

Should I Get Involved? On the Privacy Perils of Mining Software Repositories for Research Participants

Melina Vidoni
melina.vidoni@anu.edu.au
Australian National University
Australia

Nicolás E. Díaz Ferreyra
nicolas.diaz-ferreyra@tuhh.de
Hamburg University of Technology
Germany

ABSTRACT

Mining Software Repositories (MSRs) is an evidence-based methodology that cross-links data to uncover actionable information about software systems. Empirical studies in software engineering often leverage MSR techniques as they allow researchers to unveil issues and flaws in software development so as to analyse the different factors contributing to them. Hence, counting on fine-grained information about the repositories and sources being mined (e.g., server names, and contributors' identities) is essential for the reproducibility and transparency of MSR studies. However, this can also introduce threats to participants' privacy as their identities may be linked to flawed/sub-optimal programming practices (e.g., code smells, improper documentation), or vice-versa. Moreover, this can be extensible to close collaborators and community members resulting "guilty by association". This position paper aims to start a discussion about indirect participation in MSRs investigations, the dichotomy of 'privacy vs. utility' regarding sharing non-aggregated data, and its effects on *privacy restrictions* and *ethical considerations for participant involvement*.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections; Privacy-preserving protocols**; • **Software and its engineering** → **Software libraries and repositories**.

KEYWORDS

Mining software repositories, developers' privacy, ethical considerations, privacy-preserving protocols, indirect participants

ACM Reference Format:

Melina Vidoni and Nicolás E. Díaz Ferreyra. 2022. Should I Get Involved? On the Privacy Perils of Mining Software Repositories for Research Participants. In *Proceedings of Co-Located with ICSE 2022 (RoPES'22)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 POSITION MOTIVATION

Let us consider the following motivational example:

Motivational Example. A group of researchers mine several

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

RoPES'22, 1st International Workshop on Recruiting Participants for Empirical Software Engineering, Pittsburgh, USA

© 2022 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

C++ projects from GitHub with the aim of spotting security flaws and determine gaps in developers' cybersecurity training. Following existing recommendations, they craft a detailed inclusion/exclusion criteria, perform a systematic repository search, and extract developers' emails from their associated GitHub's profiles. Using the latter, they distribute a survey to elicit their knowledge and practices regarding software security. The process undergoes an Ethical Protocol examination, which is approved, and requests participants' data (e.g., emails, names) to be anonymised.

Problem Statement. Several issues arise from the motivational example. Due to open science standards, it is increasingly common to provide replication packages, which may disclose the analysed sources by name, or through non-aggregated data (e.g., code snippets, function signatures). How this data is compiled and reported can be unclear for the study participants as this information is often absent in the Participant Information Sheet (PIS). Moreover, sometimes is not even explained to them nor addressed within the Ethical Protocol. By disclosing non-aggregated (sometimes *raw*) data, researchers may allow participants' identities to be traced back, even for those who did not respond the survey.

The latter can be considered 'indirect participants', as in medical research: "those whose rights and welfare may be affected by the intervention through their routine exposure to the environment in which the intervention is being deployed" [11]. Within the context of MSRs, the 'intervention' can be interpreted as the analyses or studies conducted using the data being mined from the software repositories; disclosing the findings and the repositories by name (or with traceable data) allows the re-identification of 'indirect participants'.

This situation can nonetheless lead to 'guilty by association' issues. For instance, let us consider the case of a project A in which many developers collaborated but only a few of them contributed largely (the rest just had a secondary role). Imagine that one of these active members (who is very security unsavvy) takes the survey described in the motivational example. Now imagine that, after analysing the mined repositories and processing the survey results, the researchers conclude that security flaws in open-source projects are closely related to gaps in the cybersecurity training of their contributors. If the researchers disclose project A's name alongside the identified security flaws, one could think that all developers who were involved in the project are unsavvy when, in fact, only a bunch of them are responsible for the security flaws and the rest are just 'guilty by association'.

Results may negatively affect practitioners' careers if expose (as in the example above) less-than-ideal practices (e.g., anti-patterns, security breaches) along with attributes allowing participant's re-identification. Data linkage is a well-known privacy threat in the

Internet of Things as general information about a community or its members can be easily inferred from the data collected by different devices [15]. This poses an issue regarding the *ethical considerations of participant involvement* (e.g., informed consent) and whether or not the research may negatively impact their careers. Information leakage in data mining and its potential privacy breaches is a known risk [9], which has reportedly compromised participants' privacy through indirect data sharing [8].

On the other hand, presenting just aggregated data (to prevent these issues or due to the Ethical Protocol) may limit the study's reproducibility and lead to peer-review criticism. However, providing non-aggregated/raw data may also affect *privacy restrictions* (e.g., anonymity and confidentiality) or even bypass the data-protection boundaries enforced by *ethical committees*. Although some privacy-preserving data-linkage protocols have been proposed within the current literature [1, 7], their expansion to MSRs and mixed-methods remains unclear. Similarly, the analysis of outliers in datasets may lead to the identification of private information regarding those particular cases, as most anonymisation techniques (both semantic and syntactic) do not consider outliers in their process [14].

2 PERILS OF INVOLVEMENT

The problems derived from the motivational example may have unwanted and negative consequences.

PERIL I. The likelihood of data being retraced and identities uncovered may hinder participants' willingness to partake, by (a) not returning surveys (affecting *self-selection bias* and causing *sample sizes issues* [2]), (b) providing false information in the survey, (c) requesting not to be contacted in the future (affecting future works), or (d) requesting their repository to be withdrawn.

PERIL II. Data traceability may discourage developers from contributing to OSS (open-source software) repositories or persuade them to enact 'barriers' to prevent indirect participation [11]; e.g., licences limiting which investigations can be conducted [5]. This is an *ethical consideration* regarding *participant involvement*, because the consent over who can use community resources in OSS is crucial, reverting to several questions posed by prior research [3]:

Is it necessary for each individual in a community to give consent to research before research can be conducted, or can a few individuals represent the group? This raises two important distinctions: (1) how can a community give consent before they are researched, and (2) should a community be required to consent to the publication of the results of research.

PERIL III. Researchers are the 'guardians' of the participants' privacy, and may incur in indirect disclosure of personal data (related to *privacy restrictions* and *ethical considerations*). Current trends in open science endorse sharing a primary dataset alongside articles¹. However, as Erb et al. [4] stated, "while an informed consent given by study participants may satisfy legal requirements, and the removal of obvious identifiers prevents trivial re-identification, we argue that this perspective is typically too short-sighted to address actual ramifications of public data releases in terms of anonymity". This topic is not new, and many disciplines

have analysed it [12, 13]. Nevertheless, it remains foreign to MSRs studies and direct/indirect participants.

3 THE PATH FORWARD

Preserving participants' privacy while maximising the utility of their personal data is an old trade-off conundrum. [6]. Still, it has repercussions on individuals' privacy [4] that go beyond the sinless use of contact details in MSR studies. The nuances of ethical mining and data disclosure [5] can harm MSRs participants (both direct and indirect) while limiting the recruitment and engagement in future research. However, the impact of MSRs and open science on participants' privacy must be closely examined. Hence, this position proposes three areas for future works:

PROPOSAL I. Other disciplines posed privacy-aware methodologies to prevent deanonymisation through data linkage [1, 7], and guidelines for handling data disclosure on research depending on information retrieval and participant recruitment [10]. Investigating their transfer to mixed-method MSR studies is a future line of work, relevant to the participants' *privacy restrictions* and the *ethical considerations* of their involvement. This can provide guidelines for data sharing and presentation related to *Peril III*.

PROPOSAL II. Participant recruitment in mixed-method studies extends beyond contacting people to complete a survey, partake in workshops or interviews. Further assessing the nuances of participation, consent and involvement [3], may lead to the evolution of ethical mining approaches [5] to mitigate the risk of retracing and deanonymising data [8, 15]. This involves assessing indirect participation in MSRs [11], and issues posed in *Perils I-II*.

PROPOSAL III. Balancing the 'privacy vs utility' dichotomy [6], to investigate trends in MSRs studies and their effects on participants (both direct/indirect). Open science in software engineering, particular MSRs, needs to bring the discussion of participants' data-protection rights and information disclosure to the forefront [4]. This would enable *institutional review boards* to have guidelines for the assessment of Ethical Protocols and would assist the peer-review process of ethically-obscured methodologies. This relates to *Perils II-III*.

4 CONCLUSIONS

This position aims to present the problem of OSS contributors as an *ethical consideration of participant involvement* in MSR studies while discussing examples of existing situations that affect *privacy restrictions*. The overall goal is to highlight the problems from the lens of the 'privacy vs utility' spectrum and propose further lines of research that will contribute to the robustness and ethical conduction of MSRs (both pure and mixed-method) investigations.

Several lines of work are available as a path forward. First, to investigate what are the current perceptions of researchers regarding privacy considerations of MSR studies, and compare them to the 'actual' privacy measures they follow (if any). Second, evaluating the 'privacy vs utility' dichotomy in MSR papers, to produce guidelines that would help researchers decide when to prioritise reproducibility over privacy (and vice-versa).

REFERENCES

- [1] Dalal Al-Azizy, David Millard, Nigel Shadbolt, and Kieron O'Hara. 2014. Deanonymisation in Linked Data: A Research Roadmap. In *World Congress on*

¹See: <https://conf.researchr.org/track/icse-2022/icse-2022-open-science-policies>

- Internet Security*. IEEE, UK, 48–52.
- [2] Sebastian Baltes and Paul Ralph. 2020. Sampling in Software Engineering Research: A Critical Review and Guidelines. *ArXiv abs/2002.07764* (2020), 1–21.
 - [3] David Berry. 2004. Internet research: Privacy, Ethics and Alienation: An Open Source Approach. *Internet Research* 14 (09 2004), 323–332.
 - [4] Benjamin Erb, Christoph Bosch, Cornelia Herbert, Frank Kargl, and Christian Montag. 2021. Emerging Privacy Issues in Times of Open Science. <https://doi.org/10.31234/osf.io/u236e> PsyArXiv.
 - [5] Nicolas E. Gold and Jens Krinke. 2020. Ethical Mining: A Case Study on MSR Mining Challenges. In *17th International Conference on Mining Software Repositories* (Republic of Korea). ACM, USA, 265–276.
 - [6] Tiancheng Li and Ninghui Li. 2009. On the Tradeoff between Privacy and Utility in Data Publishing. In *15th International Conference on Knowledge Discovery and Data Mining* (France). ACM, USA, 517–526.
 - [7] Christine M. O’Keefe, Ming Yung, Lifang Gu, and Rohan Baxter. 2004. Privacy-Preserving Data Linkage Protocols. In *ACM Workshop on Privacy in the Electronic Society*. ACM, USA, 94–102.
 - [8] Layla Pournajaf, Daniel A. Garcia-Ulloa, Li Xiong, and Vaidy Sunderam. 2016. Participant Privacy in Mobile Crowd Sensing Task Management: A Survey of Methods and Challenges. *SIGMOD Rec.* 44, 4 (2016), 23–34.
 - [9] M. Prakash and G. Singaravel. 2015. An Approach for Prevention of Privacy Breach and Information Leakage in Sensitive Data Mining. *Comps. & Electrical Eng.* 45 (2015), 134–140.
 - [10] Elissa M Redmiles, Yasemin Acar, Sascha Fahl, and Michelle L. Mazurek. 2017. *A Summary of Survey Methodology Best Practices for Security and Privacy Researchers*. Digital Repository at the University of Maryland.
 - [11] Jaye Bea Smalley, Maria W. Merritt, Sana M. Al-Khatib, Debbie McCall, Karen L. Staman, and Carl Stepnowsky. 2015. Ethical Responsibilities Toward Indirect and Collateral Participants in Pragmatic Clinical Trials. *Clinical Trials* 12, 5 (Oct 2015), 476–484.
 - [12] Adrian Thorogood. 2018. Canada: Will Privacy Rules Continue to Favour Open Science? *Human Genetics* 137, 8 (01 Aug 2018), 595–602.
 - [13] Colin G. Walsh, Weiyi Xia, Muqun Li, Joshua C. Denny, Paul A. Harris, and Bradley A. Malin. 2018. Enabling Open-Science Initiatives in Clinical Psychology and Psychiatry Without Sacrificing Patients’ Privacy. *Advances in Methods and Practices in Psychological Science* 1, 1 (2018), 104–114.
 - [14] Hui Wang and Ruilin Liu. 2015. Hiding Outliers Into Crowd: Privacy-Preserving Data Publishing with Outliers. *Data & Knowledge Engineering* 100 (2015), 94–115.
 - [15] Xu Zheng, Zhipeng Cai, and Yingshu Li. 2018. Data Linkage in Smart Internet of Things Systems: A Consideration from a Privacy Perspective. *IEEE Communications Magazine* 56, 9 (2018), 55–61.