<https://helda.helsinki.fi>

Towards Utility-Based Prioritization of Requirements in Open Source Environments

Felfernig, A.

IEEE

2018-08-01

Felfernig , A , Stettinger , M , Atas , M , Samer , R , Nerlich , J , Scholz , S , Tiihonen , J & Raatikainen , M 2018 , Towards Utility-Based Prioritization of Requirements in Open Source Environments . in G Ruhe , W Maalej & D Amyot (eds) , 2018 IEEE 26th International Requirements Engineering Conference (RE) : 20-24 August Banff, Alberta Canada . Proceedings of the ... IEEE International Symposium on Requirements Engineering , IEEE , pp. 406-411 , IEEE International Requirements Engineering Conference , Banff , Alberta , Canada , 20/08/2018 . <https://doi.org/10.1109/RE.2018.00-17>

<http://hdl.handle.net/10138/321554>

<https://doi.org/10.1109/RE.2018.00-17>

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Towards Utility-based Prioritization of Requirements in Open Source Environments

Alexander Felfernig*, Martin Stettinger*, Müslüm Atas*, Ralph Samer*,
Jennifer Nerlich†, Simon Scholz†, Juha Tiisonen‡, and Mikko Raatikainen‡

*Graz University of Technology, Austria, Email: {afelfern,mstettinger,muatas,rsamer}@ist.tugraz.at,

†vogella GmbH, Hamburg, Germany, Email: {jennifer.nerlich,simon.scholz}@vogella.com,

‡University of Helsinki, Finland, Email: {juha.tiisonen,mikko.raatikainen}@helsinki.fi

Abstract—Requirements Engineering in open source projects such as ECLIPSE faces the challenge of having to prioritize requirements for individual contributors in a more or less unobtrusive fashion. In contrast to conventional industrial software development projects, contributors in open source platforms can decide on their own which requirements to implement next. In this context, the main role of prioritization is to support contributors in figuring out the most relevant and interesting requirements to be implemented next and thus avoid time-consuming and inefficient search processes. In this paper, we show how utility-based prioritization approaches can be used to support contributors in conventional as well as in open source Requirements Engineering scenarios. As an example of an open source environment, we use BUGZILLA. In this context, we also show how dependencies can be taken into account in utility-based prioritization processes.

I. INTRODUCTION

In software projects, resources are typically limited which requires the prioritization of requirements [16]. Prioritization is often interpreted as a part of *strategic planning* where the focus is to select and prioritize requirements that should be included in releases (long-term release planning) [4], [19]. Decision support in prioritization is extremely important since especially when dealing with large assortments of requirements, manual prioritization processes tend to become very costly [3], [28]. In this context, suboptimal prioritizations trigger time wasting due to the implementation of unimportant requirements.

There are *two basic approaches* to prioritize requirements – for an in-depth related analysis we refer to Achimugu et al. [1]. *First*, requirements prioritization can be interpreted as an *optimization task* where the overall objective is to identify the *middle ground*, i.e., an aggregation of individual prioritizations into a global prioritization that reflects the least possible level of dissimilarity from all stakeholder-individual prioritizations [14]. *Second*, in contrast to approximating individual prioritizations on the basis of optimization functions, *utility-based approaches* focus on (1) establishing agreement with regard to the evaluation of individual requirements and (2) thereafter determining prioritizations [2], [11], [26].

Prioritizations following the *optimization approach* are determined on the basis of individual prioritizations of stakeholders. When following a *utility-based approach*, preferences of stakeholders are first aggregated and a prioritization is

determined thereafter. In the line of basic approaches to determine group recommendations [9], the first approach is based on *aggregated prioritizations* where stakeholder-individual prioritizations are known and a recommendation minimizes dissimilarities between the given prioritizations (preferences). The second approach is based on *aggregated models* where stakeholder requirement evaluations are aggregated first and a prioritization is determined on the basis of a group profile (model) derived from requirements evaluations.

Aggregated models have the advantage that stakeholders are encouraged to focus their evaluations on specific relevant aspects of a requirement (e.g., dimensions such as profit, risk, and effort) and thus contribute to stable preferences and a higher degree of consensus [22]. Aggregated prioritizations trigger scalability issues since each stakeholder has to provide, for example, a ranked list of requirements as input for the optimization process. Furthermore, due to the computational complexity of the underlying problem, an optimal solution cannot be guaranteed and is often only approximated on the basis of local search algorithms [14], [24], [29]. Utility-based approaches as discussed in this paper focus on evaluations of individual requirements on the basis of different evaluation dimensions (e.g., profit, effort, and risk). This way, stakeholders can focus on evaluating requirements they have knowledge about and the focus of prioritization is first on establishing *consensus* and thereafter on figuring out the most relevant prioritizations [22].

Different algorithmic approaches can be used to support requirements prioritization – for an overview, see, for example, Achimugu et al. [1]. Examples thereof are constraint-based reasoning [25], incremental preference learning [17], evolutionary algorithms [14], machine learning [3], [23], and pairwise preference-based decision making [20]. *Optimization-based approaches* focus on minimizing the distance between the preferences of individual stakeholders (e.g., in terms of the distance between individual prioritizations and the prioritization determined by the optimization approach). A similar problem also solved on the basis of optimization approaches is the *next release problem* [6], [27] where a subset of a given set of requirements has to be selected in such a way that predefined cost limits are taken into account and the chosen set of requirements represents the optimum choice in terms of criteria such as market value. The focus in this

context is more to identify subsets of requirements but not to prioritize a given list of requirements. In contrast to next release problems, prioritization tasks in open source scenarios do not necessarily require (and often do not allow) a global optimum but more focus on relevant recommendations for individual stakeholders. Also in contrast to existing release planning tasks, developers in open source scenarios in most of the cases do not explicitly define their preferences, i.e., preferences have to be derived from given interaction data (in our case, interaction data collected by BUGZILLA¹).

Utility-based prioritization based on *multi-attribute utility theory* [8] can be implemented in different variants. First, requirements are simply evaluated with regard to a set of predefined interest dimensions and the overall utility of a requirement is determined as a sum of interest dimension specific utilities. Second, weights can be introduced to emphasize on specific interest dimensions (e.g., a lower risk is more important than high profits). Third, stakeholders can be enabled to define their personal evaluations and utility-based approaches should then be able to aggregate these evaluations and take into account stakeholder weights. Stakeholder weights can be interpreted as "global", i.e., there is a global weighting of stakeholders independent of a specific dimension or requirement. If weights are interpreted as "local", the importance of a stakeholder can be defined on the level of individual requirements or dimensions. Utility-based prioritization can also be implemented on the basis of analytic hierarchy process (AHP) [13]. A major disadvantage of this approach is that requirements have to be evaluated pairwise which does not scale well when the number of requirements increases.

Prioritization criteria differ depending on the requirements engineering scenario. The criteria *effort*, *risk*, and *profit* are often used in settings where a group of stakeholders engaged in the same project is in charge of completing a prioritization task [1]. In contrast, in open-source settings, developers are in most of the cases engaged in different projects and also work for different companies. In such scenarios, prioritization is less focusing on establishing consensus between individual stakeholders but more on supporting stakeholders in identifying requirements of relevance to them and to prioritize the important ones by also taking into account global criteria. Examples of criteria in such scenarios are *personal expertise* of a developer and *importance* of a requirement for the community of the stakeholder and the open source community as a whole. Thus, open source platform related prioritization processes completely differ from conventional software projects. A major focus of this paper is to introduce prioritization concepts especially applicable in open source development contexts.

BUGZILLA is an open-source based issue tracking system which supports users from different geographical locations to report their findings with regard to a given set of software components. Users can submit textual descriptions of issues and corresponding meta-information, for example, associated

components, keywords, and dependencies. Reported issues can be selected by contributors to work on. In BUGZILLA, issues can be requirements but also reported bugs. Distinguishing between these can be performed on the basis of a meta-attribute (issue type) that can be specified for BUGZILLA issues. There are different related approaches to support machine learning based requirements prioritization. The approaches operate on datasets including historical data of previous requirement (bug report in BUGZILLA) selections and try to predict future requirement selections thereof. Utility-based prioritization can be used in interactive scenarios (stakeholders are engaged in an interactive prioritization process) as well as scenarios where requirements are recommended but no further stakeholder interaction is needed for determining a prioritization.

The *contributions* of this paper are the following. We provide an overview of different application scenarios of utility-based requirements prioritization and discuss specific aspects of requirements prioritization in open source projects. For scenarios that include dependencies between requirements, we show how such dependencies can be taken into account on the basis of the concepts of model-based diagnosis [10], [18]. With this approach, we tackle the following *research gaps*. In contrast to existing prioritization and release planning approaches, we introduce model-based diagnosis concepts that also support re-prioritization and re-planning while not completely omitting already existing stakeholder preferences which is still an open issue in most of the existing prioritization and release planning approaches (these approaches focus on taking into account dependencies but do not support the aforementioned re-prioritization and re-planning scenarios). Furthermore, we present a first version of a user interface developed to support prioritization tasks in open source environments – in our case, for BUGZILLA users. This approach has the potential to reduce the workload of developers in open source platforms by automatically proposing requirements that should be implemented next instead of forcing users to analyze in detail a large number of requirements. Furthermore, the introduced utility-based approach does not require the "manual" evaluation of individual requirements but automatically derives utility models by analyzing given interaction logs taking into account interaction data such as *number of comments* related to a requirement. This is a major difference compared to existing requirements prioritization approaches which do not support automated prioritization based on background data. Finally, we discuss issues for future work to further advance the state of the art in utility-based prioritization.

The remainder of this paper is organized as follows. In Section II, we introduce variants of implementing *utility-based prioritization*. Thereafter, we introduce our variant of utility-based prioritization implemented for the BUGZILLA environment and provide a sketch of a related BUGZILLA user interface (Section III). In Section IV, we show how to extend utility-based prioritization in such a way that dependencies between requirements can be taken into account. The paper is concluded with a discussion of future work in Section V.

¹www.bugzilla.org.

II. UTILITY-BASED PRIORITIZATION

Utility-based prioritization allows stakeholders to prioritize a requirement with regard to different interest dimensions $D = \{d_1, d_2, \dots, d_n\}$. Examples of such interest dimensions are *profit*, *risk*, and *effort*. Utility-based prioritization is based on the idea to first evaluate each requirement with regard to the set of interest dimensions (see Table I) and thereafter calculate the individual utility of each requirement (see Formula 1). In general, the *priority* is associated with the *utility* of a requirement r which results from its total contributions to all of each individual interest dimensions d (denoted as $contribution(r, d)$) combined with the corresponding importance weights of individual interest dimensions (denoted as $weight(d)$).

TABLE I
CONTRIBUTION OF REQUIREMENTS $R = \{r_1, r_2, r_3\}$ TO THE INTEREST DIMENSIONS $D = \{profit, risk, effort\}$.

interest dimension	r_1	r_2	r_3
profit	10	5	4
risk	7	2	8
effort	2	3	7

TABLE II
PREDEFINED WEIGHTS FOR THE INTEREST DIMENSIONS $D = \{profit, risk, effort\}$.

interest dimension	weights
profit	0.3
risk	0.5
effort	0.2

$$utility(r, D) = \sum_{d \in D} (contribution(r, d) \times weight(d)) \quad (1)$$

Applying Formula 1 to the entries in Tables I and II results in the ranking depicted in Table III (the higher the utility with regard to the given interest dimensions, the higher the corresponding priority of the requirement).

TABLE III
RANKING OF REQUIREMENTS WITH STATIC WEIGHTS.

requirement	r_1	r_2	r_3
utility	6.9	3.1	6.6
priority (ranking)	1	3	2

In the previous example, the evaluation of requirements with regard to interest dimensions and the weighting of interest dimensions are assumed to be predefined (e.g., by a single stakeholder). However, requirements prioritization is often a group decision process [9] where different stakeholders are evaluating requirements (see, e.g., Table IV) and define importance weights with regard to interest dimensions (see, e.g., Table V). Both, stakeholder-individual evaluations of interest

dimensions and importance weights have to be aggregated. Formula 2 shows the aggregation of stakeholder-individual evaluations of requirements where S refers to the set which includes all m stakeholders (i.e., $S = \{s_1, s_2, \dots, s_m\}$).

$$contribution(r, d, S) = \frac{\sum_{s \in S} eval(d, r, s)}{|S|} \quad (2)$$

Formula 3 shows how to aggregate the stakeholder-specific importance weights (denoted as $w(d, s)$) which are related to individual interest dimensions d . Previous calculations did not take into account potential different degrees of stakeholder expertise, for example, a stakeholder s_a could have more expertise with regard to estimating the market potential of a requirement in terms of profit as estimating the corresponding development efforts. To take into account this aspect, Formula 3 includes a factor that represents the *expertise* of a stakeholder s with regard to a specific interest dimension d .

$$weight(d, S) = \frac{\sum_{s \in S} w(d, s) \times expertise(d, s)}{|S|} \quad (3)$$

Similar to the basic approach, the utility of a requirement (Formula 4) is determined as a combination of the contributions of a requirement to the given interest dimensions and related interest dimension importance evaluations of stakeholders.

$$utility(r, D, S) = \sum_{d \in D} (contribution(r, d) \times weight(d, S)) \quad (4)$$

The result of applying Formulae 2–4 to the evaluation data contained in Tables IV and V is depicted in Table VI.

III. UTILITY-BASED PRIORITIZATION IN BUGZILLA

In Section II, we took a look at different variants of utility-based prioritization. These variants were discussed on the basis of interest dimensions (evaluation criteria) typically occurring in software projects where a group of stakeholders is in charge of jointly defining and prioritizing requirements. In this section, we focus on open source scenarios where individual users (e.g., contributors in an open source platform) follow their individual interests regarding requirements without necessarily taking into account the preferences of other users. This can be considered a major difference compared to conventional software projects where stakeholders commonly develop a "global" prioritization (see Section II). We now show how utility-based prioritization can be applied in such contexts.

Table VII represents a BUGZILLA-specific evaluation of requirements (bugs) with regard to the set of interest dimensions $\{cc, geritt, blocker, comments\}$. In this context, *cc* is the number of contributors who are in the :cc list of bug-related emails, *geritt* is the number of bug-related GERITT² changes, *blocker* is the number of dependent bugs (dependent

²A code reviewing tool – gerritcodereview.com.

TABLE IV
CONTRIBUTION OF REQUIREMENTS $R = \{r_1, r_2, r_3\}$ TO DIMENSIONS $D = \{profit, risk, effort\}$ (DEFINED BY STAKEHOLDERS $S = \{s_1, s_2, s_3\}$).

interest dimension	r_1				r_2				r_3			
	s_1	s_2	s_3	AVG	s_1	s_2	s_3	AVG	s_1	s_2	s_3	AVG
profit	5	2	2	3.0	5	1	2	2.7	2	2	6	3.3
risk	3	3	4	3.3	2	5	6	4.3	3	2	2	2.3
effort	2	3	2	2.3	3	4	2	3.0	5	6	2	4.3

TABLE V
PREFERENCES OF STAKEHOLDERS $S = \{s_1, s_2, s_3\}$ WITH REGARD TO THE INTEREST DIMENSIONS $D = \{profit, risk, effort\}$.

stakeholder	s_1	s_2	s_3	weights
profit	0.5	0.3	0.6	0.47
risk	0.3	0.6	0.3	0.4
effort	0.2	0.1	0.1	0.13

TABLE VI
RANKING OF REQUIREMENTS WITH GROUP WEIGHTS.

requirement	r_1	r_2	r_3
utility	3.03	3.57	3.03
priority (ranking)	2	1	2

requirements), and *comments* refers to the number of bug-related comments. These interest dimensions do need to be evaluated manually as it is often the case in other scenarios [15], [21] but can directly be determined from corresponding user interaction data.

Formula 5 supports the calculation of the contribution of a requirement r to a specific interest dimension d . In sharp contrast to the previous scenarios, the contribution is not directly specified by stakeholders but derived from BUGZILLA specific information (e.g., #comments related to a requirement).

$$contribution(r, d) = eval(r, d) \quad (5)$$

Formula 6 supports the determination of the expertise of a stakeholder which is represented in terms of the similarity between the keywords stored in the stakeholder profile and those extracted from the requirement description and corresponding meta-information (e.g., name of associated component/system). The similarity between requirement-related keywords, meta-information, and contributor profile information is interpreted as *expertise* (see Formula 6).

$$weight(r, s) = expertise(r, s) \quad (6)$$

In the line of the previously discussed utility functions, the overall utility of a requirement is interpreted as a combination of (1) the contributions of a requirement to a set of interest dimensions and (2) the expertise level of a stakeholder (in this context interpreted "globally", i.e., not on the level of individual interest dimensions).

$$utility(r, s) = \sum_{d \in D} contribution(r, d) \times weight(r, s) \quad (7)$$

TABLE VII
CONTRIBUTION OF REQUIREMENTS (BUGS) $R = \{r_1, r_2, r_3\}$ TO THE INTEREST DIMENSIONS $D = \{cc, geritt, blocker, comments\}$.

interest dimension	r_1	r_2	r_3
cc	5	2	2
geritt	3	3	4
blocker	2	3	2
comments	2	3	2

TABLE VIII
EXPERTISE OF STAKEHOLDER s_1 WITH REGARD TO THE REQUIREMENTS $\{r_1, r_2, r_3\}$ DETERMINED, FOR EXAMPLE, ON THE BASIS OF THE SIMILARITY BETWEEN THE STAKEHOLDER PROFILE AND INFORMATION ASSOCIATED WITH A REQUIREMENT.

stakeholder	s_1
r_1	0.5
r_2	0.3
r_3	0.2

TABLE IX
RANKING OF BUGZILLA BUGS WITH STATIC WEIGHTS.

requirement	r_1	r_2	r_3
utility	6.0	3.3	2.0
priority	1	2	3

IV. TAKING INTO ACCOUNT DEPENDENCIES

Blocking Factor. Utility-based recommendation approaches per se do not explicitly take into account dependencies between requirements (e.g., requirement A must be implemented before requirement B). In the discussed open source prioritization scenario, this aspect is taken into account by prioritizing requirements on the basis of the number of related dependencies, i.e., the higher the number of requirements dependent from a requirement x , the higher the "global" relevance of x . In this context, the *blocking factor* (i.e., how many requirements depend on the implementation of requirement x) can be considered as interest dimension that has an impact on prioritization. In other words, this requirement should be implemented as soon as possible since it otherwise blocks the implementation of other requirements. This approach can also be applied in software development scenarios where a group of stakeholders (e.g., an in-house software development project) is in charge of prioritizing requirements. Such an approach helps to avoid situations where prioritizations violate dependency constraints.

Automated Repair. An alternative approach is to apply repair mechanisms from model-based diagnosis [10] that

Fig. 1. BUGZILLA view on bugs (requirements). Based on the presented utility-based prioritization approach, bugs are presented to BUGZILLA contributors in the order of the determined priority.

Id	Summary	Priority	Product	Component
398925	After prolonged usage Eclipse becomes unusable : java.lang.Object cannot be cast to org.eclipse.e4.	100%	Platform	UI
229823	Provide notification API in Eclipse platform	92%	Platform	UI
201589	[Contributions] visibleWhen has no effect on toolbar	82%	Platform	UI
344290	[KeyBindings] mac delete (to the left) and arrow keys stop working in java editing windows	59%	Platform	UI
465456	E4 should not store reference to view and editor icon contribution value	48%	Platform	UI
430090	[Workbench] Provide a way to clear the persisted state when application is updated	48%	Platform	UI
124530	[EFS] Need EFS support for file import	47%	Platform	UI
272794	[DataBinding] TitleAreaDialogSupport should use TitleAreaDialog.setDialogComplete (new in 3.6)	45%	Platform	UI
473278	PartRenderingEngine's limbo is sometimes visible	42%	Platform	UI
303868	[DataBinding] Caching of old value makes Eclipse-Databinding unusable for CDO	42%	Platform	UI
317465	Implement context menu equivalents for 3.x context menus on Views/Editors/Perspectives (in the s	41%	Platform	UI
385278	[KeyBindings] Shortcuts like Ctrl+C/Ctrl+V intermittently stop working	41%	Platform	UI
103045	[EditorMgmt] [RCP] Need way to disable "New Editor" action	40%	Platform	UI
386648	[EditorMgmt] "No editor descriptor for id org.eclipse.ui.internal.emptyEditorTab" after workbench c	39%	Platform	UI
469595	When using org.eclipse.ui.contexts.window binding context E4 application fails to open	38%	Platform	UI

help to adapt already determined prioritizations in such a way that all defined dependencies are taken into account. In the following, we will shortly sketch our approach. In order to trigger a diagnosis process, we are in the need of a pre-defined set of dependencies between requirements (denoted as $DEP = \{dep_1, dep_2, \dots, dep_n\}$). Furthermore, we assume that a prioritization (represented as sequence) $P = [p_1, p_2, \dots, p_m]$ determined by a utility-based prioritization approach is *inconsistent* with the given set of dependencies. In order to apply model-based diagnosis, we assume that both, the pre-defined set of dependencies and the requirement prioritization is represented in terms of constraints [25], for example, $DEP = \{dep_1 : r_3 < r_1, dep_2 : r_3 < r_2\}$ and $P = \{p_1 : r_1 < r_2, p_2 : r_2 < r_3, p_3 : r_3 < r_4, p_4 : r_4 < r_5, p_5 : r_5 < r_6\}$. As can be easily seen, $DEP \cup P$ is inconsistent. As variable domains we assume $[1.. \#requirements]$.

Following the principles of model-based diagnosis [10], [18], we need to detect all *minimal conflicts* [12] induced in P by the dependencies defined in DEP . In this context, a conflict set is defined as follows.

Definition: Conflict Set (CS). A conflict set $CS \subseteq P$ is a set of individual prioritization elements that are inconsistent with the elements of DEP , i.e., $inconsistent(CS \cup DEP)$. A conflict set CS is minimal if $\neg \exists CS' : CS' \subset CS$.

On the basis of a set of identified minimal conflict sets, a corresponding diagnosis includes a set of prioritization elements in P that have to be adapted such that the consistency of $DEP \cup P$ is restored (see the following definition).

Definition: Diagnosis (Δ). A diagnosis $\Delta \subseteq P$ is a set of individual prioritization elements that have to be deleted from P such that $consistent(P - \Delta \cup DEP)$.

In our example, $CS : \{p_2\}$ is the only conflict induced in P by the dependencies defined in DEP . CS is minimal, i.e., we need to adapt only one of the prioritization elements in CS such that a global prioritization can be found that is consistent with the elements in DEP [12]. A corresponding diagnosis Δ is $\{p_2\}$. In our example, we could decide to replace $p_1 : r_1 < r_2$ with the corresponding repair $r_1 > r_2$. This is a repair action that helps to restore the consistency of $DEP \cup P$.

Our approach to the repair of inconsistent prioritizations can be used for both, *interactive prioritization* where stakeholders receive feedback on the consistency of prioritizations, and *automated prioritization* where repairs for inconsistent prioritizations are determined in an automated fashion. Important issues to improve our approach are discussed in the following.

V. CONCLUSION AND FUTURE WORK

Conclusion. In this paper, we showed how to support utility-based requirements prioritization. These scenarios range from *single user prioritization* where one stakeholder is in charge of completing prioritization tasks to *group-based prioritization* where the preferences/evaluations of different group members have to be taken into account. On the basis of these scenarios, we showed how utility-based prioritization can be applied in the context of open source development projects. In this context, we sketched our initial implementation currently provided in the BUGZILLA environment. This implementation serves as a first version to support prioritization in BUGZILLA.

Future Work. Since prioritization is a repetitive process, we will include mechanisms that are capable of learning stakeholder weights and also the weights of individual requirements. This approach will help to further increase the prediction quality of prioritizations in terms of the probability that stakeholders accept the proposed prioritizations. In this context, we will also compare the predictive quality of utility-based approaches (i.e., approaches based on aggregated models) with machine learning based approaches and approaches that determine rankings on the basis of aggregated prioritizations. Furthermore, we will analyze which further features (interest dimensions) are useful to improve prediction quality. For example, the number of redundant bugs (issues) can be a further important relevance indicator.

A major challenge in requirements prioritization is the provision of persuasive user interfaces that increase the preparedness of stakeholders to actively engage in requirements engineering processes [5]. Consequently we will focus on a further extension/improvement of the existing BUGZILLA requirements prioritization user interface. Furthermore, we will analyze in which way recommended prioritizations have to

be *explained* to support specific group decision goals such as *consensus*, *fairness*, and *decision quality* [7], [9]. Finally, we will conduct a detailed empirical study regarding the impact of our prioritization approaches in the ECLIPSE community.

ACKNOWLEDGMENT

The work presented in this paper has been conducted within the scope of the EC-funded Horizon2020 project OPENREQ.

REFERENCES

- [1] P. Achimugu, A. Selamat, R. Ibrahim, and M. Mahrin. A systematic literature review of software requirements prioritization research. *Information and Software Technology*, 56(6):568–585, 2014.
- [2] G. Adomavicius, N. Manouselis, and Y. Kwon. *Recommender Systems Handbook*, chapter Multi-Criteria Recommender Systems, pages 769–803. Springer, 1st edition, 2010.
- [3] M. Alenezi and S. Banitaan. Bug reports prioritization: Which features and classifier to use? In *12th International Conference on Machine Learning and Applications*, pages 112–116, 2013.
- [4] D. Ameller, C. Farre, X. Franch, D. Valerio, and A. Cassarino. Towards continuous software release planning. In *24th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 402–406, 2017.
- [5] M. Atas, A. Felfernig, M. Stettinger, and TNT. Tran. Beyond item recommendation: Using recommendations to stimulate knowledge sharing in group decisions. In *9th International Conference on Social Informatics (SocInfo 2017)*, pages 1–10, Oxford, UK, 2017.
- [6] A. Bagnall, V. Rayward-Smith, and I. Whitley. The next release problem. *Information and Software Technology*, 43:883–890, 2001.
- [7] G. Du and G. Ruhe. Does explanation improve the acceptance of decision support for product release planning? In *3rd International Symposium on Empirical Software Engineering and Measurement*, pages 56–68. IEEE, 2009.
- [8] J. Dyer. Multi attribute utility theory. *International Series in Operations Research and Management Science*, 78:265–292, 1997.
- [9] A. Felfernig, L. Boratto, M. Stettinger, and M. Tkalcic. *Group Recommender Systems – An Introduction*. Springer, 2018.
- [10] A. Felfernig, M. Schubert, and C. Zehentner. An Efficient Diagnosis Algorithm for Inconsistent Constraint Sets. *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing (AIEDAM)*, 26(1):175–184, 2012.
- [11] S. Huang. Designing utility-based recommender systems for e-commerce: Evaluation of preference elicitation methods. *Electronic Commerce Research and Applications*, 10(4):398–407, 2011.
- [12] U. Junker. QUICKXPLAIN: Preferred Explanations and Relaxations for Over-Constrained Problems. In *19th National Conference on AI (AAAI04)*, pages 167–172, San Jose, CA, 2004.
- [13] J. Karlsson and K. Ryan. A Cost-Value Approach for Prioritizing Requirements. *IEEE Software*, 14(5):67–74, 1997.
- [14] F. Kifetew, A. Susi, D. Mutante, A. Perini, A. Siena, and P. Busetta. Towards multi-decision-maker requirements prioritisation via multi-objective optimisation. In *Forum and Doctoral Consortium Papers Presented at the 29th International Conference on Advanced Information Systems Engineering (CAiSE’17)*, pages 137–144, Essen, Germany, 2017.
- [15] P. Laurent, J. Cleland-Huang, and C. Duan. Towards automated requirements triage. In *15th International REquirements Engineering Conference*, pages 131–140, Delhi, India, 2007.
- [16] L. Lehtola, M. Kauppinen, and S. Kujala. Requirements prioritization challenges in practice. pages 497–508, Kansai Science City, Japan, 2004.
- [17] A. Perini, A. Susi, and P. Avesani. A machine learning approach to software requirements prioritization. *IEEE Transactions on Software Engineering*, 39(4):445–461, 2013.
- [18] R. Reiter. A theory of diagnosis from first principles. *AI Journal*, 23(1):57–95, 1987.
- [19] G. Ruhe and M. Saliu. The art and science of software release planning. *IEEE Software*, 22(6):47–53, 2005.
- [20] T. Saaty and L. Vargas. *Models, Methods, Concepts & Applications of the Analytic Hierarchy Process*. Springer, 2000.
- [21] F. Shao, R. Peng, H. Lai, and B. Wang. DRank: A semi-automated requirements prioritization method based on preferences and dependencies. *The Journal of Systems and Software*, 126:141–156, 2017.
- [22] M. Stettinger, A. Felfernig, G. Leitner, S. Reiterer, and M. Jeran. Counteracting Serial Position Effects in the CHOICLA Group Decision Support Environment. In *20th ACM Conference on Intelligent User Interfaces (IUI2015)*, pages 148–157, Atlanta, Georgia, USA, 2015.
- [23] Y. Tian, D. Lo, X. Xia, and C. Sun. Automated prediction of bug report priority using multi-factor analysis. *Empirical Software Engineering*, 20(5):1354–1383, 2015.
- [24] P. Tonella, A. Susi, and F. Palma. Interactive requirements prioritization using a genetic algorithm. *Information and Software Technology*, 55(1):173–187, 2013.
- [25] E. Tsang. *Foundations of Constraint Satisfaction*. Academic Press, London, 1993.
- [26] K. Wiegers. *Software Requirements*. Microsoft Press, 2003.
- [27] J. Xuan, H. Jiang, Z. Ren, and Z. Luo. Solving the large scale next release problem with a backbone-based multilevel algorithm. *IEEE Transactions on Software Engineering*, 38(5):1195–1212, 2012.
- [28] J. Xuan, H. Jiang, Z. Ren, and W. Zou. Developer prioritization in bug repositories. In *34th International Conference on Software Engineering (ICSE)*, pages 25–35, Zürich, Switzerland, 2012.
- [29] Y. Zhang, M. Harman, and S. Mansouri. The multi-objective next release problem. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, pages 1129–1137. ACM, 2007.