

How Do Communities of ML-Enabled Systems Smell? A Cross-Sectional Study on the Prevalence of Community Smells

Giusy Annunziata
gannunziata@unisa.it
University of Salerno
Salerno, Italy

Stefano Lambiase
slambiase@unisa.it
University of Salerno
Salerno, Italy

Fabio Palomba
fpalomba@unisa.it
University of Salerno
Salerno, Italy

Gemma Catolino
gcatolino@unisa.it
University of Salerno
Salerno, Italy

Filomena Ferrucci
fferrucci@unisa.it
University of Salerno
Salerno, Italy

Abstract

Successful software development depends on effectively managing both collaboration and technology. However, socio-technical challenges can disrupt team dynamics and lead to technical debt. Despite the interdisciplinary nature of teams working on ML-enabled systems, research on their socio-technical dynamics remains limited compared to the emphasis on technical aspects. This study aims to address this gap by examining the prevalence, evolution, and interrelations of “community smells”, i.e., social anti-patterns that indicate dysfunctional collaboration, in open-source ML projects. We conducted an empirical study on 188 repositories from the NICHE dataset. Leveraging the CADOCs tool, we identified and analyzed community smells within these repositories. Our analysis focused on three key aspects: (1) the prevalence of community smells, (2) their correlations, and (3) their variations over time. Our findings indicate that while some community smells are more prevalent than others, their overall occurrence remains relatively stable over time; Prima Donna Effects, Sharing Villainy, and Solution Defiance are particularly prominent in ML-enabled projects compared to Radio Silence or Organizational Skirmish. These insights provide valuable guidance for project managers and team leads in ML-focused communities, helping them mitigate social challenges more effectively, allocate resources more effectively, and improve collaboration and team dynamics.

CCS Concepts

• **Software and its engineering** → Open source model.

Keywords

Social Aspects in Software Engineering; Community Smells; Software Engineering for Artificial Intelligence; Mining Software Repositories; Empirical Software Engineering.

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 the author(s) 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.

EASE 2025, 17–20 June, 2025, Istanbul, Türkiye

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXXX.XXXXXXX>

ACM Reference Format:

Giusy Annunziata, Stefano Lambiase, Fabio Palomba, Gemma Catolino, and Filomena Ferrucci. 2018. How Do Communities of ML-Enabled Systems Smell? A Cross-Sectional Study on the Prevalence of Community Smells. In *Proceedings of The 29th International Conference on Evaluation and Assessment in Software Engineering (EASE 2025)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

The success of a software development project depends not only on the technical aspects of the product but also on effective collaboration and communication within the development team [16, 25]. Social dynamics can directly impact technical processes, influencing both productivity and software quality [19]. However, such social challenges, being abstract factors rooted in communication and collaboration, are inherently complex to detect and quantify [24]. The literature has investigated “*Community Smells*”, which are measurable patterns of organizational problems that hinder cooperation and affect project health [26]. These smells often lead to increased social debt, where unresolved social issues accumulate over time, degrading software quality and team cohesion [7, 19].

Machine Learning (ML)-enabled systems introduce additional complexity when analyzing social dynamics. By their nature, these systems rely on multidisciplinary teams composed of developers, software engineers, and data scientists, each bringing distinct educational background, perspectives, and expertise [18]. While this diversity can foster innovation and improve problem-solving, it can also create social tensions, misalignments, and communication gaps that may exacerbate the emergence of community smells [5].

Despite the critical role of social aspects in the shaping of ML-enabled systems, they remain largely underexplored compared to technical aspects. Most research to date has primarily focused on issues such as technical debt, suboptimal design decisions, or neglected maintenance tasks [9, 22], often leaving the social dimension insufficiently explored [14].

An initial investigation in this direction was conducted by Mailach and Siegmund [14], who explored socio-technical anti-patterns identified by the communities of ML-enabled systems, analyzing their causes and potential mitigation strategies. Building on this, Annunziata et al. [4] examined the correlation between these anti-patterns and known community smells, providing the first insight into community smells within communities of ML-enabled systems.

Their findings highlight that the interaction between diverse teams and the complex workflows characterizing those projects can exacerbate community smells and social issues, leading to increased social debt and weakening collaboration and productivity. Therefore, analyzing these socio-technical issues is critical to developing strategies that improve team cohesion, mitigate social tensions, and improve the overall health and results of software projects.

Our study aims to deepen the understanding and the examination of the socio-technical well-being of communities developing ML-enabled systems, providing a first investigation of the prevalence and evolution of community smells in real open-source ML-enabled systems.

We conducted our study on 188 GitHub ML-enabled open-source projects from the NICHE dataset [30], identifying and detecting 10 distinct community smells, e.g., Prima Donna Effect and Toxic Communication show the adoption of unhealthy and rude communication, superiority, and uncooperative behavior, using the CADOCS tool [28].

Next, we performed a Cross-Sectional Study to examine the prevalence of each community smell in the analyzed projects and explore correlations between them. Finally, we analyzed the evolution of community smells over time through a Longitudinal Study, examining each project in 3-month intervals from its inception to 2 years of age.

The results obtained show that Prima Donna Effect (PDE)—which refers to a superiority, disagreement, and uncooperative behavior—is the most prevalent smell, with a prevalence of 92.6%. Furthermore, the longitudinal analysis reveals that PDE remains consistently high over time (95–96%). In contrast, Radio Silence (RS)—where formal communication barriers between sub-communities cause time loss—is the least frequent, observed in 18.6% of the cases. Meanwhile, other community smells tend to fluctuate over time, for example, Unhealthy Interaction (UI)—the adoption of unhealthy behavior among team members—decreased by 24% to 6%, suggesting improvements in communication practices. Finally, the results suggest a strong correlation between PDE and other community smells, in particular the Organizational Silo Effect (OSE)—which describes isolated communities that communicate with each other only through a few members.

Our findings offer valuable insights for project managers in ML-enabled projects, equipping them with a deeper understanding of the most common social challenges and practical strategies to address them. Additionally, these results serve as a strong foundation for researchers, providing an initial exploration and quantitative analysis of community smells in the ML-enabled context. By highlighting their evolution during the early stages of software projects, this work sets the stage for future investigations into the interplay between human dynamics and their impact on the technical success of ML-enabled projects.

Structure of the paper. Section 2 presents background and related work on ML-enabled systems and community smells. Section 3 outlines the design of the study and the methodology for each research question. Section 4 reports the results of the study. Section 5 discusses the insights of the paper and practical implications, while Section 6 the threats to validity. Finally, Section 7 concludes the paper and outlines future investigations.

2 Background and Related Work

In this section, we discuss the background and related work that underpin our study.

2.1 ML-Enabled Systems

Machine Learning-enabled systems (ML-Enabled Systems) are a specialized category of software systems that integrate machine learning models to automate and enhance decision-making processes [15, 27]. These systems are increasingly deployed across various domains, including autonomous vehicles and healthcare, leveraging ML algorithms to analyze data, recognize patterns, and continuously improve through learning. Given the complexity of developing and integrating ML components into a system, both technical and social challenges arise throughout the development of an ML-enabled system. Technical debt is a clear example of a technical challenge, referring to the accumulation of suboptimal design decisions or deferred maintenance tasks that ultimately increase the cost of future system improvements and maintenance [32]. A significant aspect of technical debt is the presence of code smells and data smells—indicators of underlying issues in software and data management [9, 13, 22]. Code smells manifest in various ways and become particularly relevant in the context of ML models, data pipelines, and experimentation workflows. Similarly, data smells—issues related to data collection, storage, and preprocessing—represent a significant source of technical debt, potentially leading to biases, reduced model accuracy, and compromised fairness [9]. Recupito et al. [22] further underscores the impact of data smells in AI systems, linking them to challenges such as biased predictions and degraded performance.

Social factors also play a critical role in the accumulation of technical debt in ML-enabled systems. In particular, social challenges are significant due to the *heterogeneous nature* of the teams involved, which often include diverse roles such as data scientists and software engineers. While this diversity can foster innovation, it can also introduce socio-technical issues. Misunderstandings regarding data requirements or deployment constraints can lead to inefficiencies, rework, and increased technical debt, especially when changes to a system component are not effectively communicated across the team [17].

For instance, Busquim et al. [5] conducted semi-structured interviews with software engineering and data science professionals, revealing challenges such as disparities in technical expertise, ambiguous roles, and inadequate documentation. These challenges exacerbate inefficiencies and technical debt, underscoring the need for improved collaboration and communication strategies [17].

Similarly, Nahar et al. [17] investigated collaboration and communication challenges during ML-enabled system development. They found that miscommunication, rooted in the differing technical vocabularies among the various team members, was a primary issue. Such communication gaps hinder the understanding between developers and managers, further complicating project workflows.

In another study, Mailach and Siegmund [14] explored socio-technical aspects in ML-enabled systems by analyzing 73 relevant

videos from the ML-Ops community.¹ They identified 17 socio-technical anti-patterns, which describe problematic developer behaviors, along with 16 causes and 15 organizational strategies to address them. However, the study raises an open question about such anti-patterns, asking whether they are specific to ML-enabled systems or are simply manifestations of socio-technical challenges already investigated in the context of software engineering.

2.2 Community Smells in the Context of ML-enabled Systems

Community smells, introduced by Tamburri et al. [26], are measurable socio-technical patterns that negatively impact team collaboration. These issues manifest themselves through fragmented communication, isolated work practices, and imbalanced decision-making, contributing to social debt—organizational inefficiencies that hinder productivity. One notable example is the Prima Donna Effect, where dominant individuals disregard others’ input, reducing collaboration and team morale [2, 25].

The relevance of community smells is reinforced by the concept of socio-technical congruence, which highlights the interdependence between social and technical structures in software projects [25]. Studies suggest that unresolved social issues can influence software quality, as shown by Palomba et al. [19], who found that community smells are linked to persistent code smells and technical debt. Developers often attribute software quality issues to communication breakdowns and coordination challenges within teams. Several factors influence the emergence of community smells [3, 6, 8, 12]. Gender diversity, for instance, has been associated with reduced occurrences of these issues, as diverse teams tend to foster more inclusive and communicative environments [8]. Similarly, cultural and geographical diversity can introduce both benefits and challenges: while it encourages innovation, it may also lead to coordination difficulties in distributed teams [12]. Additionally, leadership roles such as architects and team leaders often become knowledge bottlenecks, where critical expertise is concentrated in a few individuals, leading to inefficiencies if knowledge is not effectively shared [6].

Mitigation strategies have been explored to reduce the impact of community smells. Catolino et al. [7] identified approaches such as improved communication, structured collaboration, and inclusive decision-making to address these socio-technical challenges.

While community smells have been extensively studied in traditional software development, their role in ML-enabled systems remains underexplored. Given the heterogeneous nature of ML communities—often composed of data scientists, software engineers, and domain experts—the risk of socio-technical misalignment is particularly high. Initial investigations suggest that community smells may manifest differently in ML projects. Annunziata et al. [3] found that in open-source projects using Python, a language dominant in ML, the Prima Donna Effect is more prevalent, indicating potential social challenges in ML development. Expanding on this, Annunziata et al. [4] identified correlations between community smells and identified socio-technical anti-patterns in the communities of ML-enabled systems, positioning these patterns as both causes and consequences of organizational inefficiencies.

¹ML-Ops Community: <https://mlops.community/>

These findings suggest that community smells could serve as indicators for assessing the social health of teams of ML-enabled systems, offering a new lens to study and mitigate challenges unique to this context. However, the extent to which heterogeneous teams influence the persistence of community smells in ML-enabled projects remains largely unexplored. Addressing this gap is crucial to understanding the social dynamics of ML-enabled system development and improving team collaboration, system maintainability, and overall software quality.

Research Gap

Social aspects in software development often influence technical components, ultimately affecting system quality. In ML-enabled systems, the interdisciplinary nature of heterogeneous teams amplifies these challenges. While crucial, the study of socio-technical issues—such as measurable patterns like ‘community smells’—remains partially underexplored compared to technical challenges. This gap underscores the need to better understand their impact on ML teams and system quality. Investigating the prevalence of community smells could initially provide valuable insights to mitigate these challenges and guide future research on possibly improving team performance in ML environments.

3 Research Method

The *goal* of the study was to understand and examine the socio-technical well-being of communities developing ML-enabled systems using community smells. The *purpose* gains knowledge about the role of social anti-patterns in such communities to inform new research trends and provide useful knowledge to practitioners. The study is primarily conducted from the perspective of project managers, who seek to identify and address social anti-patterns within their heterogeneous ML-enabled teams to prevent potential negative impacts on technical outcomes. A secondary perspective is that of researchers, offering them insights into the prevalence and consequences of community smells in ML projects and laying the groundwork for further investigations. To achieve the above-mentioned *objective*, we defined a set of research questions that guided the research process. First, we wanted to characterize the extent to which socio-technical anti-patterns—represented using community smells—are present in ML-enabled open-source communities, i.e., their prevalence.

Thus, we formulated the following research question:

Q RQ₁ *What is the prevalence of Community Smells in ML-enabled projects?*

The initial analysis offered a broad overview of the prevalence of community smells across these projects. To build on this, we are conducting a more detailed time-slice analysis to achieve two primary objectives: (1) identify which community smells are most prevalent over time, and (2) develop a deeper understanding of these social issues through a more comprehensive knowledge base over time. To guide this investigation, we have formulated the following research question:

Q RQ₂ *How does the Prevalence of Community Smells in ML-enabled projects vary over time?*

Then, we focus on analyzing the correlations between different smells. With this analysis, we wanted to deepen our understanding of community smells by shedding light on how these issues interact with each other and shape team dynamics. More specifically, the analysis highlights correlations among the smells that suggest that the presence of one smell may indicate the likelihood of others occurring more or less frequently. In other words, the presence of a positive correlation between two smells informs us that when one occurs, there is a high probability that the other smell will also occur. In contrast, the presence of a negative correlation suggests that, upon the occurrence of one smell, the probability of the other smell occurring is lower than normal. This insight can be valuable to practitioners, enabling them to anticipate and address related smells once one is identified.

Thus, we formulated the following research question:

Q RQ₃ *What are the relationships between different community smells in the ML-enabled projects?*

To answer our research questions, we conducted a cross-sectional study in combination with a mining study on GitHub. A *Cross-sectional* studies provide a snapshot of a population at a specific point in time, capturing the prevalence of a condition or phenomenon and relevant exposures by collecting data simultaneously from all participants [29]. Specifically, we analyzed 188 open-source ML-enabled projects belonging to the NICHE dataset [30]. As for RQ₁, we calculated the prevalence of community smells within the 188 GITHUB ML-enabled open-source projects. For RQ₂, we performed a longitudinal analysis, dividing the timeline into 3-month intervals to examine the evolution of these community smells over time. Finally, for RQ₃ we conducted a cross-sectional study using the *Prevalence Odds Ratio* (POR) to identify potential correlations between different types of community smells. Figure 1 presents an overview of our study. In particular, reports the steps, the methods, and activities to conduct it.

To conduct our analysis, we follow the *Empirical Standards* by Ralph et al. [21]. Specifically, we employed the “*Repository Mining*”, the “*Longitudinal*”, the “*Engineering Research*” and the “*General Standard*” definitions and guidelines.²

3.1 Context of the Study and Dataset Creation

To investigate socio-technical issues in the context of ML-enabled systems, we began by creating a dataset containing community smells on ML-enabled open-source projects. As a foundation for this dataset, we used the NICHE dataset developed by Widyasari et al. [30]. The NICHE dataset, validated in 2022, consists of 572 open-source machine learning (ML) projects implemented in Python and labeled manually. Of these, 441 are classified as “engineered” because they clearly adhere to good software engineering practices, and 131 are classified as “non-engineered” because they do not meet engineering standards [30]. The dataset obtained is available in the online appendix [1].

²Empirical Standards: <https://www2.sigsoft.org/EmpiricalStandards/docs/standards>

Selection Criteria. In this section, we discuss the selection criteria applied to obtain our dataset. Our criteria are based on the objective of the study and rely on criteria adopted in similar previous studies on the matter [7].

Engineered Project: The dataset NICHE [30] was filtered to exclude non-engineered projects. This decision was driven by the objective of studying human factors and behavioral aspects in environments involving data scientists and software engineers [5, 18]. ML-enabled projects inherently feature an ML component, insinuating the presence of data scientists into the development team, while engineered projects suggest the inclusion of software engineers. This combination aligns with our goal of analyzing socio-technical interactions specific to ML-enabled development.

Age of the Project: Since the second research question focuses on analyzing the evolution of community smells over time, project age was used as a key selection criterion. Projects were categorized into 3 groups based on their lifespan: Young (<24 months), Established (24–32 months), and Mature (>32 months). Most of the selected projects fell into the Established and Mature groups.

We analyzed projects in 3-month windows, following the approach used in previous studies [7]. For consistency, each age group was analyzed over a number of windows corresponding to the youngest project in that category: 6 months (2 windows) for Young, 24 months (8 windows) for Established, and 32 months (10 windows) for Mature. This ensured comparability across projects within each age group.

Due to the low number of projects and time windows in the Young category, we decided to exclude it from the dataset, focusing only on Established and Mature projects. Furthermore, because the Established category had fewer projects than the Mature category, we merged these two groups into a single category. This consolidation allowed for a consistent analysis timeframe of up to 8 time windows (equivalent to 2 years) for all selected projects, ensuring sufficient data while maintaining analytical robustness.

Size of the Team: Following conventions in the literature [7, 26], projects were categorized based on the number of contributors: *Small* (Fewer than 10 contributors), *Medium* (10–50 contributors), *Large* (50–150 contributors), and *Very Large* (More than 150 contributors). The dataset had a low number of small projects, with very small teams, less than five or even one team member. A low number of team members is insufficient to analyze community smells on these projects, so the projects identified as small were excluded from the dataset. Our focus shifted to medium, large, and very large projects to ensure meaningful socio-technical insights.

3.2 Community Smells Detection

We began with an initial dataset of 572 ML-enabled projects sourced from the validated NICHE dataset [30] (2022). Following the selection criteria explained, we refined the pool to 188 projects. On those projects we used CADOCS [28] to detect community smells. CADOCS is a client-server conversational agent integrated into Slack, designed to extend and facilitate the use of the community smell detection tool ‘CSDETECTOR’. The goal of CADOCS is to provide practitioners with more social problem analysis support from

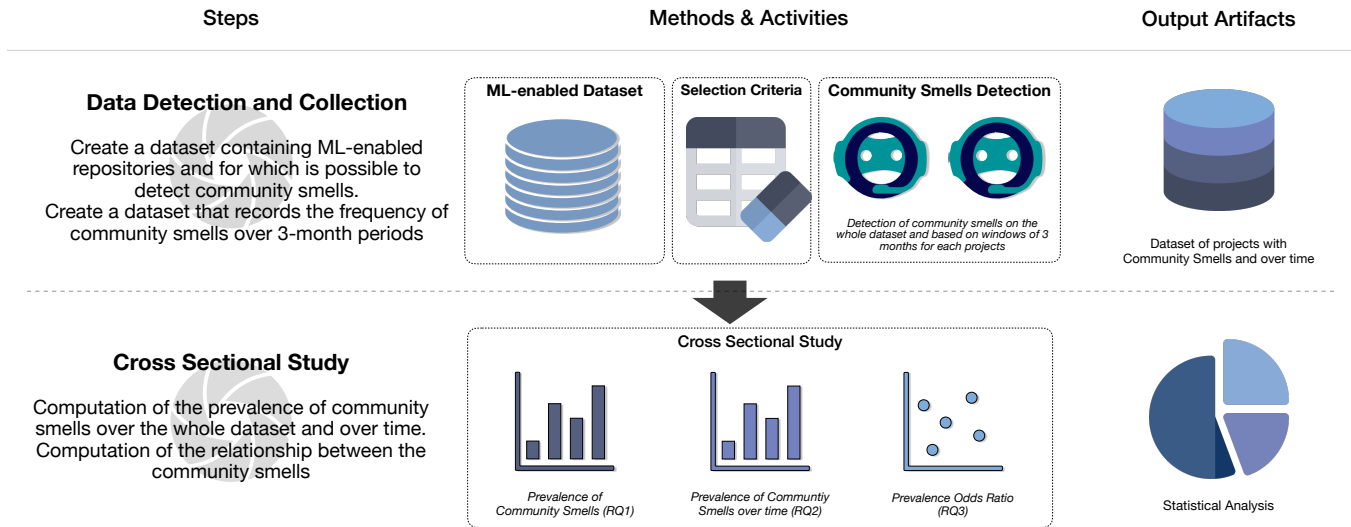


Figure 1: Research Method

a GitHub repository, leveraging a graphical interface and client-server communication. CADOCS, such as CSDETECTOR is able to detect 10 different Community Smells, illustrated in Table 1.

Following the initial detection, we conduct a longitudinal analysis by segmenting the lifespan of each project into eight consecutive 3-month windows, covering a total period of 2 years. Within each time window, we apply CADOCS to generate individual datasets, enabling a detailed temporal examination of community smells and their evolution over time.

The final dataset obtained starts from the validated data from the NICHE dataset [30], filtered according to the selection criteria explained, and expanded with the community smells for the projects that CADOCS [28] was able to detect in December 2024. Through this process, we created our dataset with 188 projects. Detailed reports on this dataset, including the time-windowed datasets (one for each project) are available in our online appendix [1].

3.3 Data Analysis

In this section, we provide a general overview of cross-sectional studies, explaining their current use in the context of software engineering, and then detail how we applied them in our specific study to address our research questions.

3.3.1 Overview on Cross-Sectional Studies. The field of software engineering research has experienced a notable surge in studies leveraging mining software repositories (MSRs), fueled by the growing popularity of online code hosting platforms like GITHUB. This trend has prompted the development of best practices and guidelines to mitigate common pitfalls and enhance the reliability of these studies and their findings [11].

Despite their utility and relative simplicity, MSR studies face a key limitation: they are unable to establish causal relationships for the phenomena they observe. To overcome this shortcoming,

Saarimäki et al. [23] have advocated for the use of observational methods, particularly cohort studies, which provide a more robust foundation for generating high-quality scientific evidence. For this reason, the adoption of a cross-sectional study is a good starting point for addressing this investigation of mining repositories in the context of community smells in ML-enabled systems. Since ML-enabled devices are considered an emerging discipline, the social-technical aspects are still not well explored.

Cross-sectional studies aim to capture and describe the characteristics of a population at a specific point in time. They offer a snapshot of the prevalence of a particular condition or phenomenon and the distribution of relevant exposures within the population [29]. Unlike other observational study designs, cross-sectional studies collect data simultaneously from all participants. The *Prevalence* represents the proportion of individuals in a population who exhibit a specific condition or characteristic at a given moment [29]; in our study, the presence of community smells. Moreover, the *Prevalence Odds Ratio (POR)* is often employed to assess the strength of the association between exposure and a condition. The POR compares the odds of exposure among those with the condition to the odds of exposure among those without it, offering a valuable measure of the relationship in the sample [29].

Cross-sectional studies provide a cost-effective and pragmatic approach for initiating the exploration of socio-technical issues in ML-enabled systems. In this context, such problems are represented by community smells and their interrelationships, given by their heterogeneity, such as they are composed by data scientists and software engineers [5]. Research on the socio-technical dimensions of ML-enabled systems is a relatively new and unexplored field [17]. Cross-sectional studies can deliver critical insights into the current state of this domain, revealing conditions and potential associations. These studies are particularly effective for generating hypotheses

Table 1: Community Smells.

Community Smells	AC	Definition [2]
Organizational Skirmish	OS	Team members with different levels of competence lead to a drop in productivity, impacting time and costs.
Black Cloud Effect	BC	A lack of structured communications or cooperative governance can lead to information overload.
Radio Silence	RS	The use of formal communication between multiple sub-communities penalizes flexibility and causes loss of time.
Prima Donnas Effect	PDE	Team members that expose condescending behavior, superiority, constant disagreement, and uncooperativeness.
Sharing Villainy	SV	Failure to exchange information can lead team members to share incorrect, obsolete, and unconfirmed information.
Organizational Silo Effect	OSE	Siloed community that do not communicate with each other except through one or two members.
Solution Defiance	SD	Having a community with different backgrounds of experience and culture can lead to subgroups with conflicting opinions.
Truck Factor Smell	TFS	Developer turnover may cause a significant loss of knowledge as it is concentrated in a minority of developers.
Unhealthy Interaction	UI	Slow, light, and short conversations and discussions caused by long delays in stakeholder communication.
Toxic Communication	TC	Developers may negatively interact with their colleagues, leading to frustration, stress, and project abandonment.

and identifying possible connections. Although they do have limitations, cross-sectional studies serve as a practical and foundational step in the advancement of research in this area.

3.3.2 Answering RQ₁. To answer to RQ₁, we calculate the *Prevalence* of each detected community smell as the ratio of repositories affected by a specific smell to the total number of repositories analyzed [29]. Formally, it is defined as:

$$P(X) = \frac{\text{Number of repositories affected by smell X}}{\text{Number of all Repository}} \quad (1)$$

where $P(X)$ is the prevalence of the smell X in the dataset.


3.3.3 Answering RQ₂. To answer to RQ₂, we calculate the *Prevalence* for each community smell over time; we analyze the projects in the created dataset using CADOCS with a 3-month temporal granularity, segmenting the data into distinct time windows. These windows span from the first commit of each project up to 2 years of its lifecycle, as explained in Section 3.1. The 3-month granularity was chosen based on socio-technical metrics, which consider this interval sufficient to capture potential turnover and provide an optimal window for studies of this nature [7, 10, 20]. Once we obtained the detection of the smells in the time windows for each project, we calculated the prevalence of each community smells for each time window.

3.3.4 Answering RQ₃. To answer to RQ₃, we calculate the *Prevalence Odds Ratio (POR)*, which measures the association between the presence or absence of a specific community smell and the presence or absence of another community smell [29]. This metric helps evaluate whether certain community smells are more likely to co-occur within ML-enabled projects, providing insights into their interrelationships. It is calculated as:

$$POR = \frac{AD}{BC} \quad (2)$$

To illustrate this with an example, consider the calculation of POR between Organizational Silo Effect (OSE) and Prima Donna Effect (PDE). In this context, A represents the number of cases where both community smells, OSE and PDE, are present. B denotes instances where OSE is present, but PDE is absent. C refers to cases where OSE is absent, but PDE is present. Finally, D represents instances where both OSE and PDE are absent.

Regarding the interpretation of the results, when the value of POR is less than 1, it indicates that there is a negative correlation,

**Figure 2: Prevalence of Community Smells in ML Projects**

in our case, that the presence of the second smell decreases the likelihood of occurrence of the first smell. In the case of a POR of 1 or higher, it indicates a strongly positive association in which the presence of the second smell increases the presence of the first one. For example, a POR value of 2 indicates that the community smell OSE is twice as likely to occur in communities where the community smell PDE is present compared to those where it is absent. Conversely, a POR value of 0.5 suggests that OSE is half as likely to occur in communities with the presence of PDE, implying that PDE may have a mitigating effect.

4 Analysis of Results

This section presents the main findings of our study, detailing the results of the cross-sectional analysis and providing insights into the observed patterns and trends.

4.1 RQ₁—Prevalence of Community Smells

To answer to RQ₁ we calculated the prevalence of each community smell across the 188 GitHub ML-enabled open-source projects of our dataset. Figure 2 shows the distribution of community smell prevalence within the analyzed projects.

Examining the prevalence distribution, we observe that half of the community smells have a prevalence greater than 50%. In particular, Prima Donna Effect (PDE) stands out with a prevalence of 92.6% indicating its presence in nearly all the analyzed projects. The second most prevalent smell is Sharing Villainy (SV), with a

prevalence of 83.5%. Another noteworthy smell is Solution Defiance (SD), which registers a significant 76.1%. Conversely, some community smells exhibit low prevalence, indicating their limited presence within the analyzed projects. For example, Radio Silence (RS) shows a prevalence of 18.6%. Similarly, Organizational Skirmish (OS) appears in only 30.9% of the projects, indicating its rarity compared to the more widespread community smells.

4.2 RQ₂—Prevalence in Time Windows

Figure 3 shows a heatmap where each row represents a community smell, and each column corresponds to the time window analyzed. Each time window represents 3 months and ranges from the beginning of the project to 2 years. Each cell represents the prevalence of a community smell within a specific time window.

From an initial analysis, we note that the heatmap reflects the overall prevalence trends of the projects, where Prima Donna Effect (PDE), Sharing Villainy (SV), and Solution Defiance (SD) remain the most frequent smells, while Organizational Skirmish (OS) and Radio Silence (RS) continue to exhibit low prevalence.

Analyzing them in detail, we noticed interesting results.

PDE emerges as the most persistent smell, with a prevalence ranging between 95% and 96%, higher than the 92.6% reported in the overall project analysis.

However, for other smells the trend is less stable. Sharing Villainy (SV), Solution Defiance (SD), and Unhealthy Interaction (UI) exhibit a steady decrease over time. More in detail, we see that SV started with a prevalence of 88% in the first quarter and dropped to 73% after 2 years, SD decreases from 91% to 72%, and UI shows a decline from 24% to just 6%.

A different discussion can be made for the community smell Organizational Skirmish (OS). It started at 25% in the first quarter, peaked at 32% in the second, and then declined steadily to 19% in the last quarter analyzed.

In contrast, Organizational Silo Effect (OSE), Radio Silence (RS), Truck Factor (TF), and Toxic Communication (TC), exhibit a continuous increase over time: OSE rose from 40% to 58%, RS increased from 22% to 39%, TF climbed from 63% to 85%, and TC grew from 52% to 68%.

A different pattern emerges for Black Cloud Effect (BCE), which started at 57%, decreased in the third quarter to 50%, and then increased to 59%.

4.3 RQ₃—Prevalence Odds Ratio Between Community Smells

Figure 4 presents a heatmap displaying the Prevalence Odds Ratio (POR) values among the various community smells analyzed in the 188 projects. The first notable insight is that most correlations between community smells are negative. However, although inferior, the positive correlations tend to be more intense. An example is the correlation between Prima Donna Effect (PDE) and Organizational Silo Effect (OSE), which stands out with a POR of 4.34. This strong association suggests that it is very probable to find the community smell PDE when the community smell OSE occurs. Similarly, Black Cloud Effect (BC) and Radio Silence (RS) also show a high correlation with OSE, with POR values of 3.44 and 3.19, respectively. Other notable positive correlations, although more moderate, are

between Organizational Skirmish (OS) and OSE, with a POR of 2.24, between Solution Defiance (SD) and BCE, with a POR of 2.12, and between Truck Factor (TC) and Organizational Skirmish (OS) with a POR of 1.97. Other moderate positive correlations are observed between PDE and ECB (POR 1.32), OS and ECB (POR 1.77), OS and PDE (POR 1.69), SD and OSE (POR 1.65), RS and ECB (POR 1.84), RS and OS (POR 1.42), Unhealthy Interaction (UI) and Sharing Villainy (SV) (POR 1.20), UI and Solution Defiance (SD) (POR 1.33), UI and TF (POR 1.38), Toxic Communication (TC) and BCE (POR 1.74), TC and SV (POR 1.80), and TC and RS (POR 1.74).

Correlations with POR values below 1 indicate a negative relationship. For instance, UI and OS show a strong negative correlation (POR 0.39), meaning that the presence of UI reduces the likelihood of OS occurring. Other high negative correlations are between OS and SV smells (POR 0.40), IU and OSE (POR 0.40), SV and ECB (POR 0.41), IU and ECB (POR 0.45), TF and PDE (POR 0.46), TC and OS (POR 0.43), IU and RS (POR 0.47), and TC and PDE (POR 0.45).

5 Discussions and Implications

In this section, we discuss the results of the study and focus on their implications, both for researchers and practitioners.

5.1 Prevalence of Community Smells

The results highlight the common occurrence of community smells in ML-enabled projects, with certain types being especially pervasive. Prima Donna Effect (PDE) stands out with a prevalence of 92.6%, indicating that dominant, non-collaborative behaviors among individual developers are a persistent issue in the analyzed projects. This suggests a critical socio-technical challenge, where influential team members impose their decisions, potentially stifling collaboration and innovation. The high prevalence of this behavior in ML-enabled projects might suggest that it may often be adopted by Data Scientists [14, 18]. Prior studies [4] associated PDE with the tendency of data scientists to assert expertise, which can lead to conflicts within team members and frustration among project managers, who struggle to manage these dynamics effectively [4, 14]. Sharing Villainy (SV), with a prevalence of 83.5%, highlights the widespread lack of adequate information sharing. This often results in the spread of incorrect or outdated knowledge, which can negatively impact project quality and decision-making. Similarly, Solution Defiance (SD), occurring in 76.1% of projects, indicates frequent subgroups conflicts, resulting in decision-making bottlenecks [5] and resistance to established best practices.

The high prevalence of PDE, SV, and SD reflects the inherent challenge of aligning the workflows and goals of data scientists and software engineers. This misalignment fosters power imbalances and conflicting goals, leading to community breakdown.

Conversely, lower prevalence rates of Radio Silence (RS) (18.6%) indicate that explicit breakdowns in communication or role-based conflicts are less pervasive but still significant. Organizational Skirmish (OS) (30.9%) also presents a low prevalence. In the ML-enabled context, this community smell can arise when data scientists and software engineers lack the skills to adequately understand one another. Their distinct backgrounds can hinder effective collaboration and impact project outcomes [14, 18]. However, the relatively lower occurrence of OS and RS suggests that differences in knowledge and




Figure 3: Prevalence of Community Smells in ML Projects in Windows of 3 Months

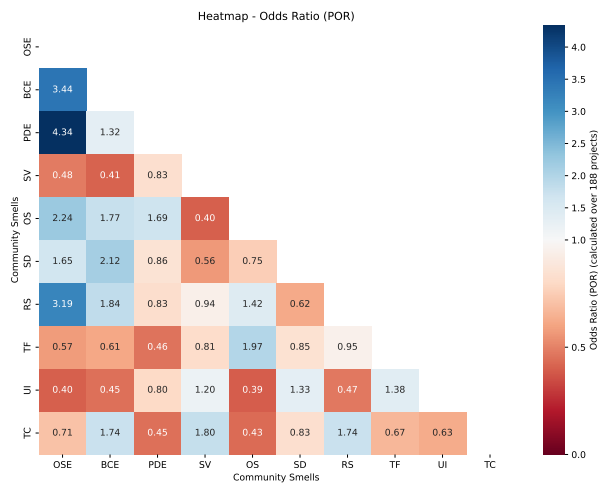


Figure 4: Heat map of Prevalence Odds Ratio

expertise may not directly contribute to social issues or technical challenges within the project [4]. The relatively lower occurrence of these smells may reflect the presence of structured workflows, although informal knowledge sharing gaps persist.

Implications for research: Further investigation is needed to understand how the distinct goals of data scientists and software engineers contribute to the emergence of community smells. Developing socio-technical frameworks tailored to ML teams could help mitigate the prevalence of PDE and SV.

Implications for practice: Encouraging structured collaboration between data scientists and software engineers can help reduce community smells. For example, rotating responsibilities between data scientists and software engineers or adding intermediate figures between the two roles could reduce the occurrence of such smells in teams [14].

5.2 Evolution of Community Smells Over Time

The longitudinal analysis reveals nuanced trends in the evolution of community smells. While PDE maintains a consistently high

prevalence over time (between 95% and 96%), indicating persistent dominance behaviors, SV and SD show gradual declines. This result indicates a consistent pattern among one or more team members exhibiting superiority and limited collaboration with the rest of the team. Literature suggests that such behaviors are often attributed to data scientists, who tend to prioritize their individual goals over the project objectives [4, 14]. A persistent and significant presence of these smells highlights the recurrence of such behavior throughout the software development process [4].

SV drops from 88% to 73% over 2 years, while SD decreases from 91% to 72%. This suggests that initial conflicts and isolated knowledge practices may reduce as teams establish shared lines of conduct and governance structures [14]. However, smells such as Organizational Silo Effect (OSE) and Toxic Communication (TC) increase over time, with OSE rising from 40% to 58% and TC from 52% to 68%. These results suggest that when projects become larger, it may be the case that split subgroups between data scientists and software engineers come into existence; these subgroups may operate independently, leading to fragmented communication and may emerge a separation of roles and responsibilities [4]. These problems often result in less efficient collaboration and hinder the overall success of software development projects.

Implications for research: Researchers can study strategies for managing the changing dynamics of heterogeneous teams in the ML-enabled context. Conducting more in-depth longitudinal studies and considering causal factors of such community smells, could provide a more complete and detailed picture of the presence and evolution of such smells in ML-enabled development teams.

Implications for practice: The findings suggest an increase in OSE and TC smells over time; to counteract this, practitioners could adopt regular team reviews and promote open communication channels.

5.3 Correlation Between Community Smells

The cross-sectional analysis identifies different positive correlations between community smells. Prima Donna Effect (PDE) correlates strongly with Organizational Silo Effect (OSE) (Prevalence Odds Ratio, POR = 4.34), suggesting that teams dominated by a lead data scientist or software engineer are more prone to isolation and lack of cross-group communication. This result suggests a strong coexistence of the two community smells. It might be reasonable to think that the adoption of Prima Donna Effect behavior, adopted mostly by data scientists when they aim for their individual goals, tends to estrange other team members, causing the team to “break” into uncooperative sub-teams and isolated workflows driven by data scientists prioritize experimentation or software engineers who emphasize delivery timelines [4].

Similarly, the Black Cloud Effect (BCE) correlates with Radio Silence (RS) (POR = 3.19), highlighting scenarios where insufficient governance leads to information bottlenecks. Negative correlations, such as between Unhealthy Interaction (UI) and Organizational Skirmish (OS) (POR = 0.39), suggest that while direct conflict may occur, it can deter large-scale division by forcing collaboration. BCE and RS often result from gaps in communication channels,

where the absence of structured leadership contributes to limited information sharing.

Implications for research: More research is needed to uncover more details about the rationale behind the behaviors of data scientists and software engineers in determining the smells of the correlated community. Understanding such a rationale may lay the basis for further work related to correlation or association rules between different community smells occurring in the ML-enabled context.

Implications for practice: Professionals should apply strategies to mitigate social issues; e.g., apply pair programming or introduce team members who can communicate and actively collaborate with both data scientists and software engineers.

5.4 Differences in Prevalence of Community Smells between RQ1 and RQ2

Interestingly, a discrepancy emerges when comparing the prevalence of some community smells between the cross-sectional and longitudinal analyses, i.e., the prevalence of the whole projects and the prevalence in windows of 3 months from birth to 2 years of life for each project. Prima Donna Effect (PDE) reported at 92.6% in the initial analysis (RQ1), increasing to 95-96% in the longitudinal study (RQ2). This slight increase may indicate that the dominance of some developers, often Data Scientists or Principal Software Engineers, intensifies as projects evolve and key collaborators consolidate their roles. In contrast, Unhealthy Interaction (UI) shows a reverse trend. Initially, UI showed a prevalence of 52% in RQ1, but over time, it fluctuated between 24% and 6% in the longitudinal analysis. This decline might suggest that, while communication problems are evident early in the project life cycle, teams gradually implement mechanisms to reduce their frequency, such as regular stand-ups or communication protocols. This discrepancy in prevalence poses an excellent baseline question to entice longitudinal studies, with the goal of studying in detail the evolution and motivations behind community smells in the ML-enabled context.

» Take Away Results

- The constant and high prevalence of the smell Prima Donna Effect suggests misaligned workflows between data scientists and software engineers, which hinders team collaboration and productivity.
- Community smells are not static, they can decrease (e.g., SV, SD, and UI), or increase (e.g., OSE, RS, TF, and TC).
- Most correlations between community smells are negative, but there are cases of strong positive correlation, such as that between PDE and OSE.

6 Threats To Validity

This section discusses how we prevent and mitigate the threats to the validity of the study, according to Wohlin et al. [31].

Threats to Construct Validity. Construct validity concerns arise when there is a mismatch between the conceptual framework and the measurements used [31]. We used the NICHE dataset [30],

which has been previously validated, to mitigate this risk. Additionally, we relied on CADOCS [28] to compute community smells. Although CADOCS focuses on a subset of community smells, potentially overlooking certain socio-technical challenges in ML-enabled developer communities, it offers valuable insights into the influence of these smells on software development.

Threats to Conclusion Validity. Conclusion validity pertains to the reliability of the relationships identified between treatments and outcomes [31]. A potential threat is the possibility of statistical errors, such as drawing incorrect conclusions due to insufficient sample size or variability in the dataset. To address this, we ensured the dataset was balanced and representative by applying specific selecting criteria established in the literature for similar studies [7, 26], to create our dataset to conduct the study. This approach enhances the reliability of our findings. Future research could further reduce this risk by replicating the analysis on larger, more diverse datasets and experimenting with alternative statistical methods to validate and extend the study’s conclusions. Another potential threat to the validity of the conclusions concerns the measures used to evaluate our research questions. We adopted a cross-sectional study and calculated prevalence and prevalence odds ratios to quantify the occurrence of community smells and their relationships. Although these measures are well established in similar studies, potential limitations remain. For example, the cross-sectional nature of RQ_1 and RQ_3 may not capture causal dynamics, and the choice of 3-month time windows in RQ_2 may affect trend observations. Future research could explore alternative statistical techniques, such as causal inference methods, to further validate and extend our results.

Threats to External Validity. External validity concerns the generalizability of the results [31]. To address this, we built a comprehensive dataset of ML-enabled projects using the validated NICHE dataset [30] as a starting point. Our study analyzed 188 projects, employing selection criteria aligned with prior research to enhance relevance and reliability. In addition, projects were examined in 3-month intervals, spanning from inception to 2 years of activity, ensuring longitudinal insights. While this approach supports satisfactory generalizability, future work will expand the sample size and diversity to further validate and generalize the findings.

Threats to Internal Validity. Threats to internal validity involve factors that could have influenced the study’s results [31]. A potential concern is selection bias from the chosen of NICHE dataset [30], which might limit the generalizability of the findings to all ML-enabled projects. To address this, we employed established selection criteria from the literature to enhance the dataset’s representativeness [7, 26]. Another potential issue is measurement bias introduced by the CADOCS tool [28], which may not capture all community smells, but only a subset of 10 community smells.

7 Conclusion

This study investigated human factors and behavioral aspects of software engineers and data scientists, analyzing the prevalence, evolution, and correlation of community smells in ML-enabled projects, and emphasizing social interactions between data scientists and software engineers. The study aimed to provide key insights into the challenges of heterogeneous teams developing

ML-enabled systems. The study analyzes the prevalence and evolution over time of social behaviors that impact the technical aspects of software, such as dominance behaviors, knowledge silos, and communication breakdowns. The key contributions of the study are:

- *Prevalence analysis:* We established the overall prevalence of community smells, and we see that PDE, SV, and SD are the most dominant in ML-enabled projects.
- *Longitudinal analysis:* We analyzed the evolution of community smells over time, revealing the constant persistence of some smells, such as PDE, and the decline of others, such as UI.
- *Correlation analysis:* We identified correlations between community smells, e.g., PDE and OSE, and provided insights into how these smells interact and reinforce each other.

The results underscore the socio-technical complexity inherent in ML-enabled projects. The high prevalence of community smells PDE underscores the need to address power imbalances within teams. Heterogeneous teams composed of data scientists and software engineers should prioritize collaborative workflows and integrated development practices to mitigate these smells.

Future work should expand the dataset to analyze a wider range of ML-enabled projects. Conducting qualitative studies, such as interviews or case studies, can provide insights into the root causes of community smells and how they manifest in different contexts. Design of socio-technical frameworks that foster better collaboration, improve communication pipelines and bridge the gap between data scientists and software engineers in ML-based environments. Finally, longitudinal studies that follow the entire lifecycle of ML projects could shed light on the evolution of socio-technical issues at different stages, offering targeted solutions at critical moments.

Data Availability

The dataset created, the material and the script used for running the analysis are available in the online appendix [1].

Acknowledgment

This work has been partially supported by the *EMELIOT* national research project, which has been funded by the MUR under the PRIN 2020 program (Contract 2020W3A5FY), and *QUAL-AI* national research projects funded by the EU - NGEU and the MUR under the PRIN 2022 program (Contracts 2022B3BP5S).

References

- [1] 2025. *Online Appendix - How Do Communities of ML-Enabled Systems Smell? A Cross-sectional Study on the Prevalence of Community Smells*. <https://figshare.com/s/64909f10a362ba9b3318>
- [2] Nuri Almarimi, Ali Ouni, Moataz Chouchen, and Mohamed Wiem Mkaouer. 2021. csDetector: an open source tool for community smells detection. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1560–1564.
- [3] Giusy Annunziata, Carmine Ferrara, Stefano Lambiase, Fabio Palomba, Gemma Catolino, Filomena Ferrucci, and Andrea De Lucia. 2024. An Empirical Study on the Relation Between Programming Languages and the Emergence of Community Smells. In *2024 50th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 268–275.
- [4] Giusy Annunziata, Stefano Lambiase, Damian A. Tamburri, Willem-Jan van den Heuvel, Fabio Palomba, Gemma Catolino, Filomena Ferrucci, and Andrea De Lucia. 2025. Uncovering Community Smells in Machine Learning-Enabled Systems:

- Causes, Effects, and Mitigation Strategies. *ACM Transactions on Software Engineering and Methodology* (Jan. 2025). <https://doi.org/10.1145/3712198> Just Accepted.
- [5] Gabriel Busquim, Hugo Villamizar, Maria Julia Lima, and Marcos Kalinowski. 2024. On the Interaction Between Software Engineers and Data Scientists When Building Machine Learning-Enabled Systems. In *Software Quality as a Foundation for Security*, Peter Bludau, Rudolf Ramler, Dietmar Winkler, and Johannes Bergsmann (Eds.). Springer Nature Switzerland, Cham, 55–75.
 - [6] Marcelo Cataldo, Patrick A Wagstrom, James D Herbsleb, and Kathleen M Carley. 2006. Identification of coordination requirements: Implications for the design of collaboration and awareness tools. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. 353–362.
 - [7] Gemma Catolino, Fabio Palomba, Damian Andrew Tamburri, and Alexander Serebrenik. 2021. Understanding community smells variability: A statistical approach. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE, 77–86.
 - [8] Gemma Catolino, Fabio Palomba, Damian A. Tamburri, Alexander Serebrenik, and Filomena Ferrucci. 2019. 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)*. 11–20. <https://doi.org/10.1109/ICSE-SEIS.2019.00010>
 - [9] Harald Foidl, Michael Felderer, and Rudolf Ramler. 2022. Data smells: Categories, causes and consequences, and detection of suspicious data in ai-based systems. In *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*. 229–239.
 - [10] Zijie Huang, Zhiqing Shao, Guisheng Fan, Jianhua Gao, Ziyi Zhou, Kang Yang, and Xinguang Yang. 2021. Predicting community smells' occurrence on individual developers by sentiments. In *2021 IEEE/ACM 29th International Conference on Program Comprehension (ICPC)*. IEEE, 230–241.
 - [11] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M German, and Daniela Damian. 2014. The promises and perils of mining github. In *Proceedings of the 11th working conference on mining software repositories*. 92–101.
 - [12] Stefano Lambiase, Gemma Catolino, Damian A. Tamburri, Alexander Serebrenik, Fabio Palomba, and Filomena Ferrucci. 2022. 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 '22)*. Association for Computing Machinery, New York, NY, USA, 67–78. <https://doi.org/10.1145/3510458.3513015>
 - [13] Zengyang Li, Paris Avgeriou, and Peng Liang. 2015. A systematic mapping study on technical debt and its management. *Journal of Systems and Software* 101 (2015), 193–220.
 - [14] Alina Mailach and Norbert Siegmund. 2023. Socio-Technical Anti-Patterns in Building ML-Enabled Software: Insights from Leaders on the Forefront. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. 690–702. <https://doi.org/10.1109/ICSE48619.2023.00067>
 - [15] Silverio Martínez-Fernández, Justus Bogner, Xavier Franch, Marc Oriol, Julien Siebert, Adam Trendowicz, Anna Maria Vollmer, and Stefan Wagner. 2022. Software Engineering for AI-Based Systems: A Survey. *ACM Trans. Softw. Eng. Methodol.* 31, 2, Article 37e (April 2022), 59 pages. <https://doi.org/10.1145/3487043>
 - [16] Antonio Martini and Jan Bosch. 2017. Revealing social debt with the CAFFEA framework: An antidote to architectural debt. In *2017 IEEE International Conference on Software Architecture Workshops (ICSAW)*. IEEE, 179–181.
 - [17] Nadia Nahar, Shurui Zhou, Grace Lewis, and Christian Kästner. 2022. Collaboration Challenges in Building ML-Enabled Systems: Communication, Documentation, Engineering, and Process (ICSE '22). Association for Computing Machinery, New York, NY, USA, 413–425. <https://doi.org/10.1145/3510003.3510209>
 - [18] Ipek Ozkaya. 2020. What is really different in engineering AI-enabled systems? *IEEE software* 37, 4 (2020), 3–6.
 - [19] Fabio Palomba, Damian Andrew Tamburri, Francesca Arcelli Fontana, Rocco Oliveto, Andy Zaidman, and Alexander Serebrenik. 2021. Beyond Technical Aspects: How Do Community Smells Influence the Intensity of Code Smells? *IEEE Transactions on Software Engineering* 47, 1 (Jan. 2021), 108–129. <https://doi.org/10.1109/TSE.2018.2883603>
 - [20] Fabio Palomba and Damian Andrew Tamburri. 2021. Predicting the Emergence of Community Smells Using Socio-Technical Metrics: A Machine-Learning Approach. *Journal of Systems and Software* 171 (Jan. 2021), 110847. <https://doi.org/10.1016/j.jss.2020.110847>
 - [21] Paul Ralph, Nauman bin Ali, Sebastian Baltes, Domenico Bianculli, Jessica Diaz, Yvonne Dittrich, Neil Ernst, Michael Felderer, Robert Feldt, Antonio Filieri, et al. 2020. Empirical standards for software engineering research. *arXiv preprint arXiv:2010.03525* (2020).
 - [22] Gilberto Recupito, Raimondo Rapacciuolo, Dario Di Nucci, and Fabio Palomba. 2024. Unmasking Data Secrets: An Empirical Investigation into Data Smells and Their Impact on Data Quality. In *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering-Software Engineering for AI*. 53–63.
 - [23] Nytyi Saarimäki, Valentina Lenarduzzi, Sira Vegas, Natalia Juristo, and Davide Taibi. 2020. Cohort studies in software engineering: A vision of the future. In *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. 1–6.
 - [24] Damian Andrew Tamburri, Rick Kazman, and Hamed Fahimi. 2016. The architect's role in community shepherding. *IEEE Software* 33, 6 (2016), 70–79.
 - [25] Damian A Tamburri, Philippe Kruchten, Patricia Lago, and Hans van Vliet. 2015. Social debt in software engineering: insights from industry. *Journal of Internet Services and Applications* 6 (2015), 1–17.
 - [26] Damian A. Tamburri, Fabio Palomba, and Rick Kazman. 2021. Exploring Community Smells in Open-Source: An Automated Approach. *IEEE Transactions on Software Engineering* 47, 3 (2021), 630–652. <https://doi.org/10.1109/TSE.2019.2901490>
 - [27] Hugo Villamizar, Marcos Kalinowski, and Hélio Lopes. 2022. Towards Perspective-Based Specification of Machine Learning-Enabled Systems. In *2022 48th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. 112–115. <https://doi.org/10.1109/SEAA56994.2022.00025>
 - [28] Gianmario Voria, Viviana Pentangelo, Antonio Della Porta, Stefano Lambiase, Gemma Catolino, Fabio Palomba, and Filomena Ferrucci. 2022. Community Smell Detection and Refactoring in SLACK: The CADOCs Project. In *2022 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 469–473.
 - [29] Xiaofeng Wang and Zhenshun Cheng. 2020. Cross-sectional studies: strengths, weaknesses, and recommendations. *Chest* 158, 1 (2020), S65–S71.
 - [30] Ratnadira Widayarsi, Zhou Yang, Ferdian Thung, Sheng Qin Sim, Fiona Wee, Camellia Lok, Jack Phan, Haodi Qi, Constance Tan, Qijin Tay, et al. 2023. Niche: A curated dataset of engineered machine learning projects in python. In *2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR)*. IEEE, 62–66.
 - [31] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in software engineering*. Springer Science & Business Media.
 - [32] Haiyin Zhang, Luis Cruz, and Arie Van Deursen. 2022. Code smells for machine learning applications. In *Proceedings of the 1st international conference on AI engineering: software engineering for AI*. 217–228.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009