

Continuous Rationale Identification in Issue Tracking and Version Control Systems

Anja Kleebaum^a, Barbara Paech^a, Jan Ole Johanssen^b and Bernd Bruegge^b

^aInstitute for Computer Science, Heidelberg University, Heidelberg, Germany

^bTechnical University of Munich, Munich, Germany

Abstract

During the elicitation and implementation of requirements, the involved stakeholders make decisions and build up important decision knowledge, which they should make explicit and share. Issue tracking and version control systems offer opportunities for a lightweight capture of decision knowledge but currently lack techniques for making this knowledge explicit. In this paper, we present techniques to make decision knowledge explicit in the text of issue tracking and version control systems. We present features for its manual documentation as well as for its automatic identification using a text classifier. The text classifier identifies implicit decision knowledge in natural language texts, in particular, in the description and comments of tickets in the issue tracking system and commit messages.

Keywords

Rationale Management, Decision Knowledge, Text Classification, Natural Language Processing, ConDec

1. Introduction

Decision knowledge, which is also referred to as *rationale*, is built up by requirements engineers, product owners, developers, and stakeholders of other roles during their daily tasks, i. e., while they elicit, prioritize, document, validate, manage, implement, and verify requirements [1]. Decision knowledge covers decision problems, their solution decisions, alternative solution options, and arguments [2]. The success of a software development project is closely linked to the stakeholders' ability to document and share their decision knowledge. Otherwise, the knowledge vaporizes and is no longer accessible for other stakeholders or even for the same stakeholders in the future [3]. However, documenting decision knowledge, i. e., making it explicit, is not an easy task [4]. Thus, decision knowledge often stays tacit or—if captured—is implicitly hidden in development artifacts such as commit messages and issue tracking data. We argue that the reason for this is the lack of techniques for making decision knowledge explicit in issue tracking systems (ITS) and version control systems (VCS).

Our overall goal is to develop techniques and workflows for a *continuous decision knowledge management*. In our previous work, we stated requirements [5, 6] and presented tool support

In: F.B. Aydemir, C. Gralha, S. Abualhaija, T. Breaux, M. Daneva, N. Ernst, A. Ferrari, X. Franch, S. Ghanavati, E. Groen, R. Guizzardi, J. Guo, A. Herrmann, J. Horkoff, P. Mennig, E. Paja, A. Perini, N. Seyff, A. Susi, A. Vogelsang (eds.): *Joint Proceedings of REFSQ-2021 Workshops, OpenRE, Posters and Tools Track, and Doctoral Symposium, Essen, Germany, 12-04-2021*

✉ kleebaum@informatik.uni-heidelberg.de (A. Kleebaum); paech@informatik.uni-heidelberg.de (B. Paech); jan.johanssen@tum.de (J. O. Johanssen); bruegge@in.tum.de (B. Bruegge)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

called ConDec [7]. With ConDec, requirements engineers and developers continuously document, share, and exploit decision knowledge in relation to other knowledge. In ConDec, all software artifacts, i. e. requirements and code, are referred to as knowledge elements. ConDec builds up a knowledge graph consisting of knowledge elements and trace links, which can be exploited by requirements engineers and developers for various purposes, e. g., for change impact analysis and during meetings.

In this paper, we focus on the Natural Language Processing (NLP) aspects of ConDec. We present tool features to document decision knowledge in the text of ITS and VCS, i. e., to make rationale explicit. In particular, we present how the documentation can be supported by an automatic text classifier. The text classifier is embedded in the ConDec Jira plug-in¹ that integrates into the ITS Jira. It classifies both Jira issue text (description and comments) as well as commit messages. It can be applied retrospectively on an existing set of Jira issues and commit messages. It is also of use during active development, i. e., while developers incrementally and collaboratively document decision knowledge. Thus, developers can directly validate the correctness of the identified decision knowledge. This distinguishes the ConDec classifier from existing work. The training data for the classifier covers relevant decision knowledge elements as well as irrelevant text. While ConDec offers a default training data set, the data of the current development project can also be used for the training. Besides, the classifier provides online training possibilities. We present first evaluation results on the quality of the text classifier.

In Section 2, we present the approach to document decision knowledge in the text of ITS and VCS. In Section 3, we present the text classifier to support the documentation. Subsequently, Section 4 presents the evaluation, Section 5 related work, and Section 6 concludes the paper.

2. Explicit Rationale in the ITS and VCS

In this section, we present the manual decision knowledge documentation in the ITS and VCS. We use developers as an example for stakeholders of different roles. Developers document decision knowledge when they 1) capture it, i. e., write it down, 2) annotate it, and 3) link it to other knowledge elements in the knowledge graph. Knowing this documentation approach is necessary to understand what we aim to support with the automatic text classifier (Section 3).

2.1. Process Assumptions: Requirements in the ITS and Trace Links

We make the following assumptions regarding requirements and trace links:

A1 Explicit Requirements: Requirements are documented in the ITS in an explicit way. That means that there are specific issue types, e. g., for epics and user stories or scenarios.

A2 Trace Links Between Requirements and Code: Trace links between requirements in the ITS and code files in VCS are created and maintained. This is commonly done in a commit-based way using issue identifiers in commit messages [8] or with other techniques [9].

A3 Trace Links Between ITS Issues: Also, the issues in the ITS are linked amongst each other. For example, a requirement is linked to its development tasks or—when using a hierarchical requirements model—epics are linked to user stories.

¹<https://github.com/cures-hub/cures-condec-jira>

2.2. Features For Making Rationale Explicit in the ITS and VCS

We present the ConDec tool features (F1 – F5) that enable manual decision knowledge documentation in the ITS and VCS. For every **tool feature**, we describe the *task of the developer* and some details, such as high-level design decisions that we made during its implementation.

F1 Easy capture: *Developers capture (write down) decision knowledge within their current development context, in particular, within the text of ITS and VCS.* We needed to decide where to capture decision knowledge. We decided to support the four documentation locations possible in the ITS and VCS: 1) documentation as entire ITS issues, similar to A1 for requirements, 2) documentation in the description and comments of existing ITS issues, e. g., in requirements or development tasks, 3) documentation in commit messages, and 4) documentation in code comments. These documentation locations enable decisions to be traced to requirements in the ITS and code in the VCS and can be used interchangeably. With the trace links in assumption A2 and A3, it is only a minor difference whether a decision is documented within the ITS, e. g., in the comment of a development task, or in the VCS, e. g., in a commit message linked to the development task. In either location, the decision can be accessed from the requirement.

F2 Explicit capture: *Developers annotate text as decision knowledge elements to make the decision knowledge explicit.* If decision knowledge elements are documented as entire ITS issues, they are explicit because of their type. In Jira issue text and commit messages, developers use a markup syntax, e. g., [decision] ... [/decision], whereas in code comments, they use a JavaDoc-like syntax, e. g. @decision ..., to annotate parts of text as decision knowledge elements. To integrate the parts of text as nodes in the knowledge graph, we decided to store them in the Jira database via object-relational mapping including their start and end positions in the text.

F3 Linking: *Developers link decision knowledge elements to other knowledge elements so that they can be accessed within the knowledge graph.* Explicit decision knowledge elements are nodes in the knowledge graph that can only be accessed from other nodes (knowledge elements) if they are linked. For this purpose, ConDec supports manual linking between all knowledge elements. Besides, ConDec offers techniques to automatically link knowledge elements within the knowledge graph, e. g., decisions to decision problems and code to requirements via commits. ConDec uses both Jira issue links and custom links, e. g., to link a requirement and the decision knowledge elements in its comments. The custom links are stored via object-relational mapping.

F4 Change: *Developers change and delete knowledge elements as well as trace links between them.* To maintain a consistent knowledge graph, developers need to ensure that the knowledge elements and links are correctly documented. For instance, developers need to change the annotations from F2, for example, to pick an alternative as the decision. Commit messages are different from the other three documentation locations as they are frozen by nature. They could be changed using a technique called (reword) rebasing, but this should only be done on the mainline. ConDec transcribes commit messages into ITS issue comments so that they can be changed (even just for spelling corrections) and so that they can be annotated.

F5 Attribute assignment: *Developers provide attributes for knowledge elements.* Similar to requirements, decision knowledge elements need attributes to document important metadata. For example, attributes can be the state, e. g., idea, decided, rejected, a decision level [10] or group [11]. We needed to solve the issue of how to store attributes for the parts of natural language text. ConDec stores attributes for parts of text via object-relational mapping (see F2).

2.3. Example for Explicit Rationale and Views on the Knowledge Graph

We use a decision problem that we needed to solve for the development of the text classifier (see Subsection 3.4) as an example: 📌 *Which library should we choose to implement the classifier?* This decision problem could be documented in each of the four documentation locations of F1. Figure 1 shows that it is documented in a Jira issue comment of the requirement *Automatically identify decision knowledge within the text of Jira issues*. The discussion on how to solve this decision problem is scattered in further comments and a commit message for the requirement as well as in a linked development task. The state of the first decision is rejected as it was replaced by a new decision documented during the implementation of the development task. ConDec offers various views on the knowledge graph [7]. Figure 2 shows a different tree view than in Figure 1. The developers filtered this view to only see decisions and their arguments concerning the requirement. Filtered-out knowledge elements are replaced by transitive links.

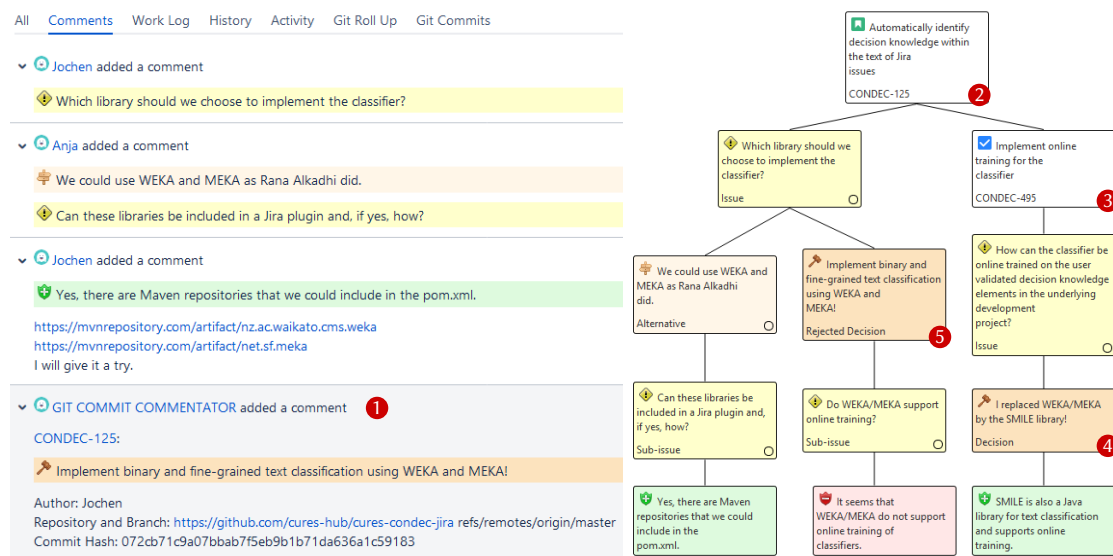


Figure 1: **Left:** Decision knowledge captured in the comments of a requirement in Jira. ① is a commit message that ConDec automatically transcribed into the new comment. **Right:** Knowledge tree visualization with the requirement ② being the root element. The subtree ③ on the right shows decision knowledge documented for a later development task. The decision ④ replaces the former decision ⑤.

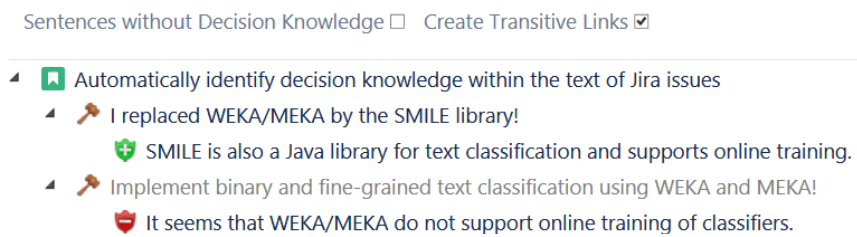


Figure 2: Filtered knowledge tree visualization for the requirement so that only the decisions and their arguments are shown. The rejected decision is colored in gray.

3. Automatic Text Classification

This section presents the automatic text classification. First, we describe ConDec’s tool features and their integration in the development process. Then, we describe the ground truth data used for training and evaluation, the preprocessing, the classifier(s), and implementation details.

3.1. Features for the Automatic Text Classification to Identify Rationale

ConDec offers the following **tool features** for the automatic identification of rationale:

F6 Automatic annotation: *The classifier predicts whether the textual parts in the Jira issue description, comments, or commit messages are relevant decision knowledge elements and—if yes—it annotates the parts accordingly.* Thus, the classifier supports developers to make decision knowledge elements explicit, i. e., it automates F2 for the explicit capture. We decided to not automatically classify code comments because they are unlikely to be used as locations for natural language discussions. The automatic annotation is part of the daily development. Besides, it can be applied retrospectively to existing projects.

F7 Easy training: *Developers train the classifier on training data from other projects and also on the decision knowledge documented for their current development project.* The text classifier is supervised, i. e., it needs training data to learn how to make predictions. We decided to add a default training file to enable the classification in new or existing projects without explicit rationale management. The training is performed in Jira. Online training is possible, i. e., newly documented knowledge elements are directly used for training. Besides, developers can create training data files from their development projects as well as import and export such files. The preprocessing (Subsection 3.3) is automatically performed.

F8 Easy evaluation: *Developers evaluate the classifier on the data of their current development project but also on the data of other projects.* The evaluation is directly possible in Jira. The training data can also be used for evaluation but is then split via k-fold cross-validation.

F9 Manual approval: *Developers manually approve the classification result, i. e., developers decide whether the annotations are correct or not.* The manual approval of classified parts of text is important to determine which parts of text can be used as training and evaluation data, i. e., as the ground truth for F7 and F8. Only the parts of text that are correctly annotated or correctly identified as irrelevant should be used as the ground truth. The information whether or not a classified part of text was manually approved by a human is stored in an attribute (F5).

3.2. Ground Truth Data

The ground truth for the text classifier comprises parts of text. The parts of text are either relevant decision knowledge elements or irrelevant. Each part of text is associated with a binary and a fine-grained label for its relevance and decision knowledge type, respectively. Five types are distinguished: issue, decision, alternative, pro-, and con-argument. As one preprocessing step for the classification is sentence splitting (see Subsection 3.3), the text classifier predicts new classifications (labels) for single sentences. However, the developer could also manually annotate sub-sentences or multiple sentences as one decision knowledge element. Then, the respective training data entry, i. e., the part of text, comprises a sub-sentence or multiple sentences. Table 1 shows a subset of the default training data included in the ConDec Jira plug-in.

Table 1

Examples for binary and fine-grained labeled parts of text in the default training data of ConDec.

Binary Label	Fine-grained Label	Part of Text (Sentence)
Irrelevant	-	Here's a small screenshot of the current state.
Relevant	🚩 Issue (Decision Problem)	What configuration possibilities should users have?
Relevant	🔧 Decision (Solution)	We decided to allow the following configuration ...
Relevant	🔄 Alternative	The user could configure the classifier algorithm.
Relevant	🟢 Pro	This will improve the performance.
Relevant	🔴 Con	I think this may be a bit confusing to the user.

The parts of text in the training data are taken from Jira issue text (description and comments), commit messages, and code comments. Code comments only contribute relevant decision knowledge elements. One decision problem in ConDec is how to deal with text from different documentation locations: Should we treat Jira issue text and commit messages equally or should there be different text classifiers? We decided to treat the documentation locations equally. In the future, it needs to be evaluated whether two different classifiers for Jira issue text and commit messages would perform better than one unique classifier.

3.3. Preprocessing

We use the following preprocessing steps: 1) sentence splitting, 2) conversion of sentences to lowercase, 3) tokenizing, 4) vector representation of words via Global Vectors for Word Representation (GloVe), 5) n-gram generation (with n=3). These preprocessing steps are commonly applied in NLP applications for requirements engineering [12] except for step 4: GloVe aims to represent the semantics of words [13] and is similar to Word2Vec, which is a rather novel technique [12]. We do not apply stemming and stop word removal because the words in GloVe are not stemmed and because stop words might be important parts of rationale [14, 15].

3.4. Classifiers and ConDec Implementation

We use two classifiers in a row: one for binary and one for fine-grained predictions, similar to Alkadhi *et al.* [14]. Only if the binary classifier predicts a sentence as relevant, the sentence serves as the input for the fine-grained classifier. As shown in the figures of Subsection 2.3, we decided to integrate the Statistical Machine Intelligence and Learning Engine (SMILE) library² into the ConDec Jira plug-in, because SMILE is implemented in Java and provides online-learning capabilities. For the preprocessing, we use the SMILE NLP component. ConDec offers a settings page for the automatic text classification, where developers can train and evaluate the binary and fine-grained classifiers. For instance, they can configure the machine learning algorithm, e. g., choose between logistic regression and support vector machine (SVM). The code for the ConDec Jira plug-in including the implementation of the automatic text classification as well as the data used for the evaluation (Section 4) can be found on GitHub³.

²<https://haifengl.github.io>

³<https://github.com/cures-hub/cures-condec-jira/blob/master/doc/features/automatic-text-classification.md>

4. Evaluation of the Automatic Text Classification

We evaluated the performance of the automatic text classifiers on the decision knowledge documentation of the ConDec Jira plug-in. During the development of the plug-in, we document decision knowledge in Jira and git as described in Section 2. We used irrelevant parts of text as well as decision knowledge elements from the ConDec Jira development project as the ground truth to train and evaluate the text classifier. The data set comprises 1468 relevant decision knowledge elements and 220 irrelevant parts of text. From the decision knowledge elements, 392 are issues (decision problems), 332 are decisions, 218 are alternatives, 288 pro-arguments, and 238 con-arguments. We used random undersampling to balance 1) relevant and irrelevant parts of text and 2) the decision knowledge elements by their type. We then used 10-fold cross-validation to train and evaluate 1) the binary and 2) fine-grained classifiers. Table 2 shows the results of the evaluation for logistic regression (LR) classifiers.

Table 2

Evaluation results of the binary and fine-grained LR classifiers on the data of the ConDec Jira project.

	Binary (Relevant/Irrelevant)	Issue	Alternative	Decision	Pro	Con
Precision	0.65	0.51	0.42	0.25	0.57	0.51
Recall	0.97	0.40	0.18	0.85	0.11	0.13
F1-score	0.78	0.45	0.25	0.39	0.19	0.20

The results of our evaluation are partly worse compared with the results reported by others. For example, Bhat *et al.* [16] and Li *et al.* [15] achieved an F1-score above 0.7 for decisions, whereas we currently achieve 0.39. There might be various reasons for this: First, alternatives and decisions can be hard to distinguish, even for humans [14]. Second, every entry in the ground truth needs to be manually approved (F9) but also the manual classification is subjective and might be wrong. Third, we focused on the specification and implementation of ConDec’s tool features (F1 – F9) for the manual and automatic text classification rather than on performing systematic experiments regarding the preprocessing and classifiers.

Another aspect, which needs to be discussed, is that our data set might be rather artificial because the text was already written as we want decision knowledge elements to be written. For example, most issues are formulated as questions. Other projects with no explicit decision knowledge management are likely to be different. In particular, it might be problematic that the classifier classifies entire sentences: First, one sentence could contain multiple decision knowledge elements. Second, one decision knowledge element could be distributed over a longer part of text [14]. However, as one goal—next to the retrospective application—is to integrate the classifier into the daily development and to nudge the developers to create a clear decision knowledge documentation, other training data can be similar to the ConDec data.

5. Related Work

In this section, we discuss related work regarding the integration of automatic rationale identification into software development tools and workflows. There are various approaches concerning

the automatic identification of rationale in natural language text, e. g., [17], [16], [14], [15], and [18]. Most of the existing approaches apply the automatic classification retrospectively on a ground truth that the researchers created. There are only a few approaches that integrate automatic text classification into tools and workflows for software development. As part of their future work, Rogers *et al.* [17] and Lester *et al.* [19] plan to integrate text classification into existing knowledge management tools. Lester *et al.* [19] work on text classification integration into the C_SEURAT tool for *Collaborative Software Engineering Using Rationale*. Bhat *et al.* [20] developed the ADeX tool to manage architectural design decisions that also offers text classification for decisions. They use the architecture knowledge management system AMELIE to collect and access knowledge for requirements. ConDec uses the ITS (Jira) and VCS (git) as the central repositories for decision knowledge, requirements, and code.

6. Conclusions and Future Work

ConDec enables developers to informally capture decision knowledge and integrates features for the automatic classification into their daily work instead of only applying it retrospectively. With ConDec, the developers can create the ground truth for the training and evaluation of the classifiers themselves. They are confronted with the views on the knowledge graph during their active development and are thus nudged to improve and exploit the documentation.

Decision knowledge can be captured in other locations, such as chat messages, pull requests, or wiki pages. We also work on other ConDec integrations⁴, for example, a Slack plug-in to integrate chat messages and a Bitbucket plug-in to integrate pull requests. We treat the text in other documentation locations similar to commit messages. That means that relevant chat messages or comments in pull requests would be transcribed into Jira issue comments so that the decision knowledge can easily be traced to requirements and code.

For future improvements, there are various ways to extend our text classifier and work on its evaluation. We plan to perform experiments with different preprocessing steps, different classifiers (not only logistic regression as used in Section 4), and different training data (e. g., would the classification result be better if we had different classifiers for each documentation location?). We will add data balancing using Synthetic Minority Oversampling Technique (SMOTE). We will do cross-project evaluations and evaluate how well the classifier performs on projects with and without explicit decision knowledge documentation.

Acknowledgments

This work was partially supported by the DFG (German Research Foundation) under the Priority Programme SPP1593: Design For Future – Managed Software Evolution (CURES project). We thank Rana Alkadhi for valuable hints during the development of the text classifier and for sharing her data and knowledge. We thank the students who developed the ConDec tools, in particular, Jochen Clormann and Philipp De Sombre who contributed to the text classifier.

⁴<https://github.com/cures-hub>

References

- [1] A. H. Dutoit, R. McCall, I. Mistrik, B. Paech, *Rationale Management in Software Engineering: Concepts and Techniques*, Springer, 2006.
- [2] T.-M. Hesse, *Supporting Software Development by an Integrated Documentation Model for Decisions*, Ph.d. thesis, Heidelberg University, 2020. doi:10.11588/heidok.00028713.
- [3] R. Capilla, A. Jansen, A. Tang, P. Avgeriou, M. A. Babar, 10 years of software architecture knowledge management: Practice and future, *Journal of Systems and Software* 116 (2016) 191–205. doi:10.1016/j.jss.2015.08.054.
- [4] A. Kleebaum, J. O. Johanssen, B. Paech, B. Bruegge, How do Practitioners Manage Decision Knowledge during Continuous Software Engineering?, in: 31st International Conference on Software Engineering and Knowledge Engineering, SEKE'19, KSI Research Inc., 2019, pp. 735–740. doi:10.18293/SEKE2019-206.
- [5] A. Kleebaum, J. O. Johanssen, B. Paech, B. Bruegge, Tool support for decision and usage knowledge in continuous software engineering, in: 3rd Workshop on Continuous Softw. Engineering, 2018, pp. 74–77. doi:10.11588/heidok.00024186.
- [6] A. Kleebaum, M. Konersmann, M. Langhammer, B. Paech, M. Goedicke, R. Reussner, *Continuous Design Decision Support*, Springer, Cham, 2019, pp. 107–139. doi:10.1007/978-3-030-13499-0_6.
- [7] A. Kleebaum, J. O. Johanssen, B. Paech, B. Bruegge, Continuous Management of Requirement Decisions Using the ConDec Tools, in: 26th International Conference on Requirements Engineering Workshops, Doctoral Symposium, Live Studies Track, and Poster Track (REFSQ20), CEUR-WS.org, Pisa, Italy, 2020. doi:10.11588/heidok.00028230.
- [8] M. Rath, J. Rendall, J. L. C. Guo, J. Cleland-Huang, P. Mäder, Traceability in the wild, in: 40th International Conference on Software Engineering (ICSE), ACM Press, Gothenburg, Sweden, 2018, pp. 834–845. doi:10.1145/3180155.3180207.
- [9] P. Hübner, B. Paech, Interaction-based creation and maintenance of continuously usable trace links between requirements and source code, *Empirical Software Engineering (ESE)* 25 (2020) 4350–4377. doi:10.1007/s10664-020-09831-w.
- [10] J. S. van der Ven, J. Bosch, Making the right decision: Supporting architects with design decision data, in: *Lecture Notes in Computer Science*, volume 7957, Springer, Berlin, Heidelberg, 2013, pp. 176–183. doi:10.1007/978-3-642-39031-9_15.
- [11] U. van Heesch, P. Avgeriou, R. Hilliard, A documentation framework for architecture decisions, *Journal of Systems and Software* 85 (2012) 795–820. doi:10.1016/j.jss.2011.10.017.
- [12] L. Zhao, W. Alhoshan, A. Ferrari, K. J. Letsholo, M. A. Ajagbe, E. V. Chioasca, R. T. Batista-Navarro, *Natural Language Processing (NLP) for requirements engineering: A systematic mapping study*, arXiv:2004.01099.
- [13] J. Pennington, R. Socher, C. Manning, Glove: Global Vectors for Word Representation, in: *Conf. on Empir. Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar, 2014, pp. 1532–1543. doi:10.3115/v1/D14-1162.
- [14] R. M. A. Alkadhhi, *Rationale in Developers' Communication*, Ph.D. thesis, Technical University of Munich, 2018.
- [15] X. Li, P. Liang, Z. Li, Automatic Identification of Decisions from the Hibernate Developer Mailing List, in: *Evaluation and Assessment in Software Engineering*, December, ACM, Trondheim, Norway, 2020, pp. 51–60. doi:10.1145/3383219.3383225.
- [16] M. Bhat, K. Shumaiev, A. Biesdorf, U. Hohenstein, F. Matthes, Automatic Extraction of Design Decisions from Issue Management Systems: A Machine Learning Based Approach, in: 11th European Conference on Software Architecture (ECSA'17), Springer, Cham, Switzerland, 2017, pp. 138–154. doi:10.1007/978-3-319-65831-5_10.
- [17] B. Rogers, Y. Qiao, J. Gung, T. Mathur, J. E. Burge, Using text mining techniques to extract rationale from existing documentation, in: 6th Int. Conf. on Design Computing and Cognition, Springer, 2014, pp. 457–474. doi:10.1007/978-3-319-14956-1_26.
- [18] P. N. Sharma, B. T. R. Savarimuthu, N. Stanger, Extracting Rationale for Open Source Software Development Decisions – A Study of Python Email Archives, in: 43rd International Conference on Software Engineering (ICSE 2021), 2021. arXiv:2102.05232.
- [19] M. Lester, M. Guerrero, J. Burge, Using evolutionary algorithms to select text features for mining design rationale, *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 34 (2020) 132–146. doi:10.1017/S0890060420000037.
- [20] M. Bhat, C. Tinnes, K. Shumaiev, U. Hohenstein, A. Biesdorf, F. Matthes, ADEx: A tool for automatic curation of design decision knowledge and recommendation of alternatives and experts, in: *International Conference on Software Architecture (ICSA 2019)*, Hamburg, Germany, 2019, pp. 158–161. doi:10.1109/ICSA-C.2019.00035.