fig. 4. Distribution of the scope of the automatically detectable rules.

a certain piece of information should be placed, there are a few rules for which violations are difficult to check automatically. For example, to understand whether a certain text should be tagged as a requirement requires context understanding. We describe further reasons for rules not being automatically detectable in RQ4.

RQ3: What information is required to automatically detect rule violations?

To understand what techniques are required to automatically detect violations of guideline rules, we classified each rule with the required information for this rule. Each required information then leads to a certain technique. For example, if the lemmas of the words are required, we obviously need a lemmatization technique. Tbl. II shows the results for this analysis. The three most common techniques are the following: In 47% of the cases, lemmatization is required to detect a

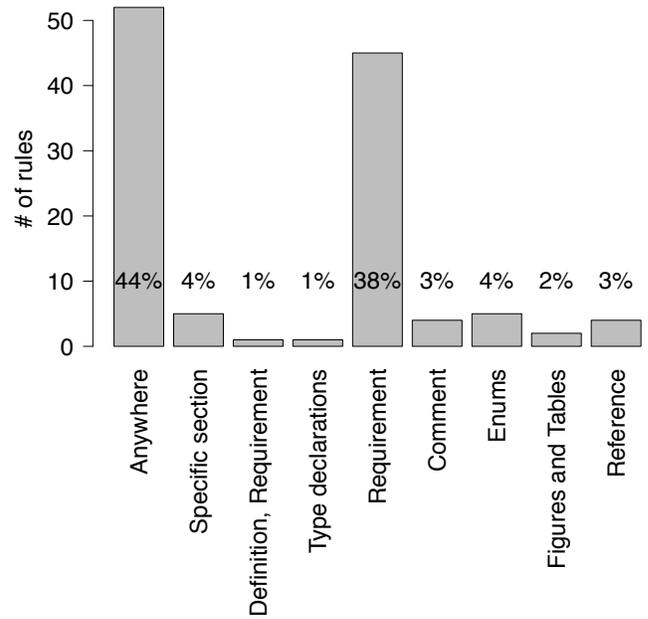


Fig. 5. Context of the automatically detectable rules.

violation of a rule. In a further 35% of cases only the pure text and regular expressions are needed. Next, formatting information is required in 22% of the cases.

Discussion: This analysis supports the hypothesis of Arendse and Lucassen [23] that in most cases, we do not need sophisticated methods to detect violations of rules.

TABLE II
FREQUENCY OF REQUIRED INFORMATION (MULTIPLE SELECTIONS)

Information	Occurrences	Share of Rules
Lemmas / Dictionaries	58	47 %
Pure Text (Reg. Expression)	43	35 %
Formatting	27	22 %
Domain Models	11	9 %
Part of Speech Tags	11	9 %
Lists of [X]	8	6 %
Morphology	5	4 %
Parse Trees	3	2 %
Word Stems	3	2 %
Tokens / Sentences	3	2 %
Named Entities	1	1 %

RQ4: Which rules resist automation and why?

When analyzing the *not automatically detectable* rules of RQ1, the reasons were distributed as shown in Tbl. III (classification extends previous work [10]). The major reason was, in our studied case, that the rules themselves are still imprecise or unclear. Examples for this are rules such as "Requirements must be accurate, unambiguous, comprehensive, consistent, modifiable, traceable." (this was one single rule) or "Requirements should contain enough information." These rules cannot be checked either manually or automatically. One could even argue that they convey little value. Such imprecise or unclear

rules are the reason for 81% of the not automatically detectable rules (see Tbl. III). In 12% of the cases, an automation would need profound domain knowledge to automatically detect a violation. An example is that requirements about certain system parts must first state that these parts exist. However, to understand which parts this refers to, we would need to know the domain. This means that only domain experts can manually detect violations to these rules. In one case, respectively, the rule requires deep semantic understanding of the text (e.g. to detect logical contradictions written in natural language in different paragraphs), the system or even the process scope.

TABLE III
SHARE OF REASONS THAT PREVENT AUTOMATED DETECTION

Reason	Frequency	Share
R_1 : Rule unclear or imprecise	34	81 %
R_2 : Deep semantic text understanding	1	2 %
R_3 : Profound domain knowledge	5	12 %
R_4 : System scope knowledge	1	2 %
R_5 : Process status knowledge	1	2 %
Sum	42	100 %

Discussion: Deep computational problems do not seem to be the major cause for why we see no chance in checking a certain rule, rather imprecise rules themselves.

V. DISCUSSION

A. Share of automatically detectable defects

In our study, we found that a substantial number of requirements writing rules can be automatically checked. This is a top-down perspective and as such helps to quantify the share of defects that can be automatically detected. However, this does not necessarily transfer to the share of defects found in reviews. This is for the following reasons: First, defects created by requirements engineers are not equally distributed over the guideline rules. Furthermore, the defects introduced by requirements engineers very much depend on the individual person, company, and project. Second, defects discovered by reviewers are not necessarily equally distributed over the guideline rules. Therefore, we argue to consider both perspectives, i.e. the share of defects based on guidelines and the share of defects existing in practice, when discussing the potential of automated requirements quality assurance.

B. The 100%-Recall Argument

There is an ongoing debate in the scientific community whether automated checks in quality assurance need 100% recall to be useful in practice. Some authors (i.a. [26], [27], [28]) argue that if an approach does not achieve perfect recall, this leads to either the reviewer does not check the rule anymore, which would lead to unchecked defects, or the reviewer has to go through the whole document anyways, and thus, the automated analysis has no benefits. We disagree with this view for two reasons. First, we argue that in industrial practice, reviewers rarely go through the artifact rule by rule.

Therefore, there is no such thing as *omitting a certain rule*. Reviewers see the guidelines rather as a supporting instrument, and thus anything that reminds them of certain rules, increases the quality. Our second argument also refers to the status quo today. The best automated quality support that is widely used are spell and grammar checks. Both do not have 100% recall. So, if recall is a problem, why do we use spell and grammar checks every day? In our experience from introducing automated analyses at various companies in industry, practitioners were more worried about precision than recall. They are convinced of the value ("*Anything helps!*"), and care more for acceptance with the end users. Here, the core aspect is usability in the form of few false positives, ergo: precision (cf. also similar discussing in static code analysis [29]).

C. Threats to Validity

There are two major threats to validity. Regarding internal validity, we classified the rules according detection accuracy. We did so because it was not feasible within the scope of this work to do a precision & recall analysis for each guideline rule. However, the first author has been translating guideline rules into automated analyses for 4 years. Thus, we are confident that the results reflect the real precision and recall after implementation. In addition, we created rough categories to gain an overview, not a precise analysis for each rule. To evaluate this aspect, we independently classified a subset of 10% of the rules and calculated a weighted Cohen's kappa of the resulting classification ($\kappa = 0.79$). This agreement fosters our confidence in the resulting classification.

The second threat relates to external validity. Since we analyzed a large guideline used at STA, we do not know whether the results generalize from this partner. We have, however, previously informally checked a guideline from another industry partner in a different domain. Here we came to the same share of not automatically detectable rules (25%). Future work should broaden the study to different guidelines.

VI. RESEARCH AGENDA

The current paper provides an estimation of the extent to which industrial requirements quality rules can be automatically checked. We plan to continue our research as follows.

Complete the rule classification. 34 of the studied rules were imprecise or unclear. Unfortunately, the authors of the writing guidelines were not available for feedback during the course of this study. We want to deepen our understanding on the nature of the imprecision of these rules. In addition, we had no access regarding the relevance, value, and frequency of violations of the rules. This could provide insights how rules that can be automatically checked potentially contribute to review effort reduction. In addition, the classification scheme used in RQ2 was beneficial for this study and worked fine regarding the first three categories (lexical, grammatical, structural). However, the scheme created some discussion around the semantic category. The reason is that most rules intertwine semantic and syntactic aspects: Since requirements artifacts are not automatically compiled like code, the point of syntactic rules is

only to prevent semantic issues. Therefore, future work should extend this classification scheme to clarify this aspect, e.g. by decoupling the two aspects.

Implement and statically validate rules. We have already begun to implement some rules that are based on dictionary lookups using an existing requirements smell detection framework [10]. While most of the rules can be implemented with simple techniques, we also plan to experiment with more advanced NLP techniques where we expect challenges in the detection accuracy. For example, violations to rule 81 in Table I could be detected by using topic models enhanced with domain knowledge [30]: requirements that contain distant topics or several closely related topics indicate candidates for rule violations. To validate the implemented rules, we can exploit the fact that at STA, the rules were developed based on experience, i.e. there exist versions of requirements that contain rule violations. We can fine-tune and validate the detection against this set. We also plan to provide an analysis of the potential benefits of using automated requirements quality control. To achieve this, we analyze historic requirements (where the current rules were not applied) and study the effort spent on discussing and repairing these violations.

Validation in Use. We plan to evaluate the efficiency and effectiveness of automated requirements quality assurance in use, i.e. in the environment of STA with the support of their requirements editors. One important question to answer is whether we can control the number of false positives, a crucial aspect for the adoption of tool support in industry that has also been observed in other areas, such as static bug detection [29].

Repository for requirements writing rules. Finally, we, as a community, should establish a repository of precise general and validated requirements rules. Such a repository can be created by replicating the work proposed in this paper in different contexts and, at the same time, advance the techniques for detecting rule violations.

VII. CONCLUSIONS

It is unclear what proportion of quality defects can be automatically detected. Therefore, in this work, we classify rules from a large, fine-grained requirements writing guideline from one of our industry partners. The results indicate that a surprisingly large proportion of rules (41%) can be automatically analyzed. 53% can be analyzed deterministically or with a good heuristic. One reason for this was that these rules contain many structural rules, which require just an analysis of formatting information or pure text. If we take also those rules into account where we have a medium heuristic, we could even tackle 64% of the rules. However, our analysis also shows that 36% of the rules have no or little chance to be automated. While being just first evidence, this analysis indicates that there is a substantial proportion of guideline rules (our intermediate for quality defects) that can be automatically checked. However, the analysis also indicates that there is little hope that we can completely replace manual reviewing with automated reviews. Combining automated and

manual quality assurance, as proposed by others [5], and also ourselves [6] could be the promising compromise.

ACKNOWLEDGEMENTS

This work was performed within the project Q-Effekt and ERSK; it was funded by the German Federal Ministry of Education and Research (BMBF) under grant no. 01IS15003 A-B and by the Swedish Transport Administration. The authors assume responsibility for the content. The authors thank Jonas Eckhardt for comments on an earlier draft of this paper.

REFERENCES

- [1] B. W. Boehm and P. N. Papaccio, "Understanding and controlling software costs," *IEEE Transactions on Software Engineering*, vol. 14, no. 10, pp. 1462–1477, 1988.
- [2] B. Lawrence, K. Wiegers, and C. Ebert, "The top risks of requirements engineering," *IEEE Software*, pp. 62–63, 2001.
- [3] D. Méndez Fernández and S. Wagner, "Naming the pain in requirements engineering: A design for a global family of surveys and first results from germany," *Information and Software Technology*, vol. 57, pp. 616–643, 2015.
- [4] L. Mich, F. Mariangela, and P. L. Novi Inverardi, "Market research for requirements analysis using linguistic tools," *Requirements Engineering Journal*, vol. 9, no. 1, pp. 40–56, 2004.
- [5] J. C. Knight and E. A. Myers, "An improved inspection technique," *Communications of the ACM*, vol. 36, no. 11, pp. 51–61, 1993.
- [6] H. Femmer, B. Hauptmann, S. Eder, and D. Moser, "Quality assurance of requirements artifacts in practice: A case study and a process proposal," in *PROFES*, 2016, pp. 506–516.
- [7] W. M. Wilson, L. H. Rosenberg, and L. E. Hyatt, "Automated analysis of requirement specifications," in *ICSE*, 1997, pp. 161–171.
- [8] F. Fabbrini, M. Fusani, S. Gnesi, and G. Lami, "An automatic quality evaluation for natural language requirements," in *REFSQ*, 2001.
- [9] G. Lucassen, F. Dalpiaz, J. M. E. van der Werf, and S. Brinkkemper, "Improving agile requirements: the quality user story framework and tool," *Requirements Engineering*, vol. 21, no. 3, pp. 383–403, 2016.
- [10] H. Femmer, D. Méndez Fernández, S. Wagner, and S. Eder, "Rapid quality assurance with requirements smells," *Journal of Systems and Software*, vol. 123, pp. 190–213, 2017.
- [11] E. Juergens, F. Deissenboeck, M. Feilkas, B. Hummel, B. Schaez, S. Wagner, C. Domann, and J. Streit, "Can Clone Detection Support Quality Assessments of Requirements Specifications?" in *ICSE*, 2010.
- [12] A. Fantechi, S. Gnesi, G. Lami, and A. Maccari, "Application of linguistic techniques for Use Case analysis," *Requirements Engineering*, vol. 8, no. 3, pp. 161–170, 2002.
- [13] E. Knauss, D. Lübke, and S. Meyer, "Feedback-Driven Requirements Engineering : The Heuristic Requirements Assistant," in *ICSE*, 2009.
- [14] G. Génova, J. M. Fuentes, J. Llorens, O. Hurtado, and V. Moreno, "A framework to measure and improve the quality of textual requirements," *Requirements Engineering*, vol. 18, no. 1, pp. 25–41, sep 2011.
- [15] D. Falessi, G. Cantone, and G. Canfora, "Empirical principles and an industrial case study in retrieving equivalent requirements via natural language processing techniques," *IEEE Transactions on Software Engineering*, vol. 39, no. 1, pp. 18–44, 2013.
- [16] J. Winkler and A. Vogelsang, "Automatic classification of requirements based on convolutional neural networks," in *3rd International Workshop on Artificial Intelligence for Requirements Engineering (AIRE)*, 2016.
- [17] B. Alchimowicz, J. Jurkiewicz, J. Nawrocki, and M. Ochodek, "Towards use-cases benchmark," *Software Engineering Techniques*, 2011.
- [18] D. M. Berry, E. Kamsties, and M. M. Krieger, "From Contract Drafting to Software Specification : Linguistic Sources of Ambiguity," 2003.
- [19] D. M. Berry, A. Bucchiarone, S. Gnesi, G. Lami, and G. Trentanni, "A new quality model for natural language requirements specifications," in *REFSQ*, 2006, pp. 1–12.
- [20] H. Femmer, D. Méndez Fernández, E. Juergens, M. Klose, I. Zimmer, and J. Zimmer, "Rapid requirements checks with requirements smells: Two case studies," in *International Workshop on Rapid Continuous Software Engineering*, 2014, pp. 10–19.
- [21] H. Yang, A. D. Roeck, V. Gervasi, A. Willis, and B. Nuseibeh, "Analysing anaphoric ambiguity in natural language requirements," *Requirements Engineering*, vol. 16, no. 3, pp. 163–189, 2011.

- [22] S. J. Körner and T. Brumm, "Natural Language Specification Improvement With Ontologies," *International Journal of Semantic Computing*, vol. 3, no. 4, pp. 445–470, 2009.
- [23] B. Arendse and G. Lucassen, "Toward tool mashups: Comparing and combining NLP RE tools," in *3rd International Workshop on Artificial Intelligence for Requirements Engineering (AIRE)*, 2016, pp. 26–31.
- [24] M. Unterkalmsteiner and T. Gorschek, "Requirements quality assurance in industry: why, what and how?" in *REFSQ*, 2017.
- [25] J. R. Landis and G. G. Koch, "The Measurement of Observer Agreement for Categorical Data," *Biometrics*, vol. 33, no. 1, pp. 159–174, 1977.
- [26] D. M. Berry, R. Gacitua, P. Sawyer, and S. F. Tjong, "The case for dumb requirements engineering tools," *Lecture Notes in Computer Science*, vol. 7195 LNCS, pp. 211–217, 2012.
- [27] N. Kiyavitskaya, N. Zeni, L. Mich, and D. M. Berry, "Requirements for tools for ambiguity identification and measurement in natural language requirements specifications," *Requirements Engineering*, vol. 13, no. 3, pp. 207–239, jul 2008.
- [28] S. F. Tjong and D. M. Berry, "The design of SREE - A prototype potential ambiguity finder for requirements specifications and lessons learned," in *REFSQ*, 2013, pp. 80–95.
- [29] N. Ayewah, D. Hovemeyer, J. D. Morgenthaler, J. Penix, and W. Pugh, "Using static analysis to find bugs," *IEEE Software*, vol. 25, no. 5, pp. 22–29, 2008.
- [30] D. Andrzejewski, X. Zhu, and M. Craven, "Incorporating domain knowledge into topic modeling via dirichlet forest priors," in *ICML*. ACM, 2009, pp. 25–32.