

Does the Tool Matter? Exploring Some Causes of Threats to Validity in Mining Software Repositories

Nicole Hoess

*Technical University of
Applied Sciences Regensburg*
Regensburg, Germany
nicole.hoess@othr.de

Carlos Paradis

No Affiliation
Sunnyvale, CA, USA
cvas@acm.org

Rick Kazman

University of Hawaii at Mānoa
Honolulu, HI, USA
kazman@hawaii.edu

Wolfgang Mauerer

*Technical University of
Applied Sciences Regensburg*
Siemens AG, Technology
Regensburg/Munich, Germany
wolfgang.mauerer@othr.de

Abstract—Software repositories are an essential source of information for software engineering research on topics such as project evolution and developer collaboration. Appropriate mining tools and analysis pipelines are therefore an indispensable precondition for many research activities. Ideally, valid results should not depend on technical details of data collection and processing. It is, however, widely acknowledged that mining pipelines are complex, with a multitude of implementation decisions made by tool authors based on their interests and assumptions. This raises the questions if (and to what extent) tools agree on their results and are interchangeable. In this study, we use two tools to extract and analyse ten large software projects, quantitatively and qualitatively comparing results and derived data to better understand this concern. We analyse discrepancies from a technical point of view, and adjust code and parametrisation to minimise replication differences. Our results indicate that despite similar trends, even simple metrics such as the numbers of commits and developers may differ by up to 500%. We find that such substantial differences are often caused by minor technical details. We show how tool-level and data post-processing changes can overcome these issues, but find they may require considerable efforts. We summarise identified causes in our lessons learned to help researchers and practitioners avoid common pitfalls, and reflect on implementation decisions and their influence in ensuring obtained data meets explicit and implicit expectations. Our findings lead us to hypothesise that similar uncertainties exist in other analysis tools, which may limit the validity of conclusions drawn in tool-centric research.

Index Terms—Mining Software Repositories; Developer Networks; Empirical Software Engineering; Research Software

I. INTRODUCTION

In software projects, social entities, such as persons or groups, interact with technical artefacts, including modules, files and functions. The relationships between these two groups can be described by socio-technical networks constructed from software repositories, which capture a project’s entire history of collaboration [1]. Temporal socio-technical networks can represent a project’s state at any point in time. This information can help managers, new contributors and others to gain a comprehensive understanding of the project, facilitating, for example, the identification of the most suitable contact persons for addressing specific issues, such as a bug in a software module. The analysis of collaboration networks has become a common tool in understanding evolutionary aspects of organisational principles, primarily based on open-source

software (OSS) projects [2], [3], [4], [5], [1], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23]. Patterns of interest include [2] the formation of developer communities [1], the emergence of (non-)hierarchical team structures [9], and the identification of core developers [7], [8]. Socio-technical factors can also aid in predicting quality aspects, including bugs, software failures and future project success [22] as well as in optimising circumstances for collaboration [24] to increase developer productivity [25] and reduce development time and costs.

To facilitate these analyses, the field of mining software repositories (MSR) offers a plethora of methods and tools to extract collaboration and communication data from a variety of sources, including version control systems (VCS), mailing lists, issue trackers and chats [26]. A series of processing stages, such as matching multiple developer identities, bot detection, file and temporal filtering, may be employed to obtain a baseline data set. To construct the collaboration network, contributors are initially linked to artefacts, mailing list threads, issues or chat messages to which they contributed. A bipartite projection eliminates the source nodes from the bipartite network, instead connecting developers who contributed to them. Alternatively, a temporal projection links developers according to the temporal order of their contributions to a source. Edge weights may reflect the strength of collaboration between two developers. In co-change networks, weights can be calculated from the number of lines of code (LoC) changed in a file, function or module [22], [7], [1].

Although existing tools and scripts perform similar steps of this complex pipeline, they differ in their underlying assumptions, algorithms, implementations and interpretations [27]. This raises the question of whether the results of previous studies would have been the same if a different tool or analysis pipeline had been employed. Studies on the promises and perils of mining version control data have identified a number of factors that can potentially lead to inaccurate analysis results, but also identified and proposed mechanisms and methods to enhance consistency. To date, there are only a few studies that quantify the effects of these implementation details on analysis results and their interpretation. In particular, there is a lack of replication studies for developer networks [2].

Our objective is to quantify the impact and threats to validity

of employing different mining tools on version control systems and utilising this baseline data in subsequent developer network analyses. To this end, we conduct a replication study to measure the similarity of the VCS data and collaboration networks mined and constructed by two independently developed tools, namely Codeface [28] and Kaiaulu [29].

Our results indicate significant uncertainties due to unknown assumptions. While overall trends largely correspond, we show that under specific circumstances, such as limitations in file filtering and identity matching, minor implementation and configuration details may have substantial impact. For instance, the percentage of developer identities jointly found by both tools may not exceed 60% for some subject projects and even most active developers can be missed. Changing tool parametrisation and implementing post-processing could overcome parts of these issues in our case. However, not all discrepancies can be eliminated by such means, and would require major re-engineering efforts. Users are likely unaware of the existence and magnitude of such limitations and could derive wrong conclusions, decisions and best practices or make predictions based on inaccurate ground truth data when trusting their pipeline. This paper deliberately focuses on the technical details responsible for this issue.

In summary, we contribute (1) quantitative insights into the impact of different mining tool implementations on the characteristics of the baseline data, (2) a qualitative analysis of the causes of these differences and (3) actionable learnings from and for the adjustment and replication of a tool to obtain outcomes as close as possible to the desired baseline data.

II. RELATED WORK

A. Replication Studies

The limited number of available replication studies is considered a threat to validity in empirical software engineering [30], [31]. In general, advances in reproducibility engineering, such as dedicated platforms [32], [33] and archived, self-contained reproduction packages bundling artefacts and pipelines [34], contributed significantly to the improvement of study reproducibility in the last years [35]. MSR studies yet remain challenging, as artefacts and implementations are often not published at all, only in parts, or in form of a diverse [36], [33], possibly unusable [35], set of scripts and tools. However, replication studies offer the potential to increase confidence in findings of previous studies [37], extend or complement their results [38] and increase the impact of the field [36]. For example, Dinh-Trong *et al.* [39] compare the development process of the Apache project investigated in the original work to the one of FreeBSD to enhance understanding of OSS development processes. Bock *et al.* [3] refine the results of a study on synchronous development, which associates developer collaboration and communication with productivity [25], by transferring the method to a different data set and abstraction level. While replication studies are more common in MSR areas such as defect and bug detection [40], [41], their availability is limited in the context of developer networks [2].

B. Validity Studies

Pitfalls in MSR that can distort results have been addressed by several influential studies. Bird *et al.* [42] addressed possible traceability issues. Kalliamvakou *et al.* [43] extend this study to include peculiar properties of GitHub. Flin *et al.* [44] studied pitfalls when working with time-based Git data and provide guidance to overcome them. Saarimäki *et al.* [45] investigate the perils of disregarding time dependencies in MSR studies. They propose methods from time series analysis and present preliminary results of their impact. Nia *et al.* [46] explore the impact of pitfalls in e-mail network construction, resulting in missing edges or edges out of temporal order. Their results show that metrics such as node centrality are stable despite the changes in topology. Meneely *et al.* [47] study developer perception of collaboration and team structure expressed by developer network edges and metrics. From a higher level perspective, Siegmund *et al.* [48] studied experts' views on the *correct* application of empirical methods in software engineering, finding that there is no consensus on the prioritisation of internal or external validity in the field.

C. Method and Tool Comparisons

In technical debt detection, Lefever *et al.* compared the results found by various commercial and open-source tools, finding that they disagree even for very common, basic measures such as LoC [49]. To the best of our knowledge, there are no similar tool comparisons for developer networks. However, differentiating aspects are highlighted by comparisons of specific methods implemented as part of MSR pipelines. Goeminne *et al.* [50] compare heuristics for developer identity matching, finding that exact substring matching performs best. Amreen *et al.* [51] propose a novel machine learning technique for identity matching and demonstrate its superiority to state-of-the-art methods. Bertoncello *et al.* [52] investigate the influence of using commits or pull-requests to distinguish core and casual contributors in a software project, concluding that pull-requests are more accurate for measuring contributions. Joblin *et al.* [1], [7] investigate differences in co-change collaboration network construction, community detection and core developer classification metrics, along with an evaluation of correspondence with the real perception of developers. Several authors study the agreement of developer networks [26], [53] and communities [6] resulting from data extracted from different communication and collaboration sources. Tymchuk *et al.* [26] find that the combination of channels is essential to obtain a comprehensive view of a project. While related studies focus on specific analysis stages in isolation, we aim to explore the impact and propagation of, possibly unknown, differences in two distinct tool implementations along the entire analysis pipeline.

III. REPLICATION STUDY DESIGN

Our replication is based on an iterative process visualised in Fig. 1. We began with tool configurations of Codeface and Kaiaulu used in previous studies, and ran the tools on ten large software repositories to analyse the results and evaluate

their agreement. Based on the identified differences and once knowing their causes, we adapted the pipeline for Kaiaulu and began the next iteration to achieve a closer replication to Codeface. For the sake of clarity, the individual iterations are not distinguished in the following sections.




Fig. 1. Overview about our approach: We mined repositories using a basis (Codeface) and replication (Kaiaulu) tool, and compared derived data. Depending on similarities and differences, we adjusted configuration parameters and post-processing methods in the replication tool. We iterated until the closest replication was achieved.

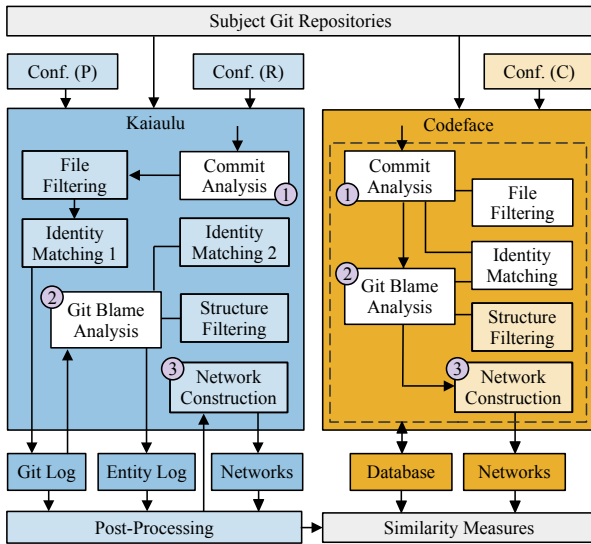


Fig. 2. Informal overview of structural components of the mining tools Codeface and Kaiaulu. Although both tools perform the same analysis steps, their interaction and data structure differ.

A. Research Questions

RQ1: *To what extent can we observe discrepancies in the data obtained from independent MSR tools?* Each tool makes assumptions on how to extract and process data, which can be influenced by configuration parameters. Implicit assumptions and implementation decisions often remain hidden from tool users. We are interested in quantifying the impact of these uncertainties on the results.

RQ2: *Which factors are responsible for the discrepancies?* Assuming we find discrepancies in RQ1, we are interested in technical causes and inherently different tool assumptions.

RQ3: *To what extent and with what effort can a tool be adapted to yield results comparable to those produced by another tool?* Based on RQ2, we evaluate how adjusting configuration and implementation achieves the closest replication of baseline data extracted by the original tool.

Fig. 2 provides an overview of the analysis pipeline.

B. Subject Selection

Our data set consists of ten relevant and active OSS projects. For our selection, we take into account previous studies relying on Codeface or Kaiaulu and drawing conclusions based on these repositories [1], [3], [4], [5], [6], [8], [9], [12], [20], [22], [54]. We consider different application domains, programming languages, ages of the project and project and team sizes. Tab. I provides an overview of the subject characteristics. To avoid a tool bias, commits and developers are extracted directly via git, without processing such as identity matching. The primary language and LoC are determined via CLOC [55].

C. Data Extraction

We chose Codeface as the basis tool because it has been used in numerous studies between 2015 and 2023 [1], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23]. We decided to use Kaiaulu as replication tool because it is relatively new with fewer studies available [27], [54], [56], [57], [58]. Kaiaulu is partly inspired by Codeface [27] and, to the best of our knowledge, the only tool implementing *all* required stages of the MSR pipeline to analyse developer collaboration in the same way as Codeface. We set up both by cloning the public, actively maintained repositories for Codeface [28] and Kaiaulu [29] and following the installation instructions.

To conduct the analyses, both tools require configuration files for the tool itself and for each project analysis run. Codeface enables its users to configure a time window size in months or specific release ranges for temporal analyses. We chose a window size of three months as proposed in Refs. [7] and [6]. In addition, Codeface offers three modes for commit analysis and network construction, which measure contributions and collaboration strengths based on files, features or code structures. For our study, we choose the structure-based *proximity* mode, which is also available in a similar form with Kaiaulu’s *entity* mode, parsing functions and classes.

Kaiaulu’s API offers a high degree of freedom for the analysis pipeline. We built on the API to define scripts for the analysis. During our study, we identified a number of decisions that influence results. Therefore, we use two distinct configurations, *prior* and *replication*. *Prior* refers to the configuration we set up based on prior studies conducted with Kaiaulu *before* the first exploration of the baseline data for RQ1. *Replication* is adjusted and complemented with additional post-processing derived from RQ2 *after* iteratively inspecting the mined data to most closely replicate Codeface’s outputs and answer RQ3.

To ensure a fair comparison, both MSR tools use identical time windows. Codeface’s time slicing starts and ends by commit timestamps instead of actual time windows, implying that a commit break of six months would not result in two empty time windows, but a larger time window ending with the commit after the break. To replicate this behaviour in Kaiaulu, we implement a new time slicing, enabling us to specify the exact window limits used by Codeface. We also add new configuration options to split time intervals by either author timestamp (*prior*) or committer timestamp (*replication*).

TABLE I
DESCRIPTIVE STATISTICS OF OUR SUBJECT PROJECTS

Project	Domain	Primary Language	Commits	Developers	LoC	Start	End
CAMEL	Middleware	Java	68,222	1352	2,252,114	2007-03-19	2024-01-11
DJANGO	Web Framework	Python	32,348	2897	637,711	2005-07-13	2024-01-23
GTK	GUI	C	78,899	1587	1,383,746	1997-11-24	2024-01-16
JAILHOUSE	Hypervisor	C	3006	75	50,593	2013-10-20	2023-01-10
OPENSSL	Secure Communication Library	C	34,691	1069	861,736	1998-12-21	2024-01-21
POSTGRESQL	DBMS	C	57,500	55	1,708,628	1996-07-09	2024-01-11
QEMU	Emulator, Virtualiser	C	109,894	2376	1,972,799	2003-02-18	2024-01-10
RSTUDIO	IDE	Java, C++	41,745	194	1,004,561	2010-12-07	2024-01-25
SPARK	Data Analytics Engine	Scala	39,503	2815	1,377,476	2010-03-29	2024-01-11
WINE	Compatibility Layer	C	168,818	1762	5,251,550	1993-06-29	2024-01-10

In Fig. 2, the lighter coloured areas represent configurable parts of the tool pipelines. In Codeface, commit analysis, subsequent code structure extraction and network construction depend on the choice of analysis mode and are executed in parallel, triggered by a single interface. Depending on the configuration, the results stored in Codeface’s database tables will differ. In Kaiaulu, each step during and subsequent to commit analysis is optional and uses independent configuration parameters to be specified by the user, such as file endings to filter and code structure tags to search for. Codeface decides internally on which parser, file filters and tags to use. In Kaiaulu’s prior configuration, we filtered the same file endings and language tags as proposed in existing configurations from prior Kaiaulu studies [27], [54], while in the replication configuration, we filtered the same file endings and language structure tags as proposed for Codeface.

Kaiaulu’s identity matching is optionally performed at multiple occasions, always requiring a specification of the table and identity columns. Codeface performs identity matching as proposed by Bird *et al.* [59] when parsing commits across all identity columns (author and committer) and stores the unified results. We implemented similar cross-column and cross-table identity matching in Kaiaulu’s replication configuration.

D. Developer Network Construction

In function-based developer networks, the collaboration strength of each pair of developers is calculated from the number of changed source lines of code $|sloc|$ reported in the git blame data of a revision. For each commit from developer b , Codeface iterates over any modified functions or code structures. For each function, it considers all commits from the collaborating developer a , who previously modified this function. For each of those previous commits, Codeface adds the sum of newly contributed $|sloc_{b,c_i}|$ from b and the remaining $|sloc_{a,c_j}|$ of the previous commit from a to the edge weight. Kaiaulu provides several edge weight schemes for the function-based analysis mode. The scheme implemented and used for prior Kaiaulu studies also sums up $|sloc|$, but, instead of adding the newly contributed $|sloc_{b,c_i}|$ to each of the previous commits of a , only adds them once. Therefore, we introduced a new edge weight scheme to replicate the nested sum from Codeface in Kaiaulu’s replication configuration.

For two developers d_1 and d_2 , the edge weight ω_{d_2,d_1} is given by the sum of their contributions measured in source lines of code $sloc$ at time of commit $\tau_{d,t}$ as described in Eq. (1) for Codeface (c) and Eq. (2) for Kaiaulu (k) with prior configuration. δ filters commits of d_1 contributing to f .

$$\omega_{d_2,d_1,c} = \sum_{f \in F} \left\{ \sum_{i=0}^n \left[\sum_{j=0}^{i-1} \left(|sloc_{d_2}(\tau_{d_2,i-1}, f)| + |sloc_{d_1}(\tau_{d_1,j}, f)| \right) \cdot \delta(|sloc_{d_1}(\tau_{d_1,j}, f)|) \right] \right\} \quad (1)$$

$$\omega_{d_2,d_1,k} = \sum_{f \in F} \left\{ \sum_{i=0}^n \left[|sloc_{d_2}(\tau_{d_2,i-1}, f)| + \sum_{j=0}^{i-1} |sloc_{d_1}(\tau_{d_1,j}, f)| \cdot \delta(|sloc_{d_1}(\tau_{d_1,j}, f)|) \right] \right\} \quad (2)$$

Our comparison uses count-based metrics for number of identified commits, changed files, changed entities and developer identities over time. Time series similarity is calculated by normalised compression distance [60], measuring the distance of the compression of two vectors individually and concatenated. Also, dynamic time warping identifies the optimal warping path aligning elements of the sequences with minimal distance [61]. We include Spearman’s rank correlation coefficient as intuitive measure, knowing that it can be misleading for time series due to dependent values [62]. To measure the similarity of the developer networks constructed by both tools, we calculated their density, mean non-zero edge weight and graph edit distance as a notion of required edit operations to transform one graph into the other.

IV. RESULTS

To present results and address research questions, space constraints require limiting visualisations to representative samples; full data is available in the reproduction package.

A. Comparison of Baseline and Derived Data

To address RQ1, we visualise the number of commits, active developers, entities and files found by each tool configuration in the defined time intervals in Fig. 3. In Tab. II, we quantitatively evaluate the visual trends, comparing the time series’ normalised compression distance (NCD), dynamic

time warping (DTW) distance and Spearman’s rank correlation coefficient (Cor.). We show the tools’ agreement on jointly identified files, entities and developer identities in Fig. 4.

Commits: Fig. 3 shows that the results from both tools describe similar trends in the time series of commits. The DTW distance and Spearman’s rank correlation coefficient presented in Tab. II confirm this quantitatively, although NCD appears more sensitive. Codeface usually reports significantly higher numbers of commits than Kaiaulu with prior configuration. The extent of discrepancies depends on the project and is particularly evident for Camel and Spark. In one of Spark’s analysis time intervals, Codeface extracted 1593 commits, while Kaiaulu only identified 331. The difference in commits is crucial, because each subsequent step in the analysis pipelines, shown in Fig. 2, relies on them.

Files: Aligned with the commit time series, the file time series in Fig. 3 indicates similar trends for Codeface and Kaiaulu. However, despite Codeface usually identifying more commits, Kaiaulu appears to find more files in certain project contexts such as GTK and Wine. The percentage of jointly identified files compared by their names is shown in Fig. 4 over all time intervals. Again, we observe a strong project-dependent intersection, with values below 20% for Spark and up to 90% for RStudio when considering the prior configuration.

Entities: The time series of changed entities shown in Fig. 3 do not allow any clear conclusions to be drawn. While we see very similar curves for the GTK and Wine projects, the similarity is not as pronounced for Spark and almost non-existent for RStudio. We can also recognise this trend from the metrics presented in Tab. II. Although correlation is high, the DTW distances of 0.34 for Jailhouse and 0.33 for RStudio are the worst in the table. The jointly identified entity names show a similar picture as the time series: in Fig. 4, we can see a high intersection of 60-80% in the GTK and Wine projects, while the intersection in Camel, RStudio and Spark is low.

Developers: Developer time series in Fig. 3 differ to a similar extent as the commit time series. Surprisingly, the direction of discrepancies is reverse for some projects (*e.g.*, Camel, RStudio and Wine). This indicates that Kaiaulu with prior configuration detects more developers than Codeface, although it identifies less commits. The time series similarity measures presented in Tab. II do not capture this phenomenon. The agreement on developer identities shown in Fig. 4 gives a better intersection than for files and entities. However, agreement on developer identities can again drop below 60%.

Discrepancies in the four dimensions are not always related, but may have different causes. For instance, for Camel, we observe a high intersection of 60-80% in developers, while the agreement on identified entities is below 20%.

Developer Networks: To determine similarities between constructed networks, we compute graph edit distance, network density and mean (non-zero) edge weight for each time interval (see Tab. III). Intuitively, graph edit distance is large for projects with several thousand developers (*e.g.*, QEMU and Django), and small for projects with several dozen to hundred developers (*e.g.*, Jailhouse, Postgres and RStudio). As for time

series, Spark represents an outlier with a graph edit distance twice as high as the second-largest observation. This indicates that discrepancies from the baseline data do propagate to derived data. Tab. III also shows that networks derived by Kaiaulu are in general denser than by Codeface, while edge weights from the latter considerably exceed the former in its prior configuration. For QEMU, the maximum edge weight found by Kaiaulu is 11,960, while Codeface reports a value of 261,362—almost 22 times larger. Manual inspection reveals that several developers were missed by Kaiaulu; especially two of the most active developers according to Codeface in one time interval of Spark (see Fig. 5). Anomalies such as duplicate developer identities persist throughout the networks.

Answering RQ1: We observe similar evolutionary trends for count statistics of commits, developers, modified files and entities. The extent of (dis-)agreement however depends on measure, project context and tool configuration. In case of the Spark project, for instance, we observe an absolute percentage difference of up to 500% in the number of commits. Substantial discrepancies in nodes and edge weights of derived developer networks confirm that baseline data differences propagate to high-level analyses, and merit close scrutiny.

B. Investigation of Discrepancies

To address RQ2, we analyse differences between baseline and derived data, and summarise findings. First, we filter instances detected by one but missed by the other tool. We then manually review a random subset of instances using VCS history, Codeface’s database, Kaiaulu’s git and entity log tables, and adjacency matrices of the developer networks.

Commits: The majority of commits missed by Kaiaulu concerned files that use programming languages not specified in the prior configuration. Unlike Codeface, commit extraction in Kaiaulu is not generic, but requires to explicitly specify file suffixes to be analysed. In prior studies, Kaiaulu’s configurations considered suffixes `.c`, `.cc`, `.ccp`, `.java`, `.js`, `.py`, and `.r`. Commits to language files such as `.scala`, primary language in Spark, and commits to non-code files such as `.html`, `.md` and `.xml` were missed by Kaiaulu, but captured by Codeface. Other commits missed by Kaiaulu indicate in their commit message that they were cherry-picked, refactorings or tests. This discrepancy can be partially attributed to the exclusion of directories that supposedly contain tests and code examples in the prior configuration. Another reason could be that Codeface’s parsing of the git blame output from step (2) in Fig. 2 iteratively updates commits in its database if a line has been changed by a commit not yet captured. Such updates of the commit table are not performed by Kaiaulu. After parsing the git log with Perceval in step (1), Fig. 2, captured commits may only be reduced, but no more data is added to the table.

The difference data set also includes commits missed by Codeface and captured by Kaiaulu. These commits were usually merge commits of pull requests or feature branches. However, not all merge commits were missed by Codeface. We also note that, occasionally, commits appear in different time intervals in Codeface and Kaiaulu. Kaiaulu’s prior configura-

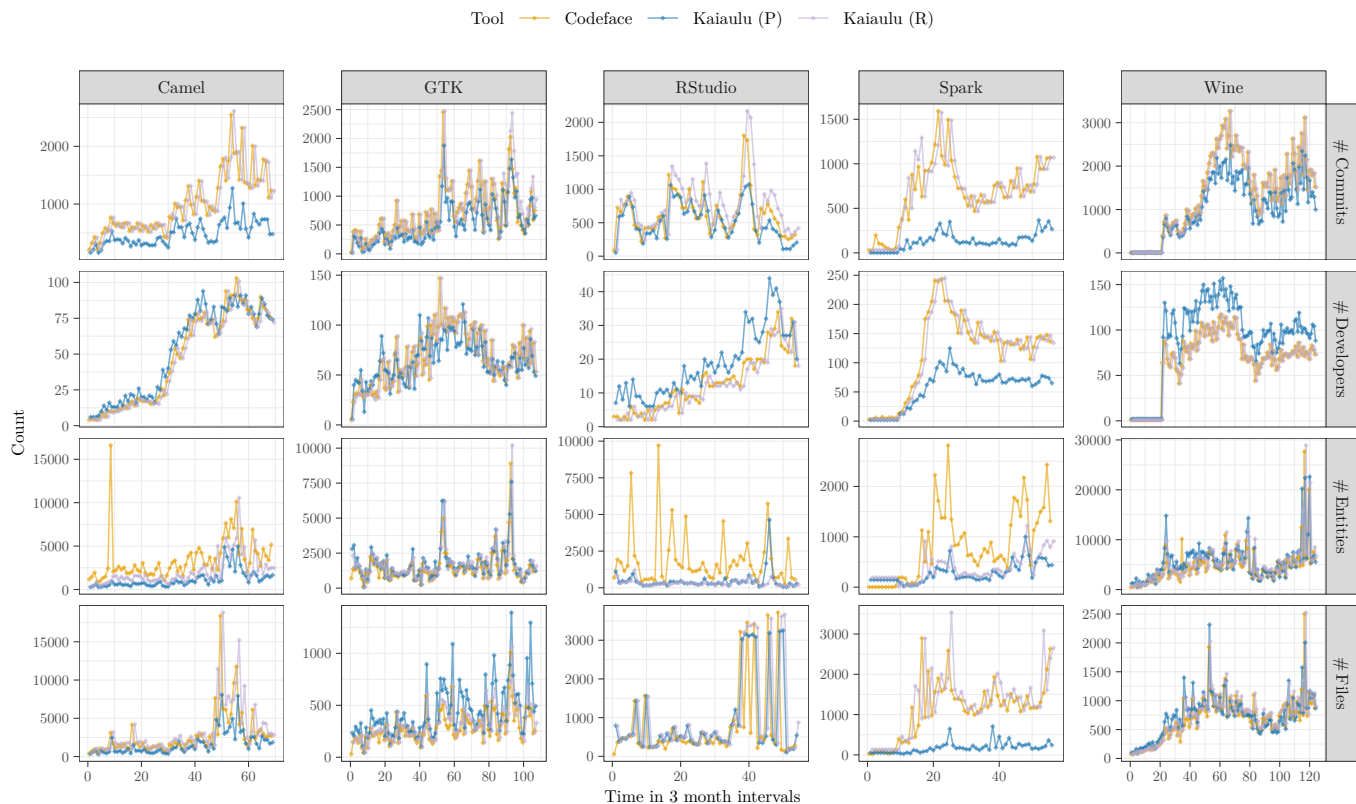


Fig. 3. Time series of count-based metrics calculated based on the git log extracted by Codeface and Kaiaulu with prior (P) configuration and Kaiaulu with replication (R) configuration. Lines are plotted with a small horizontal offset to visualise overlapping lines. We see similar trends for both tools, but notice that the extent of discrepancies is project- and configuration-specific. The most prominent differences can be observed for project Spark.

tion excludes the end time stamp from a time interval, while Codeface includes it. Further discrepancies arise as Codeface splits time intervals based on the committer timestamps, while Kaiaulu uses author timestamps, meaning that time series peaks could be shifted at the transition between two intervals.

Files: Inspecting file difference data reveals discrepancies caused by schema differences. Kaiaulu stores commits multiple times in a table, once for each modified file. This gives a direct relation between commit and file. Codeface, however, stores each commit only once in its commit database table. In proximity analysis mode, file names are recorded in the additional commit dependency table filled *after* file and code structure filtering during git blame analysis. Consequently, a file is only recorded if a code structure (*e.g.*, function) in the file was edited. Thus, Codeface occasionally misses files found by Kaiaulu. All files could be determined by a second Codeface analysis using *file* mode. However, this workaround would lead to issues such as multiple internal identifiers for identical objects, requiring a substantial merging effort. Alternatively, we could adjust Codeface’s database schema to store files and code structures separately. This would require restructuring the entire analysis process depicted in Fig. 2.

Entities: Discrepancies in identified entities can result from parsing. Codeface applies, based on file extensions, Doxygen, C-Tags or an internal SQL parser to extract structural ele-

ments. Kaiaulu uses C-Tags uniformly. However, similar to file filtering, Kaiaulu requires to configure tags. In prior studies, Kaiaulu only parsed function and class tags for C, C++, Java, Python and R. Codeface, when opting for the C-Tags parser internally, parses a default set of tags to detect structures such as functions, enumerations, namespaces, typedefs, macros, and more for *any* language. More tags for special entities (*e.g.*, constructors) are defined for individual languages such as Go.

Developers: Developers missed by Kaiaulu emphasise the relevance of file filtering, because many of these developers contributed to files with neglected extensions in the prior configuration. Comparing developer names shows that higher developer counts by Kaiaulu are caused by different means of identity matching. In Codeface, developers using multiple identities are merged *across* columns and tables considering both author and committer columns. In Kaiaulu’s prior studies, identities were matched only within the author column. This results in duplicate identities consisting of different sets of names and e-mail addresses of the same person.

Developer Networks: Identifying causes of developer network discrepancies is complex. For instance, filtering few selected programming languages in Kaiaulu’s configuration risks that developers focusing on other languages may be under-represented. Reconstructing the edge weight formulas 1 and 2 defined (implicitly) by the tool authors was especially time-

TABLE II

TIME SERIES SIMILARITY FOR CODEFACE AND KAIUALU WITH PRIOR (P) AND REPLICATION (R) CONFIGURATION MEASURED BY NORMALISED COMPRESSION DISTANCE (NCD), DYNAMIC TIME WARPING (DTW) AND SPEARMAN’S RANK CORRELATION COEFFICIENT (COR). FOR NCD AND DTW, VALUES CLOSE TO ZERO (GRAY BAR) ARE DESIRED; FOR CORRELATION, VALUES CLOSE TO ONE (ORANGE BAR) INDICATE HIGHER SIMILARITY.

Tools	Project	Commits			Files			Entities			Developers		
		NCD	DTW	Cor	NCD	DTW	Cor	NCD	DTW	Cor	NCD	DTW	Cor
Codeface/Kaiaulu (P)	Camel	0.86	0.21	0.92	0.87	0.19	0.92	0.89	0.24	0.91	0.69	0.08	0.91
	Django	0.83	0.17	0.96	0.86	0.23	0.91	0.87	0.24	0.69	0.73	0.08	0.69
	GTK	0.88	0.18	0.93	0.87	0.22	0.91	0.88	0.18	0.89	0.68	0.26	0.89
	Jailhouse	0.67	0.23	0.95	0.68	0.27	0.85	0.68	0.34	0.88	0.56	0.27	0.88
	OpenSSL	0.86	0.12	0.98	0.86	0.11	0.96	0.87	0.18	0.88	0.78	0.07	0.88
	Postgres	0.89	0.31	0.75	0.89	0.22	0.82	0.89	0.22	0.91	0.48	0.08	0.91
	QEMU	0.90	0.11	0.97	0.87	0.14	0.98	0.88	0.22	0.88	0.75	0.06	0.88
	RStudio	0.82	0.23	0.94	0.81	0.20	0.66	0.80	0.33	0.71	0.66	0.17	0.71
	Spark	0.81	0.21	0.88	0.80	0.32	0.78	0.80	0.26	0.83	0.78	0.11	0.83
	Wine	0.89	0.09	0.99	0.90	0.18	0.90	0.92	0.20	0.84	0.68	0.07	0.84
	Mean μ	0.84	0.18	0.93	0.84	0.21	0.87	0.85	0.24	0.84	0.68	0.12	0.84
	Std. Dev. σ	0.06	0.06	0.07	0.06	0.06	0.09	0.07	0.05	0.07	0.09	0.08	0.07
Codeface/Kaiaulu (R)	Camel	0.82	0.02	1.00	0.85	0.12	0.98	0.86	0.20	0.92	0.49	0.02	0.92
	Django	0.53	0.05	1.00	0.83	0.16	0.98	0.86	0.14	0.91	0.43	0.01	0.91
	GTK	0.63	0.10	0.98	0.86	0.12	0.98	0.89	0.06	0.94	0.58	0.03	0.94
	Jailhouse	0.11	0.00	1.00	0.61	0.14	0.98	0.68	0.21	0.92	0.18	0.03	0.92
	OpenSSL	0.12	0.00	1.00	0.85	0.07	0.98	0.87	0.04	0.98	0.29	0.01	0.98
	Postgres	0.07	0.00	1.00	0.89	0.12	0.93	0.88	0.17	0.94	0.10	0.00	0.94
	QEMU	0.75	0.03	1.00	0.89	0.10	0.98	0.87	0.13	0.97	0.25	0.00	0.97
	RStudio	0.78	0.12	0.97	0.77	0.21	0.67	0.80	0.34	0.64	0.58	0.08	0.64
	Spark	0.56	0.10	0.96	0.80	0.16	0.93	0.76	0.25	0.88	0.65	0.01	0.88
	Wine	0.07	0.00	1.00	0.90	0.07	0.98	0.92	0.06	0.97	0.39	0.01	0.97
	Mean μ	0.44	0.04	0.99	0.83	0.13	0.94	0.84	0.16	0.91	0.39	0.02	0.91
	Std. Dev. σ	0.30	0.05	0.01	0.08	0.04	0.09	0.07	0.09	0.17	0.02	0.02	0.09

TABLE III

GRAPH SIMILARITIES DETERMINED BY CODEFACE (C) AND KAIUALU (K) WITH PRIOR (P) AND REPLICATION (R) CONFIGURATION.

Project	Graph Edit Distance				Density			Edge Weight					
	C/K(P)		C/K(R)		C	K(P)	K(R)	C		K(P)		K(R)	
	Mean	Max	Mean	Max	Mean	Mean	Mean	Mean	Max	Mean	Max	Mean	Max
Camel	119	399	140	474	0.17	0.17	0.21	1009	474,484	298	167,093	1633	3,036,609
Django	162	533	236	744	0.13	0.20	0.24	510	188,873	47	9431	230	356,944
GTK	141	534	164	682	0.02	0.07	0.11	642	185,640	120	31,874	706	339,516
Jailhouse	5	15	4	11	0.25	0.34	0.30	539	9608	26	664	124	4943
OpenSSL	63	351	86	442	0.29	0.33	0.44	759	170,390	130	38,087	1471	1,192,273
Postgres	47	231	48	229	0.37	0.36	0.43	960	63,799	247	19,578	1677	1,316,173
QEMU	374	849	544	1436	0.20	0.18	0.20	4469	261,362	73	11,960	951	10,429,560
RStudio	33	175	46	246	0.39	0.59	0.97	7019	157,248	969	31,959	92,754	4,765,887
Spark	848	1951	873	2068	0.13	0.07	0.10	336	63,907	92	27,929	347	159,043
Wine	251	877	647	2083	0.03	0.02	0.12	4200	512,257	61	108,895	158	352,091



consuming. As shown in Fig. 2, an analysis run in Codeface passes all steps from commit analysis to network construction and large projects require run-times of multiple days. Using Kaiaulu’s mock data generator (650 LoC), which creates small test repositories, and running Codeface on these while adding targeted debugging statements to the code, revealed that Kaiaulu’s prior edge weight scheme did not account for repeated past contributions, but only considered them once.

Answering RQ2: Discrepancies arise due to interaction and aggregation of several factors: File filtering and design decisions such as the order of data extraction steps and data

schema influence identified commits and, consequently, files, entities and developers. The choice (and configuration) of code structure parsers may further limit obtained entities. Limitations in identity matching cause major discrepancies in identified developers and derived collaboration networks. Different definitions of collaboration strength affect edge weights in the developer network.

C. Adjustments for Improving Similarity

Based on RQ2 findings, we summarise adjustments in parametrisation, code and post-processing extensions, and give

Intersection of  Codeface and Kaiaulu (P)  Codeface and Kaiaulu (R)

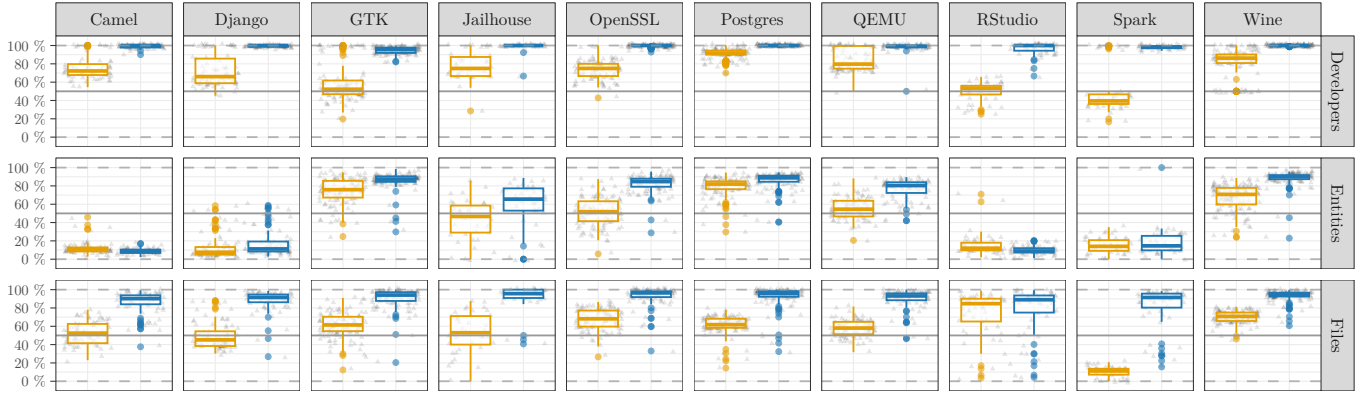


Fig. 4. Jointly identified files, entities and developers, considering prior (P) and replication (R) configurations over the entire VCS history. Again, we observe that the magnitude of discrepancies depends on the project, the considered metric and the tool configuration.

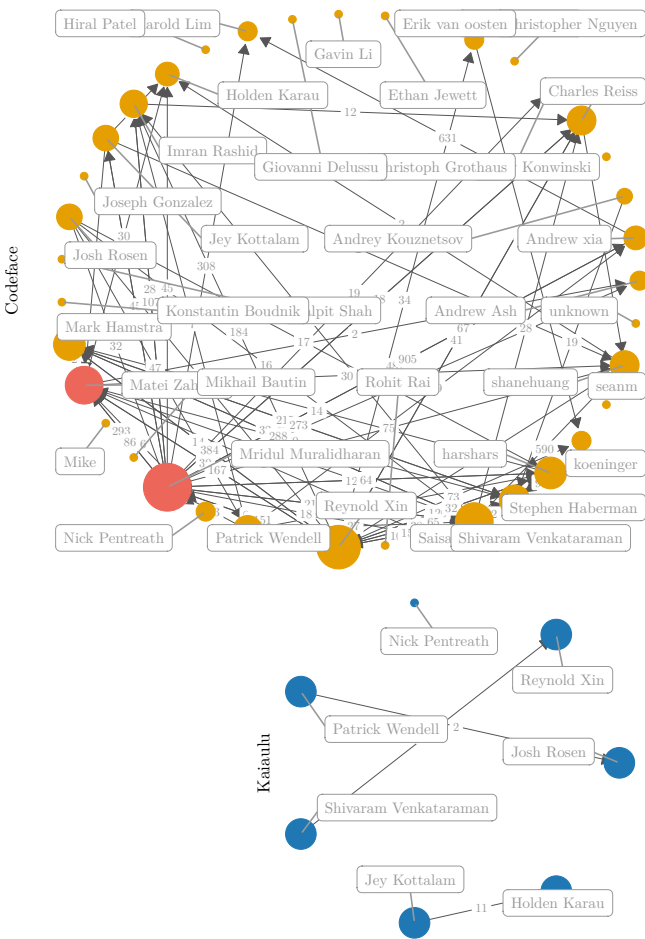


Fig. 5. Developer networks constructed by Codeface (orange) and Kaiaulu with prior configuration (blue) for Spark from March–June '13. Two very active developers missed by Kaiaulu are marked red. The discrepancies in networks are due to Kaiaulu not supporting the Scala language.

effort estimates to assess practical feasibility of the measures.

Commits: To overcome the substantial difference in extracted commits, we define identical file filters (as proposed

by Codeface) in Kaiaulu’s replication configuration, and deactivate file path filters not implemented in Codeface. While this does not require tool changes, we face a trade-off between detected commits, files and developers in Kaiaulu: Either deactivate file filtering and keep all commits with all developers, but include many non-code files such as documentation in the file count. Alternatively, activate file filtering and miss commits to these files, including developers. In both scenarios, discrepancies remain. Addressing this issue requires either changing Codeface’s data schema or Kaiaulu’s, requiring major changes. A workaround is a second Kaiaulu run *without* file filtering, and use this second table to calculate commit and developer metrics, but rely on the tables obtained *with* file filtering for file and entity metrics. Despite the use in this study, possible inconsistencies make this method unattractive.

To address commit shifts caused by time interval splitting, we introduce new configuration options to Kaiaulu that (a) allow us to include or exclude boundaries in time windows, and (b) switch between author and committer timestamps. This leads to closely matching commit time series, as seen in Fig. 3, particularly for Camel, Spark and Wine, and confirmed quantitatively by the decrease in NCD and DTW distance and the increase in correlation in Tab. II.

Files: Having two versions of Kaiaulu’s git log table, as previously motivated, provides closely matching file counts (see Fig. 3). Time series similarity increases particularly for Spark. Tab. II confirms that especially for DTW distance. With the adjustments, the percentage of jointly identified files in Fig. 4 increased to a median of over 90% in all projects.

Entities: We adapted C-Tags code structure tags in Kaiaulu’s replication configuration. However, using identical language-specific and generic C-Tags as Codeface could not increase similarity, also for previously neglected languages such as Scala, despite slight improvements seen in Figures 3, 4 and Tab. II. Adding targeted debugging statements to the code reveals that the C-Tags output provides unused information (e.g., Scala functions) when using the replication configura-

tion.¹ Overcoming this limitation would require to adjust the C-Tags parsing in Kaiaulu. Still, discrepancies could remain, given that Codeface incorporates a Doxygen and custom SQL parser for code structure identification, and both would need to be implemented in Kaiaulu for an exact replication. As we estimate an implementation time of weeks to months, we consider such adjustments impracticable.

Developers: We introduced configuration options to include several columns in Kaiaulu’s identity matching. As identity matching in Kaiaulu is still performed *within* individual columns and tables, we implemented post-processing (400 LoC) to unify identities *across* columns and tables.² As shown in Figs. 3, 4 and Tab. II for Kaiaulu’s replication configuration, this results in a significantly improved similarity, with a median percentage of jointly identified developers of over 90% for all projects.

Developer Networks: To address different interpretations of developer collaboration strength, we implemented the edge-weight scheme identified in RQ2 in Kaiaulu (100 LoC). Tab. III shows higher mean edge weights compared to the prior configuration for all subject projects. In most cases, replication weights are similar to Codeface. However, Kaiaulu can lead to outliers in edge weights, such as 10,429,560 for QEMU, compared to 261,362 with Codeface. Manually inspecting the developer networks and entities found by Kaiaulu’s replication configuration identifies a high self-collaboration on test files as cause. For the same reasons, Kaiaulu finds a mean edge weight of 92,754 compared to 7019 by Codeface for RStudio. Thus, the new weight scheme introduced in Kaiaulu cannot universally increase similarity to Codeface, as likewise observed for graph edit distance.

Discrepancies also remain as not all language C-Tags are supported by Kaiaulu. Even though Fig. 3 shows that the number of developers matches closely for the replication configuration, a developer *must* appear in the entity tables to be considered in network construction. Entity tables, however, only capture developers who contributed to entities with supported C-Tags; developers who only contributed in unsupported languages thus can still be missing. This also applies to the Scala developers in Spark (Fig. 5). After the adjustments, duplicate identities were merged and the extended structure tags resulted in more edges, leading to smaller networks with higher mean density (Tab. III). Yet, we could not reach the desired similarity to Codeface’s developer networks.

Adjustments were made iteratively, as visualised in Fig. 1, as small details such as obtaining edge weight formulae required several days of effort. The required analysis re-runs

¹Codeface uses a library parser, Kaiaulu relies on internal string matching. Since the line format reported by C-Tags for languages such as Scala differs from the format in the prior configuration, Kaiaulu cannot process them.

²Identity matching inserts identities from Kaiaulu’s git log table and all git entity tables in the respective temporal order into a new single source of truth table. If a first match by name and e-mail address fails, we try to match based on each of their e-mail addresses. If this fails, developers are matched based on full name. Only if neither an e-mail address nor their full name matches, a new identity is created. Then, all previous identities in the git and entity tables are replaced by the respective unique identities from the identity table.

took about a week. More than ten such runs were needed to achieve the presented replication configuration.

Answering RQ3: For commit, developer and file metrics, we achieve very close results by adjusting tool parametrisation, implementing tool extensions and post-processing, and applying workarounds. The identified entities as basis for network construction could not be aligned more closely; this would require a major rework of one of the tools.

V. LESSONS LEARNED

Commit Analysis: Even in extracting commits from VCS histories (step (1) in Fig. 2), tools may differ in terms of filtering, aggregation and information storage. Evolutionary analyses divide the history into time intervals by multiple approaches: Kaiaulu relies on author timestamps, while Codeface uses committer timestamps. Also time interval limits can be handled differently. More severely, different commits and files may be captured: While Codeface stores *all* commits and later filters file endings implicitly, Kaiaulu defines file filters in the configuration file in advance. Differing file path filters (*e.g.*, for test files) may lead to more discrepancies. This is caused since MSR tools are developed with initial intentions: Codeface’s commit analysis is motivated by constructing developer networks. When code structures such as functions are used as source, higher-level entities such as files are not relevant, and only file dependencies in which a code structure of interest has been found and edited are stored. However, an external user could assume that all files can be found in the corresponding database table, which can lead to false conclusions.

To avoid such misinterpretation, Kaiaulu encourages data verification prior to each analysis step. This is essential to find anomalies such as partially matched developer identities: Kaiaulu failed to merge authors and committers correctly due to different e-mail addresses and names in both columns. Since Kaiaulu used author timestamps in all analysis steps, merging identities across columns was not required, and duplicate or inconsistent identities in both columns were not an issue. However, when expecting the behaviour of Codeface, which merges identities across tables and columns, discrepancies that strongly influence count-based developer metrics arose. Creating a unified identity table as a single source of truth for all person-related information is advisable.

Data updates may also be handled differently by the tools. While Kaiaulu parses commits using Perceval and fixes its output as a baseline for commits, Codeface performs updates of its commit table in the further course of the pipeline. For instance, during the git blame analysis, it may add previously overlooked commits from the git blame output for analysis.

Git Blame Analysis: Code structures such as functions in step (2) in Fig. 2 are determined by third-party tools (C-Tags, Doxygen, ...) from git blame data. Our results show that choice of parser and detail configuration can strongly influence statistics. For instance, Kaiaulu relies on C-Tags, whereas Codeface uses it when Doxygen and the built-in SQL parser are not applicable. C-Tags configurations differ depending on programming languages and structure tags. Even though both

tools have the same amount of tags available, they don't handle them the same way. First, Codeface applies a default set of tags to *any* language file not filtered out in step (1) in Fig. 2, while Kaiaulu requires an explicit specification of tags. Secondly, third-party tool outputs are parsed differently.

If not explicitly communicated, such details can easily be lost: Developers contributing code in a particular language were overlooked by Kaiaulu, which could only be inferred from a qualitative comparison of the baseline data. This is problematic for large-scale mining without qualitative inspection. Therefore, users and tool developers should check which project properties and third-party tool settings are actually supported, and verify their data qualitatively to reduce uncertainty.

Network Construction: Herbold *et al.* [2] report a lack of guidelines for developer social network research. For instance, a question of interpretation arises from edge weighting. Original and replication tool consider collaboration strength as a measure of contributed LoC to an entity, but in the absence of a precise definition, the influence of past contributions remains unclear, and edge weights can vary substantially. Thus, subsequent analyses such as core developer and community detection for team structuring could lead to different results.

VI. THREATS TO VALIDITY

Construct Validity: The comparisons in our study can be based on a wide range of data sources and metrics. Therefore, there is a risk that we draw conclusions based on the wrong choices. We reduced this risk by using established parameter choices from previous studies and by evaluating multiple metrics simultaneously. However, due to the large number of customisable parameters, we did not test all of them in detail. For example, we fixed analysis time intervals to non-overlapping three-month windows. A larger window size could lead to other results, although studies reported a minor effect [20]. Also, we did not include communication data from e-mails and issue trackers, as this exceeds the scope of the study.

We did not evaluate if our results correctly reflect reality or whether both tools produce wrong results. Earlier studies using Codeface addressed this issue through interviews with a large group of OSS developers, finding that network construction at function level [1] is perceived as an accurate representation of collaboration. Thus, we consider this threat as minor.

Internal Validity: We identified discrepancies between tools in an automated way, but additional manual investigations were needed to explore the root causes. Since divergences were high in absolute numbers, we could not look at each of them individually. Instead, we selected a random subset for inspection. Although we are confident that major causes of differences are identified, as including more instances did not reveal new causes, it is possible that minor edge cases remain.

External Validity: Our study is based on only two MSR tools. Both were developed by different research groups, but the groups worked together at numerous prior studies and thus are assumed to have a common understanding of assumptions. While this facilitates a fair comparison, it also introduces a bias. Tools from other research groups could make

fundamentally different assumptions and lead to even more divergent results. Given the limited scope of the study, the inclusion of further tools would have not allowed us to address the diverse set of MSR pipeline stages, which we prioritised to evaluate the impact and interplay of discrepancies.

Another threat is that we only analysed a set of ten subject projects. To limit the impact on generality, we considered a wide range of application areas, programming languages, project and team sizes. Further projects written in other programming languages could have led to much more varying results. For example, for project WordPress mainly written in PHP, Kaiaulu did not find any entities as PHP is not supported. We will look at more diverse repositories in future studies.

VII. DISCUSSION AND CONCLUSION

Our study shows that uncertainties in MSR pipelines can lead to substantial discrepancies for the extracted baseline data. Minor implementation details can have considerable consequences: Two tools sharing similar goals, methods and processing steps can output substantially different results.

Some of the discrepancies are caused by technical limitations, such as missing language support. Other differences arise from subjective choices, for instance how strongly past collaborations should influence edge weights in developer networks. Many of these factors cannot be easily adjusted by changing tool parametrisation or implementing minor extensions and post-processing, but require a substantial rework of tool capabilities to ascertain result agreement. Therefore, differences remain in our replication. It is likely that discrepancies would be even more pronounced for other combinations of MSR tools not developed by collaborating groups.

Our work underscores the need for MSR replication studies to promote a common understanding and higher standardisation of mining processes. To revisit the methodology, a closer investigation of the implications of identified discrepancies on the generality of conclusions drawn in empirical software engineering is needed. Therefore, we plan a follow-up study with a systematic literature review following the guidelines proposed by Kitchenham *et al.* [63], [64] to replicate the most influential findings in recent decades with diverse tools.

Finally, we argue that assumptions made by each tool or study should be described as fully as possible and supplemented by code. Considering these technical details in replication studies and in the design and evaluation of novel tools and pipelines can improve fairness and reliability. We hope that the process and findings from this replication will be helpful as a reference for these efforts in the future.

VIII. REPRODUCTION PACKAGE

For understandability and to facilitate future comparisons, we provide containerised study tools and scripts in our GitHub repository. A self-contained reproduction package and Docker image including OS, libraries, input and results data [34] is available at Zenodo.

Acknowledgements: WM and NH acknowledge support by the High-Tech Agenda of the Free State of Bavaria.

REFERENCES

- [1] M. Joblin, W. Mauerer, S. Apel, J. Siegmund, and D. Riehle, "From developer networks to verified communities: A fine-grained approach," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, 2015, pp. 563–573.
- [2] S. Herbold, A. Amirfallah, F. Trautsch, and J. Grabowski, "A systematic mapping study of developer social network research," *Journal of Systems and Software*, vol. 171, p. 110802, 2021.
- [3] T. Bock, C. Hunsen, M. Joblin, and S. Apel, "Synchronous development in open-source projects: A higher-level perspective," *Automated Software Engineering*, vol. 29, no. 1, p. 3, 2021.
- [4] C. Hunsen, J. Siegmund, and S. Apel, "On the fulfillment of coordination requirements in open-source software projects: An exploratory study," *Empirical Software Engineering*, vol. 25, no. 6, pp. 4379–4426, 2020.
- [5] M. Joblin, S. Apel, and W. Mauerer, "Evolutionary trends of developer coordination: A network approach," *Empirical Software Engineering*, vol. 22, no. 4, pp. 2050–2094, 2017.
- [6] T. Bock, A. Schmid, and S. Apel, "Measuring and modeling group dynamics in open-source software development: A tensor decomposition approach," *ACM Trans. Softw. Eng. Methodol.*, vol. 31, no. 2, 2021.
- [7] M. Joblin, S. Apel, C. Hunsen, and W. Mauerer, "Classifying developers into core and peripheral: An empirical study on count and network metrics," in *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*, 2017, pp. 164–174.
- [8] T. Bock, N. Alznauer, M. Joblin, and S. Apel, "Automatic core-developer identification on GitHub: A validation study," *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 6, 2023.
- [9] M. Joblin, B. Eckl, T. Bock, A. Schmid, J. Siegmund, and S. Apel, "Hierarchical and hybrid organizational structures in open-source software projects: A longitudinal study," *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 4, 2023.
- [10] D. A. Tamburri and F. Palomba, "Evolving software forges: An experience report from Apache Allura," *Software: Evolution and Process*, vol. 33, no. 12, 2021.
- [11] D. A. Tamburri, F. Palomba, and R. Kazman, "Exploring community smells in open-source: An automated approach," *IEEE Trans. Softw. Eng.*, vol. 47, no. 3, pp. 630–652, 2021.
- [12] G. Catolino, F. Palomba, D. A. Tamburri, A. Serebrenik, and F. Ferrucci, "Gender diversity and women in software teams: How do they affect community smells?" in *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*, 2019, pp. 11–20.
- [13] N. Almarimi, A. Ouni, and M. W. Mkaouer, "Learning to detect community smells in open source software projects," *Knowledge-Based Systems*, vol. 204, p. 106201, 2020.
- [14] M. De Stefano, F. Pecorelli, D. A. Tamburri, F. Palomba, and A. De Lucia, "Splicing community patterns and smells: A preliminary study," in *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops (ICSEW)*. Association for Computing Machinery, 2020, pp. 703–710.
- [15] S. Lambiase, G. Catolino, D. A. Tamburri, A. Serebrenik, F. Palomba, and F. Ferrucci, "Good fences make good neighbours? On the impact of cultural and geographical dispersion on community smells," in *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*, ser. '22. Association for Computing Machinery, 2022, pp. 67–78.
- [16] F. Palomba and D. A. Tamburri, "Predicting the emergence of community smells using socio-technical metrics: A machine-learning approach," *Journal of Systems and Software*, vol. 171, p. 110847, 2021.
- [17] F. Palomba, D. Andrew Tamburri, F. Arcelli Fontana, R. Oliveto, A. Zaidman, and A. Serebrenik, "Beyond technical aspects: How do community smells influence the intensity of code smells?" *IEEE Trans. Softw. Eng.*, vol. 47, no. 1, pp. 108–129, 2021.
- [18] M. D. Stefano, E. Iannone, F. Pecorelli, and D. A. Tamburri, "Impacts of software community patterns on process and product: An empirical study," *Science of Computer Programming*, vol. 214, p. 102731, 2022.
- [19] B. Eken, F. Palma, B. Ayşe, and T. Ayşe, "An empirical study on the effect of community smells on bug prediction," *Software Quality Journal*, vol. 29, no. 1, pp. 159–194, 2021.
- [20] W. Mauerer, M. Joblin, D. A. Tamburri, C. Paradis, R. Kazman, and S. Apel, "In search of socio-technical congruence: A large-scale longitudinal study," *IEEE Trans. Softw. Eng.*, vol. 48, no. 8, pp. 3159–3184, 2022.
- [21] P. Pícha, P. Brada, R. Ramsauer, and W. Mauerer, "Towards architect's activity detection through a common model for project pattern analysis," in *2017 IEEE International Conference on Software Architecture Workshops (ICSAW)*, 2017, pp. 175–178.
- [22] M. Joblin and S. Apel, "How do successful and failed projects differ? A socio-technical analysis," *ACM Trans. Softw. Eng. Methodol.*, vol. 31, 2022.
- [23] D. A. Tamburri, K. Blincoe, F. Palomba, and R. Kazman, "The canary in the coal mine... A cautionary tale from the decline of SourceForge," *Software: Practice and Experience*, vol. 50, no. 10, 2020.
- [24] C. Bird, B. Murphy, N. Nagappan, and T. Zimmermann, "Empirical software engineering at Microsoft Research," in *ACM 2011 Conference on Computer Supported Cooperative Work (CSCW)*, 2011, pp. 143–150.
- [25] Q. Xuan and V. Filkov, "Building it together: synchronous development in OSS," in *Proceedings of the 36th International Conference on Software Engineering (ICSE)*. ACM, 2014, pp. 222–233.
- [26] Y. Tymchuk, A. Mocchi, and M. Lanza, "Collaboration in open-source projects: Myth or reality?" in *11th Working Conference on Mining Software Repositories (MSR)*, 2014, pp. 304–307.
- [27] C. Paradis and R. Kazman, "Building the MSR tool Kaiaulu: Design principles and experiences," in *Software Architecture*. Springer International Publishing, 2022, pp. 107–129.
- [28] W. Mauerer and M. Joblin, "Codeface," 2013. [Online]. Available: <https://github.com/lfd/codeface>
- [29] C. Paradis and R. Kazman, "Kaiaulu," 2020. [Online]. Available: <https://github.com/sailuh/kaiaulu/>
- [30] V. Basili, F. Shull, and F. Lanubile, "Building knowledge through families of experiments," *IEEE Trans. Softw. Eng.*, vol. 25, no. 4, pp. 456–473, 1999.
- [31] G. Robles, "Replicating MSR: A study of the potential replicability of papers published in the Mining Software Repositories proceedings," in *2010 7th IEEE Working Conference on Mining Software Repositories (MSR)*, 2010, pp. 171–180.
- [32] G. Ghezzi and H. C. Gall, "Replicating mining studies with SOFAS," in *2013 10th Working Conference on Mining Software Repositories (MSR)*, 2013, pp. 363–372.
- [33] F. Trautsch, S. Herbold, P. Makedonski, and J. Grabowski, "Addressing problems with replicability and validity of repository mining studies through a smart data platform," *Empirical Software Engineering*, vol. 23, no. 2, pp. 1036–1083, 2018.
- [34] W. Mauerer, S. Klessinger, and S. Scherzinger, "Beyond the badge: Reproducibility engineering as a lifetime skill," in *4th International Workshop on Software Engineering Education for the Next Generation (SEENG)*. ACM, 2022, pp. 1–4.
- [35] J. M. Gonzalez-Barahona and G. Robles, "Revisiting the reproducibility of empirical software engineering studies based on data retrieved from development repositories," *Information and Software Technology*, vol. 164, p. 107318, 2023.
- [36] J. T. Liang, C. Badea, C. Bird, R. DeLine, D. Ford, N. Forsgren, and T. Zimmermann, "Can GPT-4 replicate empirical software engineering research?" *Proc. ACM Softw. Eng.*, vol. 1, no. FSE, 2024.
- [37] F. J. Shull, J. C. Carver, S. Vegas, and N. Juristo, "The role of replications in Empirical Software Engineering," *Empirical Software Engineering*, vol. 13, no. 2, pp. 211–218, 2008.
- [38] J. M. González-Barahona and G. Robles, "On the reproducibility of empirical software engineering studies based on data retrieved from development repositories," *Empirical Software Engineering*, vol. 17, no. 1, pp. 75–89, 2012.
- [39] T. T. Dinh-Trong and J. M. Bieman, "The FreeBSD project: A replication case study of open source development," *IEEE Trans. Softw. Eng.*, vol. 31, no. 6, pp. 481–494, 2005.
- [40] Z. Mahmood, D. Bowes, T. Hall, P. C. R. Lane, and J. Petrić, "Reproducibility and replicability of software defect prediction studies," *Information and Software Technology*, vol. 99, pp. 148–163, 2018.
- [41] M. Di Penta, G. Bavota, and F. Zampetti, "On the relationship between refactoring actions and bugs: A differentiated replication," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM, 2020, pp. 556–567.
- [42] C. Bird, P. C. Rigby, E. T. Barr, D. J. Hamilton, D. M. German, and P. Devanbu, "The promises and perils of mining git," in *2009 6th IEEE International Working Conference on Mining Software Repositories (MSR)*, 2009, pp. 1–10.

- [43] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian, "The promises and perils of mining GitHub," in *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR)*. ACM, 2014, pp. 92–101.
- [44] S. W. Flint, J. Chauhan, and R. Dyer, "Pitfalls and guidelines for using time-based Git data," *Empirical Software Engineering*, vol. 27, no. 7, p. 194, 2022.
- [45] N. Saarimäki, S. Moreschini, F. Lomio, R. Penalzoza, and V. Lenarduzzi, "Towards a robust approach to analyze time-dependent data in software engineering," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2022, pp. 36–40.
- [46] R. Nia, C. Bird, P. Devanbu, and V. Filkov, "Validity of network analyses in Open Source Projects," in *2010 7th IEEE Working Conference on Mining Software Repositories (MSR)*, 2010, pp. 201–209.
- [47] A. Meneely and L. Williams, "Socio-technical developer networks: Should we trust our measurements?" in *Proceedings of the 33rd International Conference on Software Engineering (ICSE)*. ACM, 2011, pp. 281–290.
- [48] J. Siegmund, N. Siegmund, and S. Apel, "Views on internal and external validity in empirical software engineering," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 1, 2015, pp. 9–19.
- [49] J. Lefever, Y. Cai, H. Cervantes, R. Kazman, and H. Fang, "On the lack of consensus among technical debt detection tools," in *Proceedings of the 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE Press, 2021, pp. 121–130.
- [50] M. Goeminne and T. Mens, "A comparison of identity merge algorithms for software repositories," *Science of Computer Programming*, vol. 78, no. 8, pp. 971–986, 2013.
- [51] S. Amreen, A. Mockus, R. Zaretski, C. Bogart, and Y. Zhang, "AL-FAA: Active Learning Fingerprint based Anti-Aliasing for correcting developer identity errors in version control systems," *Empirical Software Engineering*, vol. 25, no. 2, pp. 1136–1167, 2020.
- [52] M. V. Bertoncello, G. Pinto, I. S. Wiese, and I. Steinmacher, "Pull requests or commits? Which method should we use to study contributors' behavior?" in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 592–601.
- [53] S. Panichella, G. Bavota, M. D. Penta, G. Canfora, and G. Antoniol, "How developers' collaborations identified from different sources tell us about code changes," in *2014 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2014, pp. 251–260.
- [54] C. Paradis, R. Kazman, and D. Tamburri, "Analyzing the tower of babel with Kaiaulu," *Journal of Systems and Software*, vol. 210, p. 111967, 2024.
- [55] A. Danial, "cloc: v1.92," Dec. 2021.
- [56] C. Paradis, R. Kazman, and M. Konrad, "A socio-technical perspective on software vulnerabilities: A causal analysis," *Information and Software Technology*, vol. 176, p. 107553, 2024.
- [57] H. Mumtaz, P. Singh, and K. Blincoe, "Analyzing the relationship between community and design smells in open-source software projects: An Empirical Study," in *Proceedings of the 16th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. ACM, 2022, pp. 23–33.
- [58] H. Mumtaz, C. Paradis, F. Palomba, D. A. Tamburri, R. Kazman, and K. Blincoe, "A preliminary study on the assignment of GitHub issues to issue commenters and the relationship with social smells," in *Proceedings of the 15th International Conference on Cooperative and Human Aspects of Software Engineering (CHASE)*. Association for Computing Machinery, 2022, pp. 61–65.
- [59] C. Bird, A. Gourley, P. Devanbu, M. Gertz, and A. Swaminathan, "Mining email social networks," in *Proceedings of the 2006 International Workshop on Mining Software Repositories (MSR)*. Association for Computing Machinery, 2006, pp. 137–143.
- [60] R. Cilibrasi and P. Vitanyi, "Clustering by compression," *IEEE Transactions on Information Theory*, vol. 51, no. 4, pp. 1523–1545, 2005.
- [61] D. J. Berndt and J. Clifford, "Using dynamic time warping to find patterns in time series," in *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining (AAAIWS)*. AAAI Press, 1994, pp. 359–370.
- [62] R. T. Dean and W. T. M. Dunsmuir, "Dangers and uses of cross-correlation in analyzing time series in perception, performance, movement, and neuroscience: The importance of constructing transfer function autoregressive models," *Behavior Research Methods*, vol. 48, no. 2, pp. 783–802, 2016.
- [63] B. Kitchenham, "Procedures for performing systematic reviews," *Keele, UK, Keele Univ.*, vol. 33, 2004.
- [64] B. Kitchenham and P. Brereton, "A systematic review of systematic review process research in software engineering," *Information and Software Technology*, vol. 55, no. 12, pp. 2049–2075, 2013.