

From Principle to Practice: Evaluating Inclusivity in Free and Open Source Software

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Tony Chang
Spring 2025

On my honor as a University Student, I have neither given nor received unauthorized aid on this
assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor
Kent Wayland, Department of Engineering and Society

Introduction

Free and Open Source Software (FOSS) has become a cornerstone of digital infrastructure, shaping everything from government systems and scientific research to the tools used in everyday programming. It empowers developers to freely access, modify, and share software, and has historically embodied values of openness, collaboration, and collective innovation. However, despite its decentralized and democratic ideals, participation in FOSS communities remains unequal. Marginalized groups—including individuals with disabilities, developers from non-Western countries, and people from low-income or underrepresented backgrounds—face persistent barriers to entry and retention. These barriers range from inaccessible documentation and exclusionary governance structures to unpaid labor expectations and language constraints.

Understanding how these structural challenges affect the participation of underrepresented groups is essential for shaping a more inclusive future for FOSS. This paper analyzes these dynamics by examining academic literature, empirical data from large repositories such as GitHub and GNOME, and case studies highlighting the contributions and struggles of marginalized developers. Drawing from Science, Technology, and Society (STS) frameworks, the analysis explores how power, culture, and governance operate within FOSS communities and how inclusive practices can enhance both innovation and social equity. The findings emphasize the urgent need for reforms that go beyond surface-level diversity efforts and instead address the deeper structural inequalities embedded in FOSS ecosystems.

Background & Context

While the open-source movement aspires to democratize technology, it often falls short of this goal, particularly for marginalized groups such as people with disabilities and individuals from diverse cultural backgrounds. FOSS development has historically been dominated by a specific demographic, primarily able-bodied individuals from English-speaking, resource-rich countries. This dominance has shaped the culture and practices within FOSS communities, often excluding those who do not fit the prevailing norm. For example, in the early days of Linux kernel development, the majority of contributors were male programmers from North America and Europe, leading to a culture that often alienated those outside this group. This lack of diversity stems from both technical and social barriers that hinder full participation by underrepresented groups, such as inaccessible development tools and language barriers. The technical challenges include the design and accessibility of tools, which were not originally built with the needs of disabled developers in mind. For example, visually impaired developers often face significant challenges in using text-based interfaces and software tools that do not support screen readers or other assistive technologies.

Language barriers further complicate the inclusivity of FOSS communities, as the majority of development and communication occurs in English. Non-native speakers struggle to contribute effectively to code or discussions that are conducted in English-only forums, mailing lists, and chat channels. Additionally, technical jargon and complex documentation often exclude

contributors who are not familiar with the specialized language of FOSS. This creates an environment that inadvertently discourages participation from those without access to advanced language skills or educational resources. Moreover, the issue of unpaid labor within the FOSS ecosystem disproportionately affects marginalized groups, particularly those from lower-income backgrounds. Since FOSS contributions are typically volunteer-based, individuals with limited financial flexibility are often excluded from participation, as they must prioritize paid work and other financial responsibilities over unpaid contributions.

Despite these systemic barriers, there have been ongoing efforts to address these disparities. Initiatives like Open Source Diversity and accessibility-focused development frameworks have begun to pave the way for broader participation in open-source projects. These initiatives recognize that inclusivity in FOSS is not just about providing open access to code but also about creating an environment that actively supports marginalized contributors. Research has shown that diversity in software development is a powerful driver of innovation, leading to more creative solutions to complex problems. Projects like GNOME's engagement with disabled developers have demonstrated how inclusivity can lead to significant improvements in accessibility features that benefit all users. However, while these efforts are important, structural and cultural challenges persist, and further work is needed to ensure FOSS aligns with its ideals of openness and inclusivity for all.

Literature

A key theme emerging from the literature is the role of diversity in software development, particularly within the FOSS ecosystem. Research consistently suggests that knowledge diversity among developers enhances innovation and inclusivity, leading to more creative solutions and a broader range of ideas. A study by Bogers, Foss, and Lyngsie (2018) emphasized that diverse knowledge bases within teams improve an organization's ability to absorb external knowledge and adapt to new challenges, fostering openness to collaboration. This is particularly relevant to FOSS, where contributions from developers with varied backgrounds, including those from underrepresented communities, can drive accessibility improvements, as seen in the GNOME project's engagement with disabled developers (Castro, 2003). However, ensuring that this diversity translates into meaningful participation remains a challenge, as marginalized groups continue to face multiple barriers to entry.

One significant barrier is the accessibility of development tools for people with disabilities. While the potential for software engineering to serve as a pathway for social inclusion is evident, research indicates that disabled developers face substantial technical and social challenges. A notable study by Thayssa A. da Rocha et al. (2024) highlighted the difficulties faced by developers with visual impairments, such as the lack of screen reader-compatible development environments and inaccessible collaboration tools. These challenges often deter individuals from pursuing careers or hobbies in software development despite the growing awareness of the need for accessible technologies. To address these obstacles, some researchers have proposed frameworks such as the "User challenges – heatmap"

(Kivimäk, 2023), which aims to identify the most pressing accessibility issues and guide more inclusive software engineering practices. However, the widespread adoption of such frameworks across the broader FOSS ecosystem remains slow, and there is still much work to be done to ensure that disabled developers can participate fully and effectively.

The impact of diversity on team dynamics and project outcomes is another critical area of focus in the literature. Research by Liang et al. (2007) examined how diversity can influence both task and relationship dynamics within teams. They found that while knowledge diversity can foster constructive task conflict—leading to better problem-solving—value diversity may contribute to relationship conflict, which can undermine collaboration and project success. This is particularly relevant within FOSS, where projects are often decentralized and voluntary, with contributors working asynchronously across the globe. In such environments, social tensions arising from differing cultural values, work ethics, and communication styles can create friction that hampers collaboration. Studies have shown that, despite the ideal of open collaboration, these tensions are often heightened in FOSS communities, making it difficult for diverse teams to work together harmoniously. Understanding how to manage these conflicts and foster productive relationships is key to creating more inclusive and effective FOSS projects (Ting-Peng Liang et al., 2007).

Furthermore, systemic barriers such as gender bias, socioeconomic factors, and geopolitical constraints continue to shape who can participate in FOSS development. The literature reveals significant biases in hiring and differential acceptance of code contributions based on the developer's identity (Albusays et al., 2021). Gender bias is particularly pronounced in FOSS, where studies have shown that women are underrepresented in both contributor and leadership roles. A report by Vasilescu et al. (2015) found that 75.3% of GitHub projects had no gender diversity, highlighting the structural nature of this exclusion. Similarly, developers from lower-income backgrounds often face significant challenges in contributing to FOSS, as the unpaid nature of many contributions creates barriers for those without the financial flexibility to participate. This is compounded by the fact that many FOSS projects are dominated by contributors from resource-rich countries, particularly English-speaking regions, further disadvantaging individuals from underrepresented geographies (Schwartz et al., 2024). These findings underscore the need for intentional interventions and inclusive policies that actively support the participation of marginalized groups.

Lastly, the governance structures of FOSS projects have been the subject of considerable scrutiny in the literature. Research has shown that many FOSS communities lack formal governance frameworks, relying instead on informal norms and community-based decision-making processes. While this approach has its advantages, such as flexibility and responsiveness, it can also create barriers for newcomers, particularly those from marginalized groups who may not be familiar with the unspoken rules of the community. Coleman (2013) explored this phenomenon and argued that the informality of governance structures in FOSS can lead to an exclusionary culture, where those who are not well-versed in the norms of the community may struggle to navigate decision-making processes. This lack of formalized

structure can also exacerbate issues related to power dynamics, with a small group of highly active contributors holding disproportionate influence over the direction of a project. Rethinking governance structures to be more inclusive and transparent is therefore crucial for fostering greater participation from underrepresented groups in FOSS.

Methods

This project relied on qualitative and quantitative research methods to examine the role of underrepresented minority groups in FOSS and the challenges they face. The evidence included case studies of FOSS projects, such as the GNOME accessibility toolkit repository, to assess inclusivity efforts. Additionally, existing literature and reports on knowledge diversity, accessibility in software development, and participation barriers in FOSS provided context. Data on code contributions and community participation were also reviewed, with a focus on acceptance rates, documentation accessibility, and governance structures.

The research took place over three weeks. In the first week, a comprehensive literature review was conducted to refine the research focus and identify relevant case studies. During the second week, I gathered data from FOSS repositories, including information on contributor demographics, contribution frequency, pull request acceptance rates, and documentation accessibility. The third week was dedicated to analyzing the collected data. Qualitative data from case studies and literature underwent thematic analysis to identify key challenges and best practices for inclusivity in FOSS. Quantitative data from repositories were analyzed using descriptive statistics to assess trends in participation, accessibility, and sustainability. A comparative analysis was performed to evaluate whether FOSS communities' stated inclusivity goals aligned with actual practices. This structured research process provided a clear assessment of the barriers and opportunities for fostering greater participation from underrepresented groups in FOSS.

Results

The findings indicate that while Free and Open Source Software (FOSS) communities are built on principles of openness and inclusivity, significant structural barriers persist, limiting the participation of marginalized groups. The case studies revealed that knowledge diversity enhances innovation, as contributions from individuals with varied backgrounds foster new approaches to problem-solving and lead to the development of more accessible technologies. For instance, the GNOME project's engagement with disabled developers highlighted how the inclusion of diverse perspectives led to improvements in the accessibility toolkit, making the software more usable for individuals with visual impairments. However, these benefits are not always realized in practice, as barriers such as inaccessible development tools, language constraints, and cultural norms still prevent marginalized groups from participating fully.

Repository data revealed clear disparities in contribution acceptance rates, with underrepresented groups, particularly women and developers from non-Western countries, experiencing higher rejection rates for their pull requests compared to their male and Western

counterparts. In some cases, these developers also faced difficulties in getting their contributions reviewed and merged into projects, leading to feelings of exclusion and discouragement. This is consistent with previous findings by Albusays et al. (2021), which documented biases in the acceptance of contributions based on the developer's gender, location, and identity. Additionally, the data showed that documentation accessibility was inconsistent across repositories, often lacking clear guidelines for contributors with disabilities or non-native English speakers. For example, while some projects offered translated documentation or screen reader-friendly formats, many others did not, creating unnecessary barriers for non-English speaking and disabled contributors.

Governance structures within FOSS communities also play a significant role in shaping participation. The data revealed that many FOSS projects had informal governance structures that favored contributors familiar with the unspoken norms of the community, which often aligned with the practices of developers from resource-rich, English-speaking countries. These informal structures often created power imbalances, as the most active contributors, typically from these privileged groups, had more influence in decision-making processes. The findings also indicated that these informal structures led to a concentration of power among a small group of contributors, exacerbating the "bus factor" problem, where the sustainability of a project becomes highly dependent on a few key individuals. This concentration of power further marginalized newcomers and those from less-represented backgrounds who were unfamiliar with the informal community practices. This aligns with previous research by Smith et al. (2001), which found that core activities within FOSS projects were concentrated among a small group of active contributors, leading to uneven workload distribution and contributing to sustainability risks.

In terms of team dynamics, the findings also reflected the complex relationship between diversity and project outcomes. While the inclusion of diverse team members can enhance productivity and innovation, it can also introduce social tensions that affect collaboration. For example, the data revealed that cultural differences and varying expectations around communication and work styles created challenges for some FOSS projects, particularly those with contributors from different parts of the world. This was especially true in decentralized projects, where contributors often work asynchronously and may not have opportunities for face-to-face interaction. These tensions, while often productive, could also lead to frustration and disengagement, particularly among those from marginalized groups who might not feel that their perspectives are fully appreciated or understood by the rest of the team. The findings suggest that creating a more inclusive and collaborative environment will require fostering a culture of respect and understanding, as well as implementing formal structures to support cross-cultural communication.

Analysis

The results indicate that while FOSS communities have made some progress in promoting diversity and inclusivity, these efforts have not yet been sufficient to overcome the

entrenched barriers that limit the full participation of marginalized groups. One key finding is the significant disparities in contribution acceptance rates for underrepresented developers, suggesting that the review process in many FOSS projects may be biased. These biases can result in the exclusion of contributions from individuals who are already disadvantaged by accessibility and language challenges. The higher rejection rates faced by marginalized contributors highlight the need for more inclusive review practices that ensure all developers are evaluated based on the quality of their contributions rather than unconscious biases (Schwartz et al., 2024).

The inconsistent accessibility of documentation further exacerbates these challenges. While some projects have made strides toward improving accessibility for developers with disabilities, many free and open-source software (FOSS) projects still lack basic accommodations, such as screen reader-friendly documentation or guidelines on how to contribute if one has a disability. This inconsistency underscores the importance of establishing universal accessibility standards within FOSS development. Without these standards, marginalized groups, particularly those with disabilities, will continue to face significant barriers to contributing effectively. For instance, Thayssa A. da Rocha et al. (2024) detailed how visually impaired developers encounter difficulty using common development environments and communication tools, further marginalizing them despite the potential of software engineering to offer accessible career pathways.

Furthermore, the informal governance structures that dominate many FOSS communities often favor established contributors who are already familiar with the community's unspoken norms. This presents a significant barrier to newcomers, particularly those from underrepresented backgrounds who may be unfamiliar with these implicit rules. As Coleman (2013) argued, such informal governance systems can lead to the development of exclusionary cultures, where decision-making is opaque and power is concentrated among a small group of core contributors. These imbalances discourage broader participation and hinder the inclusion of diverse perspectives in important project decisions. The persistence of such informal governance practices may also reinforce socioeconomic disparities, as only those with the time, resources, and cultural familiarity can navigate these informal networks effectively.

The "bus factor" problem revealed in the data highlights the broader implications of these exclusionary structures: sustainability risks. The reliance on a small group of core contributors makes projects vulnerable to loss of momentum or expertise, especially if these contributors leave or burn out. This concentration of work, as observed in GNOME and documented by Smith et al. (2001), not only creates bottlenecks in project management but also perpetuates power imbalances that discourage new contributors from engaging deeply with the project. Addressing this issue requires proactive redistribution of responsibilities, better onboarding procedures, and mentorship programs aimed at empowering underrepresented contributors to take on leadership roles.

In terms of the impact of diversity on team dynamics, the literature shows that the presence of varied knowledge bases can lead to more innovative and adaptive development outcomes. Research by Bogers, Foss, and Lyngsie (2018) highlights how knowledge diversity

enhances an organization's capacity to integrate external ideas and respond to challenges. This aligns with the GNOME project's experience, where collaboration with developers with disabilities resulted in practical improvements to accessibility features. However, this potential is often undercut by relationship conflicts and social tensions that arise from value diversity. As Liang et al. (2007) noted, while task-related diversity can drive innovation, differences in cultural values and communication norms can lead to interpersonal conflict, particularly in decentralized and asynchronous open-source software (FOSS) environments. These tensions can be particularly damaging when they discourage continued participation by contributors who already face systemic disadvantages.

Additionally, structural inequalities rooted in gender, geography, and class continue to shape participation in FOSS. The Vasilescu et al. (2015) study revealed that a significant majority of GitHub projects lacked gender diversity, and that the retention and participation of women in open-source communities remain low. Similarly, Albusays et al. (2021) documented how developers from non-Western countries and lower-income backgrounds are more likely to have their contributions rejected and less likely to advance to leadership positions within projects. These disparities suggest that diversity in FOSS is not just about recruitment but about retention and empowerment, ensuring that once individuals join a community, they are supported, respected, and included in meaningful ways.

Taken together, these findings suggest that achieving true inclusivity in FOSS requires structural interventions rather than simply aspirational commitments. Improvements are needed in contribution evaluation, documentation accessibility, and governance transparency. Additionally, tools like the "User challenges – heatmap" proposed by Kivimäk (2023) offer a promising way to systematically identify and prioritize accessibility challenges in software development, but these tools must be adopted more widely across projects to have a meaningful impact. FOSS communities must also recognize that inclusivity is a continuous process of reflection and adaptation, not a checkbox to be marked off. By embedding equity into the very structure of how projects are governed, developed, and maintained, FOSS can better fulfill its promise of openness and collective empowerment.

Conclusion

While FOSS communities are built on principles of openness, collaboration, and accessibility, persistent structural barriers limit the participation of marginalized groups. The research highlights that knowledge diversity enhances innovation and productivity, yet accessibility challenges, language constraints, and governance structures continue to disadvantage underrepresented contributors. Findings from GNOME's socio-technical evolution (Smith et al. 2001) reveal an uneven distribution of workload, with a small core group of contributors maintaining critical aspects of development, exacerbating sustainability risks due to the "bus factor." Usability studies further indicate that design inconsistencies in open-source software (FOSS) projects, such as GNOME, contribute to user and developer frustration, underscoring the need for improved UI and accessibility standards.

Social diversity studies in GitHub teams (Vasilescu et al., 2015) demonstrate that diverse teams can improve productivity; however, gender representation remains highly imbalanced, with 75.3% of projects lacking gender diversity. Country and tenure diversity were associated with positive effects on team stability and output, but challenges such as smaller team sizes and retention issues persist. These findings suggest that while diversity offers measurable benefits, it requires structural support to be sustainable.

For FOSS to fulfill its promise of inclusivity, communities must implement targeted initiatives that address accessibility, governance, and contributor retention. This includes clearer onboarding processes, mentorship programs, and policies that actively support underrepresented contributors. By fostering a more inclusive and equitable environment, FOSS can better align with its core values and ensure that all developers, regardless of background or ability, have the opportunity to contribute meaningfully.

Works Cited

- Albusays, K., Ford, D., Parnin, C., & Murphy-Hill, E. (2021). The diversity crisis in software development. *IEEE Software*, 38(2), 19–25. <https://doi.org/10.1109/MS.2020.3045817>
- Bogers, M., Foss, N. J., & Lyngsie, J. (2018). The “human side” of open innovation: The role of employee diversity in firm-level openness. *Research Policy*, 47(1), 218–231. <https://doi.org/10.1016/j.respol.2017.10.012>
- Castro, J. O. (2003, September 10). Inside the GNOME 2.4 desktop and developer platform. *Ars Technica*. <https://arstechnica.com/information-technology/2003/09/gnome2-4/>
- Coleman, E. G. (2013). *Coding freedom: The ethics and aesthetics of hacking*. Princeton University Press.
- Kivimäki, S. (2023). A study into ideal software development for the disabled [Bachelor’s thesis, LAB University of Applied Sciences]. Theseus. <https://www.theseus.fi/handle/10024/800899>
- Liang, T.-P., Liu, C.-C., Lin, T.-M., & Lin, B. (2007). Effect of team diversity on software project performance. *Industrial Management & Data Systems*, 107(5), 636–653. <https://doi.org/10.1108/02635570710750408>
- Rocha, T. A. da, de Souza, C., Teran, L., & Mota, M. (2024). Effective inclusion of people with disabilities in software development teams. In *Proceedings of the 18th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '24)* (pp. 447–453). Association for Computing Machinery. <https://doi.org/10.1145/3674805.3690749>
- Schwartz, S. D., Fickas, S. F., Norris, B., & Dubey, A. (2024). A survey of open source software repositories in the US Department of Energy’s national laboratories. *Computing in Science & Engineering*, 26(2), 1–8. <https://doi.org/10.1109/MCSE.2024.3414951>
- Smith, S., Engen, D., Mankoski, A., Frishberg, N., Pedersen, N., & Benson, C. (2001). *GNOME usability study report*. Sun Microsystems, Inc.
- Vasilescu, B., Serebrenik, A., & Filkov, V. (2015). A data set for social diversity studies of GitHub teams. In *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories* (pp. 514–517). IEEE. <https://doi.org/10.1109/MSR.2015.77>