

Rationale in Forum Discussions as Support for Software Upgrading of Online Course Platforms

Tiago L. P. Clementino¹, José Antônio B. Moura¹

¹Departamento de Sistemas e Computação – Universidade Federal de Campina Grande (UFCG)
Caixa Postal 10.106 – 58.429-900 – Campina Grande – PB – Brasil

tiagolucas@copin.ufcg.edu.br, antao@dsc.ufcg.edu.br

Abstract. *In online course software platforms, the forum may be an important instrument for feedback on pedagogical aspects. Students' and instructors' evaluations spontaneously shared in the forum – in terms of criticism, recommendations or requests – may also regard the platform's software tools. These evaluations resemble what is called “rationale”, a useful concept from Software Engineering. This paper presents an exploratory study about mining spontaneous user rationale in Massive Open Online Course forum discussions. Results provide qualitative evidence that rationale can be an instrument to collect requirements for the evolution of support software in online educational platforms.*

1. Introduction

Although Massive Open Online Course (MOOC) forum post logs have been explored by learning analytics to improve learning gains and pedagogical design [Khalil and Ebner 2016], MOOC platform software tools maintenance and upgrading seem to remain out of the scope of these efforts. To this end, this paper analyzes forum discussions in order to mine issues, criticisms, recommendations and requests that were spontaneously shared by MOOC participants. Such minings may be taken as a kind of “user rationale”.

“Rationale”, according to the Merriam-Webster dictionary, is “the explanation of controlling principles of opinion, belief or practice”; it is not the actual feedback from the user but her/his reasons behind it. Here we extend the definition by [Kurtanović and Maalej 2017] to include user spontaneous comments on any online course aspect which can be indirectly linked to software issues but also non-stimulated reaction to software issues explicitly and subsequent discussions as way of “user rationale”. This paper posits that online course user rationale, mined from forum discussions, may provide information for the improvement of the online course platform software.

For that, two 4-year apart sets of MOOC forum postings are considered: the first from courses offered by Stanford University (<https://www.stanford.edu/>) on the edX (<https://edx.org/>) MOOC platform from 2013 to 2014; and the second, postings from 1,034 courses offered by several providers worldwide on the same edX platform in November 2018, marked as “current” or “archived” (corresponding to a final total of 86 courses), and with some specific channel (discussion) for criticism or suggestions. Then, the edX release notes (<https://edx.readthedocs.io/projects/edx-release-notes/en/latest/index.html>) are examined in search of some correlation between mined rationale decrease and functionality additions or bug fixes in new edX platform versions.

Lastly, nine MOOC users with teaching experiences were interviewed as part of further validation effort.

2. Background and Related Work

A literature search yielded no works on user rationale from educational platforms. Two related research areas have reports on approaches which may be adapted for that purpose.

2.1. MOOC and Forum Analytics

MOOC forum analytics represent an important instrument in pursuing quality and feasible online education [Almatrafi and Johri 2018, Romero and Ventura 2017]. Although the exploration of spontaneous user rationale from MOOC forums has received scanty attention in the literature, some strategies from other areas may be used to inspire automatic mining of spontaneous user rationale. [Cui and Wise 2015] and [Wise et al. 2017] researched ways to mine content-related postings on discussions. The first work describes a classifier that identifies content-related question and answer discussions where instructor participate. The second, also based on a content-related classifier, identifies that postings strongly related to domain seem to have some recurrent text patterns non-related to the domain. [Elalfy et al. 2018] propose a way to point out the best answers. Also, ongoing research on useful spontaneous user rationale in forums may benefit from works on urgent staff intervention [Almatrafi et al. 2018].

2.2. Rationale

In software engineering, *rationale* represents the decision process performed before and after software development [Lee 1997]. On the other hand, spontaneous user rationale refers to rationale about the user experience [Kurtanović and Maalej 2017]. Here we use the adjective “spontaneous” because the MOOC forum is not explicitly meant as a channel for software reviews as it is the case of application (app) store channels. Algorithms for mining and analyzing app store – and even social networks – as a feedback channels are proposed in [Williams 2019]. Williams’ work differs from the one here for software reviews are a proper user rationale channel rather than a spontaneous one as in an educational environment – which requires more subtle and comprehensive analyses of root-causes.

Forum and education analytics have been researched [Almatrafi and Johri 2018, Romero and Ventura 2017], but qualitative data analysis and software user rationale extraction have not been addressed. Much has been done on rationale mining. For instance, [Alkadhi et al. 2018] offered an empirical study on the discussion of rationale by Open Source Software (OSS) geographically spread developers in Internet Relay Chat (IRC) channels and explored the possibility of extracting that rationale automatically from IRC messages. Another work considered rationale in IRC messages clustering to distinguish on-topic from off-topic, and thus, reduce useless discussions [Chowdhury and Hindle 2015].

[Kurtanović and Maalej 2017] presents a full definition of user rationale in the context of app reviews, besides a complete methodology on how to extract and classify them. [Dhinakaran et al. 2018] also discusses sophisticated machine learning approaches for mining rationale. [Kurtanović and Maalej 2018] presents an approach to investigate

the arguments and justifications behind decisions. In addition, [Jha and Mahmoud 2017] also considers app review requirement mining and [Khan et al. 2019] surveys the literature on crowd-based requirements engineering research, which mentions “user rationale” as a promising research area. None of these relevant studies however, considers “user rationale” in collaborative platforms as a source of software requirements as done here.

The work here is more closely related to those of [Villanes et al. 2017]; [Alkadhi et al. 2017]; and, [Chatterjee et al. 2019]. The first work focused on exploring discussion forums using a machine learning algorithm to analyze rationale on Android (<https://www.android.com/>) tests from the Stackoverflow (<https://stackoverflow.com/>) collaborative platform. The second performed an exploratory study on the frequency of rationale in chat messages by development teams, and on ways for automatic extraction and their feasibility. The third considered the usefulness and challenges of using Q&A discussions for supporting software evolution. Instead of studying rationale behind software developers’ communications, this paper deals with extracting useful rationale for educational platform software development from non-IT, course content-based communications of students. As such, this paper complements the cited related work.

3. Research Design and Execution

This section summarizes the research question, methodology and the data collection procedure used to mine and analyze rationale from MOOC forum postings. Complementary details and results for this and Sections 4 to 6 appear in <https://bit.ly/2TMLWvp>.

3.1. Research Questions and Aims

This work aims at gathering evidence that MOOC forum postings embed users’ rationale that may serve to elicit requirements for maintaining and/or upgrading the platform software. The research question of interest may be formulated as:

- RQ: Can rationale underlying forum postings support the maintenance and upgrading of educational software platforms?

3.2. Methodology

The methodology applied to answer RQ involves three phases:

1. Mine rationale from forum logs of a given MOOC platform over a period of time;
2. Find evidence the mined rationale correlates with changes in release notes for the given MOOC platform software; and,
3. Interview experienced MOOC educators to validate the relevance of the mined rationale they perceive as an instrument for eliciting requirements of the platform software they used.

3.3. Gathering of Quantitative Data

In phase 1, two 4-year apart logs of MOOC forum postings are considered: the first from 12 courses offered by Stanford University on the edX MOOC platform from 2013 to 2014; and, the second, with postings from 1,074 courses which were being offered by several providers worldwide also on the edX platform in November 2018. Henceforth, the data collected from the first set will be labelled as **D1** and those from the second, as **D2**. Rationale was manually mined from **D1** and **D2**.

D1 encompasses three knowledge areas: Humanities/Sciences, Medicine and Education, adding up to 29,604 forum postings grouped into 22,804 discussions. These discussions can further be characterized as follows: 20,268 have only one post; 2,208 with 2 to 5 postings; 179 with 6 to 10 postings; and, 48 with more than 11 postings.

D2 is not composed of data from MOOC forum log automatically recorded, but rather, by findings collected directly from the edX platform interface – totalling 1074 courses of the most varied knowledge areas. From these, only those courses which maintain an open channel for complaints and suggestions were used. Hence, only 87 courses contributed to **D2**.

3.4. Rationale Extraction

D1 and **D2** were mined for rationale separately as summarized next.

Mining D1: **D1** was mined manually using a set of initial keywords that evolved by snowballing. Searches were performed based on the current set of keywords and by means of the results, new keywords were added to this set. The initial and final sets of keywords were: **Initial Keywords** – *platform, thread, post, discuss(ion), forum, navigat(e)(ion), user, student, MOOC*; **Final Keywords** – *helpful, platform, module, review(er), learn(ing)(er), thread, post, discuss(ion), forum, navigat(e)(ion), user, student, friendly, MOOC, Coursera, (Open)edX, interface, comment, idea, categor(y)(ies), frustrat(e)(ing)(ion), disappoint(ing), maint(aining)(einance), point, automatic(ally), feedback, system, (dis)(en)courage, (un)read, assess(ment)*.

For a mined user rationale to be “accepted” as possibly supporting the maintenance and/or upgrade of the educational platform used, it had to concern an actual use case or user story – as defined in agile software development methods such as XP or Scrum [Beck and Gamma 2000, Schwaber 1997]. Accepted rationale were attributed a score:

- 1 point: If the posting user demonstrated some technical knowledge about the issue being addressed;
- 1 point: If it appears in other users’ rationale (per appearance);
- 2 point: It is a criticism;
- 4 point: If it brings a suggestion for the maintenance and/or the upgrade of the platform software.

Mining D2: The criteria that a given course had to meet to be selected from **D2** were: *be marked as current or archived, be open, have a forum and have an official forum channel for criticisms and suggestions*. The accepted, mined rationale in the case of **D2** met the same criteria and point awards as in the case of **D1**. A total of 47 postings was gathered as rationale from 87 courses.

4. Phase 1 Results

The result of phase 1 of the methodology is composed of accepted, mined rationale embedded in postings. For illustration, Table 1 presents two user rationale mined from **D1**

and their respective requirement specification structured according to software maintenance procedures, whose format we adapted from [Sommerville 2011]. The postings in Table 1 represent rationale about Usability (requirement 1) and Thread Exposure (requirement 2). Notice that requirement 1 is a suggestion, but requirement 2 is a criticism. As shown, the most complete rationale is one that brings arguments based on suggestions but not only on criticism. Criticism postings usually have gaps which must be filled by software requirements engineers.

Table 2 categorizes rationale mined from *D1* and *D2*. As it may be observed from Table 2, the pattern of rationale changed over time, from **D1** to **D2**. Specifically, demand for maintenance/upgrade in the Gamification rationale category diminishes while in the Peer Assessment category increases; the remaining categories present almost the same pattern of rationale in both datasets. The remaining phases of the methodology give us material to analyze this observations in more detail.

Table 1. Examples of rationale requirement specifications mined from D1 and structured according to adapted format from [Sommerville 2011]

Requirement 1

Description	Separate buttons Save and Final Check in questions of quizzes. Each button sends a different command to Server. Save button sends a command just to save the current status and Final Check sends the commands “Save” and “Exit” simultaneously.
Input	Any question of quiz total or partially filled
Action	A different command for each button. Save button sends a command just to save the current status and Final Check sends both commands “Save” and “Exit”.
Precondition	The user has chosen to answer the quiz and s/he has met previous credit and payment criteria.
Postcondition	The state of question of quiz is updated to “Saved” and, if the Final Check button has been pressed, the quiz process is finished and the interface is changed from “open quiz” to the “completed quiz”.
User Rationale^a	<i>“It would be helpful if each question had a separate Save and Final Check button”</i>
Type of Rationale	Suggestion

Requirement 2

Description	The user shall be able to follow the threads which s/he posts and anyone which s/he marks to follow. It shall be done using daily emails with new postings notifications.
Input	New postings written by the user or some thread marked as following.
Action	Update the status of that respective thread as following and an email shall be sent notifying every new post of that thread.
Precondition	The user opened a thread in her interface
Postcondition	The user posted in the thread or marked it as “following”. Whichever situation shall change the status of thread to “following”
User Rationale^a	<i>“If I participate in a discussion, it’s nearly impossible for me to find it again later. I don’t get notified if anyone responds to a post of mine, and I can’t find a way to follow a discussion that interests me.”</i>
Type of Rationale	Criticism

^a This line transcripts the user spontaneous rationale directly from the dataset, without any edition.

5. Phase 2 Results

The changes (increases or decreases) in rationale over time happen because some demand may met by new functionality or fixes in new releases of the MOOC platform software. In order to find which of those rationale categories were incorporated in the platform

Table 2. Frequency of rationale and its nature.

Rationale Category	D1		D2	
	Suggestions	Criticisms	Suggestions	Criticisms
Peer assessment	2	1	9	1
Thread Classification ^a	2	0	0	1
Thread Exposure ^b	5	0	3	2
Gamification	9	1	2	0
Feedback ^c	2	4	2	3
Multimedia	0	0	0	1
Plagiarism Checker	1	0	0	0
Usability Points	17	2	11	4
Group Learning ^d	1	0	0	0

a According to title and options describing the subject of discussion.

b It refers to the order, facilities to follow the discussion, the presentation interface, etc.

c Manual or automatic.

d Arrangement of students into collaborative groups instead of letting collaboration involve them all

software, the release notes for the edX platform were examined. It is useful to our research question to analyse if rationale in forum actually reflects real demands of users by checking whether the incorporation of those demands ceases the rationale.

Some of the edX rationale from **D1** was solved and decreased or disappeared by the time D2 was mined. That is the case of rationale about Gamification and Thread Exposure. This last category, although present in **D2**, had its subject of rationale changed over time. Example of correlation between a change in release notes and rationale from **D1** is given in Table 3. This rationale is no longer present in **D2**.

Table 3. Example of edX release note fulfilling a demand in a rationale category (demonstrated here by two rationale) in D1

Release date	Week of 18 April 2016
Release Title	Sorting by Date for Discussion Posts
Release description	A problem with the method used to sort discussion posts in web browsers is now fixed. Previously, when a learner sorted discussion posts by recent activity in a web browser, the discussion posts were sorted by creation date rather than last activity date. Discussion posts are now sorted correctly. (TNL-3818).
Rationale 01	“The next time Stanford talks to the people at OpenEdX, I suggest asking that they design the ability to sort threads by those which have unread posts. I, for one, follow a very large number and have to keep viewing more threads to make sure there aren’t some with posts from latecomers. Thanks!”
Rationale 02	“I didn’t think of this when filling out the exit survey, so I’ll post it here. I would like an easier (or automatic) way to follow posts and comments – you can manually star a topic you want to follow, but there is no easy way to tell if someone has replied to one of your posts within a topic, especially if you forgot to star it.”
Rationale Category	“Thread exposure”.

6. Phase 3 Results

To analyze the actual relevance of the mined rationale, nine professional from different knowledge areas with teaching experience were interviewed. Each interviewee must have

Table 4. Interviewees' profiles.

	Age	Highest degree	years past \approx	Performance ^a	Platform	Work
1	30-39	Bachelor	0.3	C	Saitestore	Medicine
2	30-39	Bachelor	1.0	B+	AVASUS	Medicine
3	21-29	Bachelor	1.0	B+	Unasus	Medicine
4	30-39	Masters	1.5	B	Coursera	Design
5	30-39	Bachelor	1.0	C	Futurelearn	Law
6	30-39	Phd	2.0	B	Coursera	Design
7	30-39	Bachelor	2.0	B+	edX	Plastic Arts
8	60+	Phd	0.6	A+	Coursera	Electric Engineering
9	30-39	Bachelor	2.0	A	edX	Culture

^a Self-evaluated performance on the last finished MOOC – Which was used as parameter.

had significant hands-on MOOC experience dating back two years at the most. Interviewees were recruited by the snowballing method, where the first participant indicates the next and so on. Interviews were conducted either in person or through video conference. Each interviewee must have finished at least two MOOCs of one year duration or longer. (That was required to show commitment by the participant towards the course and to increase the odds that s/he worked with more pedagogical tools and resources.) All participants came from Brazil, work in different research areas and are 30 years old on average with a standard deviation, $\sigma = 5$ years. Table 4 presents other demographics.

6.1. Questionnaire and Interviews

A questionnaire in Portuguese based on the rationale categories resulting from phase 1 was applied using Google Forms (<https://www.google.com/intl/pt-BR/forms/about/>) platform and is available from <http://bit.ly/2ZfOrrw>. It was asked that it be answered online 24 hours before the interview itself; the answered questionnaire was then used as a guide for interviewing. (The 24 hour delay granted time for the interviewee to reflect on the rationale categories.) Each question inquires about the interviewee's impressions about a given category's effectiveness and perspectives. Answers are in a 5-level Likert scale [Likert 1932, Boone and Boone 2012]: 1 – strongly approve; 2 – approve; 3 – undecided; 4 – disapprove; and, 5 – strongly disapprove. Some of questions use a variation of this scale, without loss of amplitude. The questionnaire is thus quantitative. The definitions of each category in Table 2 were offered during the interview only, but some questions in the questionnaire tried to elicit the interviewee's basic knowledge about that category. Remaining questions were meant to promote discussions about the rationale categories. Interviews are thus semi-structured and were conducted by the authors.

The entire procedure of interviewing, from the elaboration of the questionnaire to the actual meeting with interviewees, followed by analyses, was conducted based on the grounded theory method of Strauss and Corbin [Strauss and Corbin 1990]. These interviews were conducted assuming that *the interviewees would reproduce spontaneously the most recurrent rationale mined in forums*, and lasted between 40 to 90 minutes each and later interviews tended to last longer and were more fruitful.

The interviews focused on stimulating awareness, on interviewees, of rationale for specifying educational software tools, asking for their experiences, real use cases, com-

plaints and suggestions. To do so, the rationale mined in **D1** was compared to that from **D2** during interviews. The interviewees were profiled by asking about their behavior on regular education – in the classroom – and their last MOOC experience (see Table 4) to match educational profiles and opinions. All interviews were recorded, with the interviewee's consent, for posterior analyses.

6.2. Results

Different perspectives were identified spreading over different knowledge areas: “Medicine”, “Social Sciences” and “Exact Sciences”. For instance, automated *Feedback* models cause a lot of resistance in Social Sciences educators, as perceived in the interviews. This observation is aligned with the rationale mined in both data sets whose amount was approximately constant, relevant and close related to courses in this area. Oddly, in despite of *Multimedia* popularity, it was in user rationale just once in **D2**. This observation was confirmed by the interviewees, who did not offer any rationale for that category but only neutral opinions typically. On the other hand *Usability* category stimulates so much rationale that it does not seem feasible for software developers to work on them all. Much of it seems related to the learning curve but not to real issues. These examples support the use of rationale as a requirement source for collaborative platforms as confirmed by comments of interviewees on the mined rationale.

Besides *Usability*, the unique areas which led to rationale by interviewees were *Thread Exposure* and *Peer Assessment*. These rationale, represented mainly by criticism, may be exemplified by “*Maybe a meta assessment works to fix low experience peer assessments*” for *Peer assessment*; and, “*Threads with some high up-voted post would be exposed first*” for *Thread Exposure*.

Peer Assessment was praised by interviewees of Exact (such as Engineering) and Medical Sciences backgrounds for its feasibility and scalability. Social Sciences interviewees however, think this method may be used only in intermediate assessments due to its “imprecision”. They recommend caution with its application. Interviewees were vague about *Group Learning* and *Plagiarism Checker*, reinforcing them as outliers. On the other hand, one of them disagreed with the use of *Gamification*, probably because some gamified tasks were used to grade students in her/his course. Both theses findings match mined rationale.

6.3. Threats to Validity

A small number of interviewees bias the study, even though, for a viable qualitative study, the set of validators in the analysis must remain small. However, the manual mining method applied in phase 1 of the methodology threatens the main phase of the study. This threat was accepted keeping in mind that this study seeks evidence rationale “may” work to elicit requirements for maintenance and upgrading of MOOC software.

Although the typical survey at the end of MOOCs may replace user spontaneous rationale mining, such courses tend to have high drop out rates and many users who use the forum facility to share rationale on software do not persevere to the end (of their courses). Mining rationale from forums offers a continuous source of educational platform software requirements elicitation.

7. Conclusions and Ongoing Work

This paper reported on the investigation of the potential of mining of spontaneous user rationale in forums postings as a source / strategy for eliciting maintenance and upgrade requirements of MOOC software. Although only a few MOOCs were examined, results were relevant and in some cases, surprising. Our observations found mined rationale, according to the opinions of interviewees, to be a educational software requirements source.

Specifically to address *Usability* issues, this strategy seems to be a fertile requirement gathering method. Other requirement categories are also promising, like *Thread Exposure* and *Classification*. *Multimedia* issues, surprisingly, received less attention by users in their complaints. However, users (in this case, MOOC students) can only rationalize about what they know and have seen occasionally. Consequently, many breakthroughs which they are not yet aware about, obviously cannot appear in their postings.

For ongoing work, a method for automatically extracting and analyzing user rationale from MOOC forums would provide a better understanding. In addition, to compare (automatically, this time) mined rationale to commit series of multiple MOOC software platforms instead of just release notes, could provide better conclusions about rationale usefulness for eliciting requirements of new versions of these platforms.

References

- Alkadhi, R., Lata, T., Guzman, E., and Bruegge, B. (2017). Rationale in development chat messages: an exploratory study. In *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, pages 436–446. IEEE.
- Alkadhi, R., Nonnenmacher, M., Guzman, E., and Bruegge, B. (2018). How do developers discuss rationale? In *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 357–369. IEEE.
- Almatrafi, O. and Johri, A. (2018). Systematic review of discussion forums in massive open online courses (moocs). *IEEE Transactions on Learning Technologies*.
- Almatrafi, O., Johri, A., and Rangwala, H. (2018). Needle in a haystack: Identifying learner posts that require urgent response in mooc discussion forums. *Computers & Education*, 118:1–9.
- Beck, K. and Gamma, E. (2000). *Extreme programming explained: embrace change*. addison-wesley professional.
- Boone, H. N. and Boone, D. A. (2012). Analyzing likert data. *Journal of extension*, 50(2):1–5.
- Chatterjee, P., Damevski, K., Pollock, L., Augustine, V., and Kraft, N. A. (2019). Exploratory study of slack q&a chats as a mining source for software engineering tools. In *Proceedings of the 16th International Conference on Mining Software Repositories (MSR’19)*. ACM.
- Chowdhury, S. A. and Hindle, A. (2015). Mining stackoverflow to filter out off-topic irc discussion. In *Proceedings of the 12th Working Conference on Mining Software Repositories*, pages 422–425. IEEE Press.

- Cui, Y. and Wise, A. F. (2015). Identifying content-related threads in mooc discussion forums. In *Proceedings of the 2nd ACM Conference on Learning@ Scale*, pages 299–303.
- Dhinakaran, V. T., Pulle, R., Ajmeri, N., and Murukannaiah, P. K. (2018). App review analysis via active learning: reducing supervision effort without compromising classification accuracy. In *2018 IEEE 26th International Requirements Engineering Conference (RE)*, pages 170–181. IEEE.
- Elalfy, D., Gad, W., and Ismail, R. (2018). A hybrid model to predict best answers in question answering communities. *Egyptian informatics journal*, 19(1):21–31.
- Jha, N. and Mahmoud, A. (2017). Mining user requirements from application store reviews using frame semantics. In *International working conference on requirements engineering: Foundation for software quality*, pages 273–287. Springer.
- Khalil, M. and Ebner, M. (2016). What massive open online course (mooc) stakeholders can learn from learning analytics? *Learning, design, and technology: An international compendium of theory, research, practice, and policy*, pages 1–30.
- Khan, J. A., Liu, L., Wen, L., and Ali, R. (2019). Crowd intelligence in requirements engineering: Current status and future directions. In *International working conference on requirements engineering: Foundation for software quality*, pages 245–261. Springer.
- Kurtanović, Z. and Maalej, W. (2017). Mining user rationale from software reviews. In *2017 IEEE 25th International Requirements Engineering Conference (RE)*, pages 61–70. IEEE.
- Kurtanović, Z. and Maalej, W. (2018). On user rationale in software engineering. *Requirements Engineering*, 23(3):357–379.
- Lee, J. (1997). Design rationale systems: understanding the issues. *IEEE expert*, 12(3):78–85.
- Likert, R. (1932). A technique for the measurement of attitudes. *Archives of psychology*.
- Romero, C. and Ventura, S. (2017). Educational data science in massive open online courses. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 7(1):e1187.
- Schwaber, K. (1997). Scrum development process. In *Business object design and implementation*, pages 117–134. Springer.
- Sommerville, I. (2011). Software engineering 9th edition. ISBN-10, 137035152.
- Strauss, A. and Corbin, J. (1990). *Basics of qualitative research*. Sage publications.
- Villanes, I. K., Ascate, S. M., Gomes, J., and Dias-Neto, A. C. (2017). What are software engineers asking about android testing on stack overflow? In *Proceedings of the 31st Brazilian Symposium on Software Engineering*, pages 104–113. ACM.
- Williams, G. S. (2019). Modeling crowd feedback in the mobile app market.
- Wise, A. F., Cui, Y., Jin, W., and Vytasek, J. (2017). Mining for gold: Identifying content-related “mooc” discussion threads across domains through linguistic modeling. *Internet and Higher Education*, 32:11–28.