

On the impact of pull request decisions on future contributions

Damien Legay
Software Engineering Lab
University of Mons
Mons, Belgium
damien.legay@umons.ac.be

Alexandre Decan
Software Engineering Lab
University of Mons
Mons, Belgium
alexandre.decan@umons.ac.be

Tom Mens
Software Engineering Lab
University of Mons
Mons, Belgium
tom.mens@umons.ac.be

Abstract—The pull-based development process has become prevalent on platforms such as GitHub as a form of distributed software development. Potential contributors can create and submit a set of changes to a software project through pull requests. These changes can be accepted, discussed or rejected by the maintainers of the software project, and can influence further contribution proposals. As such, it is important to examine the practices that encourage contributors to a project to submit pull requests. Specifically, we consider the impact of prior pull requests on the acceptance or rejection of subsequent pull requests. We also consider the potential effect of rejecting or ignoring pull requests on further contributions. In this preliminary research, we study three large projects on GitHub, using pull request data obtained through the GitHub API, and we perform empirical analyses to investigate the above questions. Our results show that continued contribution to a project is correlated with higher pull request acceptance rates and that pull request rejections lead to fewer future contributions.

I. INTRODUCTION

The turn of the century saw the rise of version control systems (VCS) to support large-scale software engineering projects. Centralised VCS (e.g. CVS and Subversion) allow developers to share a common repository. Decentralised ones (e.g., Mercurial and git) allow each developer to own a local copy of the repository containing the full change history. This enables collaborative (often geographically distributed) software development on an hitherto unmatched scale. It has given birth to extremely popular online hosting platforms such as GitHub, BitBucket and Mozdev, allowing thousands of people to remotely work together on the same projects. These platforms provide additional features on top of their underlying VCS to further support distributed collaborative development. Examples of such features are issue tracking, code review, integrated discussions, team management, documentation & wiki and integration with external tools.

Today, git has become the most popular distributed VCS by a large margin¹. It will thereby be the focus of our current research. git supports two types of development processes: the *shared repository* approach, where all contributors are given write access to the central repository and can therefore con-

tribute to the project directly; and the *pull request* (PR) approach where only project *integrators* are allowed to do so.

With the PR approach, external contributions are managed indirectly: would-be contributors create a fork of the repository and, once they have addressed an issue or lack in the project, they request for their modifications to be “pulled” to the repository by submitting a *pull request*. The project integrators can decide to approve these PRs, which are then merged into the main project’s codebase. PRs are extremely valuable, as they represent a major part of the project’s continued evolution and expansion. It is therefore important to incentivise people to create pull requests, thereby contributing to the project. Previous studies have attempted to identify the factors influencing whether and when a PR will be merged [1]–[3].

We expand upon this work, by focusing on determining those patterns of PR-acceptance behaviour that are indicative of continued contribution. Our working hypothesis is that people contributing to a project repository through PRs may get demotivated (and hence stop contributing) if their submitted PRs get rejected too often, or if too many of them are left open without any decision to merge them. Evolutionary insights in such phenomena may help us to understand which of such factors tend to dissuade people to keep contributing to a given project. To this extent, we quantitatively study the following research questions using techniques based on survival analysis:

RQ₁: How are PR acceptance and rejection rates influenced by previous PRs? As a contributor accrues familiarity with a project, he becomes more able to contribute effectively, which we expect to result in a lower PR rejection rate. Similarly, as integrators become more acquainted with a contributor, they may develop a favourable bias towards his PRs, further decreasing rejection rates.

RQ₂: To which extent does PR acceptance or rejection influence further contributions? When his PRs are rejected, a developer could become discouraged and stop, temporarily or permanently, contributing to the project as a result.

RQ₃: To which extent do PRs left open influence further contributions? A PR is sometimes left open for a long period, neither rejected nor merged into the core project. We posit this may constitute a form of “soft” rejection, wherein the integrators want to avoid alienating the contributor but do not want to merge the PR. Seeing a large number of untreated

¹For anecdotal evidence, based on a 2016 survey with 881 votes, 87% of responders identified git as their VCS of choice <https://rthodcode.com/insights/version-control-systems-2016>

PRs may send an implicit message to potential contributors that the project integrators are unwilling or unable to process the volume of contributions they receive, and, therefore, that their participation to the project would not be valued or useful.

To provide preliminary evidence for these RQs, we carry out an empirical analysis on a large number of PRs in three large, popular and long-lived projects on GitHub. We focus on GitHub because it is undoubtedly one of the largest and most active online hosting services for git projects.

II. RELATED WORK

Several researchers have studied aspects related to the PR-based software development process, either qualitatively or quantitatively. Gousios and Zaidman proposed a PR dataset [4] including 900 projects and 350,000 PRs extracted using GHTorrent. Through a mixed-method analysis of 291 GitHub projects, Gousios et al. [1] established that the PR-based development approach is used as frequently as the shared repository approach on GitHub. They observed that most PRs are short, receive few comments and are processed quickly. They also found that most PR rejections are due to the distributed nature of the pull-based process (e.g., PRs that are already obsolete upon creation).

In a follow-up work [3], they interviewed 645 contributors to examine their work practices and identify the challenges they face. They found that while contributors tend to check if their intended contribution is already covered, they do not communicate their intended contributions. Interviewed contributors outlined that poor responsiveness on the part of integrators could be a barrier to attracting or retaining contributors. Contributors also stated that it is hard to accept rejection of their PRs, as rejected PRs could harm their reputation as developers. Conversely, it is hard for integrators to explain the reasons for rejecting PRs. Rejecting a PR without alienating its contributor was already identified as a challenge of the PR-based model [2]. In that paper, they evaluated PRs from an integrator's point of view by interviewing 749 project integrators in order to understand which criteria are used to determine the quality of a PR and how they prioritise the evaluation of contributions. They found that most integrators decide to merge PRs based on project's objectives, their quality as measured by compliance to the project guidelines, test coverage and passing continuous integration checks.

Yu et al. [5] studied the factors that contribute to latency in PR reviews, defining this latency as the "time interval between pull request creation and closing date". They found that PR latency is mainly affected by process-related factors such as whether a PR was assigned to a specific reviewer or not. They also found that continuous integration is a dominant factor in PR latency.

Rahman and Roy [6] categorised the technical issues discussed in PR comments and analysed information about projects and developers to obtain insights into PR acceptance or rejection. They discovered that the rate of PR rejection is highly correlated to the programming language used (e.g., Java PRs are more frequently rejected than PRs for the C

programming language), the application domain of the project (e.g., the database application domain sees fewer merged PRs than the IDE domain), the maturity of a project (older projects accept fewer PRs) and the number of developers on the project.

Tsay et al. [7] explored both technical and social factors that contribute to acceptance of PRs. They found that, although technical factors like the presence of tests in the PR and a small number of lines changed contribute to a higher probability of acceptance, social factors, such as whether the contributor follows the user that closes the PR, had stronger associations to PR acceptance than technical ones.

Terrel et al. [8] established that PR acceptance is subject to a bias against women, when their gender is identifiable. Rastogi et al. [9] built upon the factors identified in [1] and [7], adding information about the geographical location of contributors and integrators. They conclude that PR acceptance rate is higher when both contributor and integrator are from the same country, with the exception of India, and that contributors from some countries (e.g., Switzerland and Japan) see their contributions more frequently accepted than contributors from other countries (e.g., China and Germany).

III. METHODOLOGY

The main goal of our research is to study the longevity of PR-based contributions to large open source software projects. We focus on software development through GitHub, the largest and most active online hosting service for git projects. As of 2018-09-30, GitHub has hosted 96M+ repositories, 31M+ developers, and 200M+ PRs and about one third of these repositories and PRs were created in the last 12 months.²

For this exploratory research, we selected three case studies of large open source git projects on GitHub. These projects have been obtained by convenience sampling. This method is acceptable for getting preliminary research insights, and will be replaced in a later phase to obtain a bigger corpus that covers a larger set of relevant projects.

The main criteria for our selected sample were that the projects should be representative of a typical PR-based software development process. To do so, the projects needed to be mature (i.e., have a time span of several years), have an active development history with a huge number of commits and contributors, and of course contain a very large number of PRs, in order to be able to derive statistically significant results from their analysis. In addition to this, we selected projects written in three different languages to ensure sufficient diversity. The three selected projects are `ansible`, `rails` and `kubernetes`. Some of their characteristics are shown in Table I.

Because we have observed problems of missing or inconsistent data when using GHTorrent, we decided to extract the PR data of the selected projects from GitHub repositories through the GitHub API directly. For each PR, the data contains information about the PR creation date, its status (accepted, rejected or open), its closing date (for accepted and rejected PRs), the GitHub ID of its author and its PR number. This

²<https://octoverse.github.com>

PR number corresponds to a chronological ordering of issues opened in the repository, of which PRs are a subset.

TABLE I: Project repository characteristics on 24/10/2018

repository	language	start year	#contributors	#commits	#PR
ansible	Python	2012	3930	40k	27k
rails	Ruby	2010	3683	70k	22k
kubernetes	Go	2014	1861	71k	42k

IV. PRELIMINARY RESULTS

RQ₁: How are PR acceptance and rejection rates influenced by previous PRs?

To answer *RQ₁*, we examined whether repeat contributions impact a contributor’s PR acceptance rate. To that effect, for each repository we analysed the PR acceptance rate in function of the number of submitted PRs by each contributor.

Figure 1 displays, for each positive integer threshold x between 1 and 250, the PR acceptance rate (blue curve) and rejection rate (orange curve) considering the first x PRs of each contributor only, thereby discarding contributors having less than x PRs. The green curve shows the number of contributors having submitted at least x PRs. Thresholds above 250 are excluded due to the low fraction of contributors having that many submissions: 0.33% for `ansible`, 0.15% for `rails` and 1.23% for `kubernetes`.

One can observe in all three examined project repositories that, as contributors submit more PRs, their acceptance rates increase significantly. Over the first 50 PRs, we observe a rise from 54.2% to 80.0% for `ansible`, from 61.3% to 81.4% for `rails`, and from 49.1% to 74.3% for `kubernetes`. Beyond the 50 first PRs, all three projects saw continuous increase in PR acceptance rates as contributors submitted more PRs to them.

These results agree with prior findings by Tsay et al [7] and Gousios et al [1], but are to be nuanced, given the rapid decrease in number of contributors as the threshold x increases. As a consequence, in Figure 2 we looked at the PR acceptance rate of all contributors, excluding the few that made over 250 contributions. We observe that, while contributors with a high number of PRs tend to have a consistently high PR acceptance rate, the behaviour for contributors with few PRs is quite unpredictable: they can have either low or high acceptance rates. Therefore, although the number of previous PRs influences acceptance rate, this can only be verified starting from a certain threshold of PRs, below which no conclusion can be reached as to whether such an influence exists.

RQ₂: To which extent does PR acceptance or rejection influence further contributions?

While related work (e.g., [1]) has studied the impact of PR acceptance rate on future PR decision time, *RQ₂* focuses on the impact of PR acceptance rate on the likelihood of making further PRs. To do so, we compared the probability to contribute again after either a rejected or an accepted PR. The results are presented in Table II. In all three considered

TABLE II: Likelihood to contribute again

repository	after acceptance	after rejection
ansible	85.6%	73.0%
rails	83.9%	69.2%
kubernetes	95.5%	88.2%

projects, contributors are more likely to make subsequent PRs if their prior PRs were accepted.

We then used the statistical technique of survival analysis (a.k.a. event history analysis) [10]. Given a specific “event of interest” (in our case: acceptance or rejection of a PR), survival analysis models the “time to event” data during a given observation period. Survival functions model the survival rate, i.e., the expected time duration until the event of interest occurs. The models take into account the “censoring” of some observed subjects, either because they enter or leave the study during the observation period, or because the event of interest was not observed for them during the observation period. A common non-parametric statistic used to estimate survival functions is the Kaplan-Meier estimator [11].

We performed an analysis of the survival probability to submit a new PR in function of the time elapsed since the latest submission (at that time) of a PR by the same contributor. In order to assess if the PR acceptance rate influences the delay for a contributor to submit new PRs, we considered three acceptance rate classes: $[0, 33\%[$, $[33\%, 67\%[$ and $[67\%, 100\%[$. The survival curves are shown in Figure 3. We observe that the survival probability is higher for classes of higher acceptance rate, regardless of the considered projects. For instance, after ten days, the probability to submit a new PR is 72.2% in `Ansible` if the acceptance rate is over 67%, while this probability drops to 52.9% if the acceptance rate is between 33% and 67%, and even to 31.5% if the acceptance rate is below 33%. Similar patterns can be observed for the two other projects. We carried out pairwise log-rank tests to compare whether statistically significant differences could be found between the survival curves. The differences were statistically confirmed at $\alpha = 0.01$ (after a Bonferroni correction [12]), i.e., the null hypotheses, assuming that the survival curves for different acceptance rate classes are the same, were rejected.

We performed a proportional hazards regression based on Cox regression to determine to which extent the acceptance rate impacts the probability of further contribution [13]. The Cox regression is a method for investigating the effect of several variables upon a specified event’s hazard rate. For this analysis, we included the following factors: the acceptance rate of all prior PRs by the same contributor; the number of prior PRs made by this contributor, and the time elapsed since the contributor’s first PR (the contributor’s *age*).

TABLE III: Influence of acceptance rate, number of PRs, and contributor age on the time required to submit a new PR.

repository	regression coefficients			concordance
	acceptance rate	#prior PRs	contributor age	
ansible	0.5481	0.0038	-0.0009	0.689
rails	0.4630	0.0047	-0.0015	0.728
kubernetes	0.7455	0.0031	-0.0016	0.637

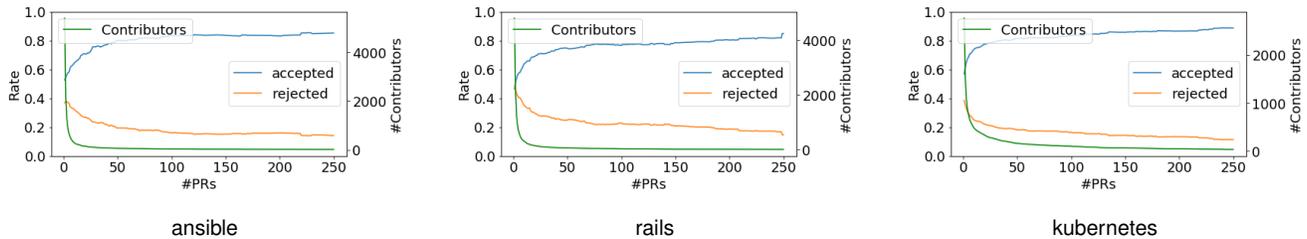


Fig. 1: Acceptance rate of the first x PRs of each contributor.

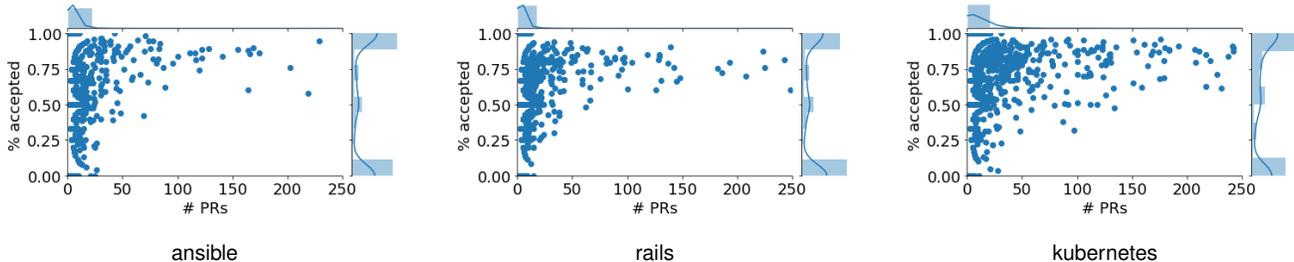


Fig. 2: Acceptance rate of all PRs by contributor.

Table III summarizes the results we obtained. The concordance (fourth column) provides the goodness of fit of the model. It is comprised between 0 (perfect anti-concordance) and 1 (perfect concordance). The table also reports the regression coefficients for the three considered factors. These coefficients measure the magnitude of the impact of the aforementioned factors on the probability to submit a new PR. All these coefficients are statistically significant ($p < 0.01$ after Bonferroni correction). Their values signify that an increase of one increment (10% in acceptance rate, 1 prior PR or 1 day since the contributor’s first PR) multiplies the probability to submit a PR by a factor $e^{\text{coefficient}}$. For instance, in the case of Kubernetes: an increase of 10% in PR acceptance rate modifies the probability to submit a new PR by a factor 1.0774, each prior PR by 1.0032 and each day since the contributor’s first PR by 0.9984.

RQ₃: To which extent do PRs left open influence further contributions?

To provide insight into RQ_3 , we looked at the proportion of PRs that were ultimately accepted or rejected given the time it took to decide (the PR’s *age*). We excluded PRs that were left open, since no decision has been reached for those. This is plotted in Figure 4. We notice that, the longer a PR remains open, the higher the probability that it will be rejected. After a threshold x , PRs have a higher probability to be rejected than accepted. The threshold is 28 days for *ansible*, 5 days for *rails* and 25 days for *kubernetes*. Presuming that contributors are aware of this phenomenon, we expect that they implicitly consider PRs left open for a too long duration as being *tacitly* rejected, producing effects similar to those identified in the previous RQ. This preliminary result needs to be confirmed with further analyses.

V. THREATS TO VALIDITY

A threat to the validity of this paper is the fact that we only selected three projects in this exploratory phase, so the preliminary findings might not generalise to bigger sets of projects. Choosing only large, popular and mature projects are also a source of bias, as Rahman and Roy [6] found that such factors affect PR acceptance rates.

Another threat is that the PR status returned by the GitHub API does not necessarily correspond to the actual fate of the PR in some repositories. One such case is *homebrew-core*, where the policy of the repository is to close most PRs without merging them, but to integrate those they deem appropriate through another mechanism, such as the integrators committing the changes themselves. If analyses were to be applied to this repository, the rate of acceptance would be artificially low due to that specific PR handling policy. Another example is that of *angular*, wherein PRs marked with specific tags (“PR action: merge” and “PR target: *” where * represents one or more branch branches to merge the PR into) will have their relevant code automatically integrated into the repository through commits. Those PRs will appear to be rejected on GitHub, even though they aren’t. It would be possible to recover the actual PR status based on those tags, which is not the case for *homebrew-core*.

Yet another threat is tied to the way we have identified contributors. It has been reported that a single individual may use multiple identities in different capacities or at different times on software repositories [14]–[16]. More specifically, it may be the case that the same author owns multiple GitHub accounts, or that multiple authors contribute under the same GitHub account. In that case, we may have erroneously attributed the PRs to an incorrect identity. This could have affected our findings. Therefore, as future work, we aim to

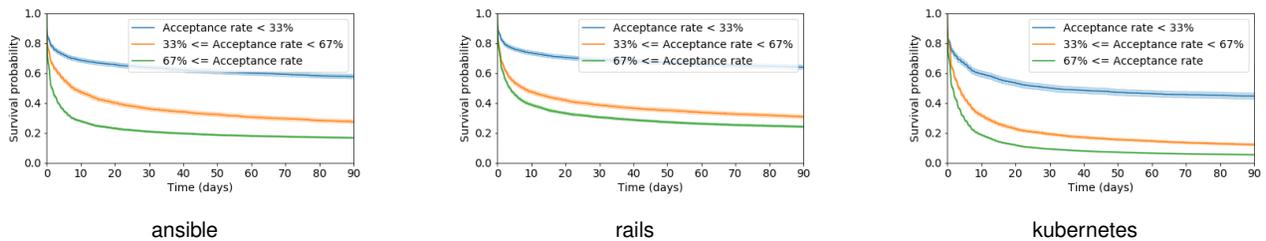


Fig. 3: Survival curves for the probability to submit a next PR, grouped by acceptance rate classes.

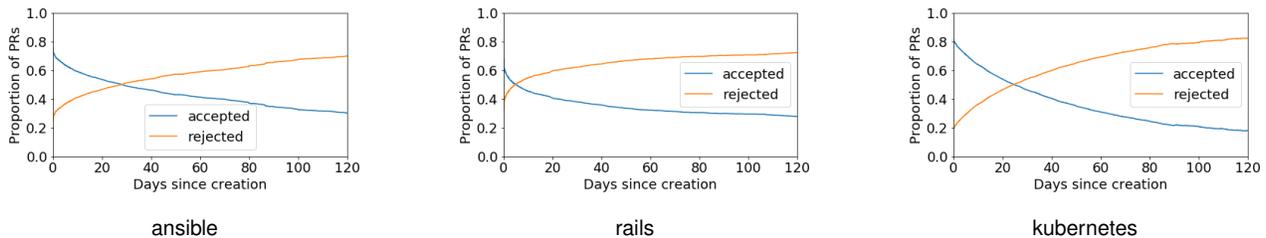


Fig. 4: Proportion of PRs that were ultimately accepted in function of their age

empirically study the impact of such incorrect identifications.

VI. CONCLUSION

The collaborative development of open-source software through a pull-based contribution process involves subtle social interactions that can influence the frequency and likelihood of contribution to a repository, or even its ability to retain contributors. Recent qualitative results have highlighted that contributors do not appreciate the rejection of their PRs, and that they find poor responsiveness from integrators frustrating. Integrators, on the other hand, are wary of alienating contributors in their handling of PRs.

In this paper, we provide preliminary quantitative results showing that a contributor’s PRs are more likely to be accepted when he has submitted more PRs previously. We also reveal the impact of PR decisions on the willingness of contributors to contribute anew. Indeed, fewer contributors submit a new PR after the previous one was rejected than when the previous one was accepted. This highlights the importance for project integrators to avoid alienating contributors, lest they lose their contributions.

ACKNOWLEDGEMENTS

This research was supported by the FRQ-FNRS collaborative research project R.60.04.18.F *SECOHealth*, the Excellence of Science project 30446992 *SECO-ASSIST* financed by FWO-Vlaanderen and F.R.S.-FNRS, and F.R.S.-FNRS Grant T.0017.18.

REFERENCES

- [1] Georgios Gousios, Martin Pinzger, and Arie van Deursen. An exploratory study of the pull-based software development model. In *International Conference on Software Engineering*, pages 345–355. ACM, 2014.
- [2] Georgios Gousios, Andy Zaidman, Margaret-Anne Storey, and Arie van Deursen. Work practices and challenges in pull-based development: The integrator’s perspective. In *International Conference on Software Engineering*, pages 358–368. IEEE Press, 2015.
- [3] Georgios Gousios, Margaret-Anne Storey, and Alberto Bacchelli. Work practices and challenges in pull-based development: The contributor’s perspective. In *International Conference on Software Engineering*, pages 285–296. ACM, 2016.
- [4] Georgios Gousios and Andy Zaidman. A dataset for pull-based development research. In *Working Conference on Mining Software Repositories*, pages 368–371. ACM, 2014.
- [5] Y. Yu, H. Wang, V. Filkov, P. Devanbu, and B. Vasilescu. Wait for it: Determinants of pull request evaluation latency on GitHub. In *Working Conference on Mining Software Repositories*, pages 367–371, May 2015.
- [6] Mohammad Masudur Rahman and Chanchal K. Roy. An insight into the pull requests of GitHub. In *Working Conference on Mining Software Repositories*, pages 364–367. ACM, 2014.
- [7] Jason Tsay, Laura Dabbish, and James Herbsleb. Influence of social and technical factors for evaluating contribution in GitHub. In *International Conference on Software Engineering*, pages 356–366. ACM, 2014.
- [8] Josh Terrell, Andrew Kofink, Justin Middleton, Clarissa Rainear, Emerson Murphy-Hill, Chris Pamin, and Jon Stallings. Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Computer Science*, 3:e1111, May 2017.
- [9] Ayushi Rastogi, Nachiappan Nagappan, Georgios Gousios, and André van der Hoek. Relationship between geographical location and evaluation of developer contributions in Github. In *International Symposium on Empirical Software Engineering and Measurement*. ACM, 2018.
- [10] O. Aalen, O. Borgan, and H. Gjessing. *Survival and Event History Analysis: A Process Point of View*. Springer, 2008.
- [11] E. L. Kaplan and P. Meier. Nonparametric estimation from incomplete observations. *J. American Statistical Association*, 53(282):457–481, 2012.
- [12] Winston Haynes. *Bonferroni Correction*, pages 154–154. Springer, New York, 2013.
- [13] D. R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2):187–220, 1972.
- [14] Mathieu Goeminne and Tom Mens. A comparison of identity merge algorithms for software repositories. *Science of Computer Programming*, 78(8):971–986, August 2013.
- [15] Erik Kouters, Bogdan Vasilescu, Alexander Serebrenik, and Mark G. J. van den Brand. Who’s who in Gnome: using LSA to merge software repository identities. In *International Conference on Software Maintenance*, pages 592–595. IEEE, 2012.
- [16] I. S. Wiese, J. T. d. Silva, I. Steinmacher, C. Treude, and M. A. Gerosa. Who is who in the mailing list? Comparing six disambiguation heuristics to identify multiple addresses of a participant. In *International Conference on Software Maintenance and Evolution*, pages 345–355, Oct 2016.