Sustaining FOSS Communities in the Age of AI: Modeling Human-Centric Social Collaboration in the Age of Hybrid AI Software Development

Summary

This thesis studies how Free and Open Source Software (FOSS) communities change as AI agents become part of software development. Its main goal is to understand how AI-mediated collaboration affects coordination, review, responsibility, and governance in large-scale software ecosystems. Rather than focusing only on code generation or individual productivity, the project examines AI at the level of the community. It asks how human roles shift when agents take over more operational tasks and humans move toward supervision, arbitration, and collective decision-making.

The innovative aspect of the thesis is its community-first perspective. Current research on AI in software engineering mainly evaluates technical performance or developer assistance. By contrast, this project studies how AI reshapes the social and organizational structure of collaborative software production. It positions FOSS as a key empirical setting because its repositories preserve rich traces of coordination, review, and discussion over time.

Methodologically, the thesis combines large-scale trace analysis, computational modeling, and governance-oriented interpretation. It proceeds in five steps: data collection, construction of a Collaboration Knowledge Graph, pattern discovery, analysis of AI-mediated collaboration, and translation of findings into governance principles. The project uses these methods to reconstruct existing coordination patterns and analyze how they change when AI agents mediate contributions and workflows.

Expected results include new models and datasets for studying hybrid human-AI collaboration in FOSS, as well as analytical tools to make coordination dynamics more visible at ecosystem scale. The thesis will also propose governance principles to help communities preserve accountability, transparency, and sustainability under AI-augmented development. More broadly, it aims to support human-centered forms of AI integration in the digital commons.

Keywords: Free and Open Source Software (FOSS); Human-AI Collaboration; Knowledge Graphs; Graph Neural Networks; Digital Commons Governance

1. Scientific Context and Motivation

Free and Open Source Software (FOSS) ecosystems are sustained by large-scale forms of collaboration in which participants coordinate work, negotiate contributions, review changes, and maintain shared infrastructures over time. Recent advances in AI suggest not only new forms of individual assistance in software development, but a deeper transformation in how collaborative work may be organized.

In particular, AI agents are beginning to mediate activities that were previously carried out directly by human contributors, such as drafting code, adapting artifacts, proposing modifications, and assisting in review and integration tasks. This evolution raises a broader question: **if agents increasingly handle operational tasks, how are human roles, coordination processes, and governance structures reconfigured in collaborative software communities?**

This project starts from the assumption that the key challenge is no longer only how AI improves individual productivity, but how AI reshapes collective production. In this sense, FOSS provides a particularly relevant setting: it is a long-standing model of large-scale digital collaboration, and one in which transparency, accountability, maintainability, and community norms are central to sustainability.

2. Core Problem and Originality

This project adopts a community-first perspective, where the primary unit of analysis is not the individual developer or the code artifact, but the evolving structure of collaboration within FOSS ecosystems. From this perspective, AI is not studied as an isolated tool, but as an active element in the reconfiguration of community dynamics.

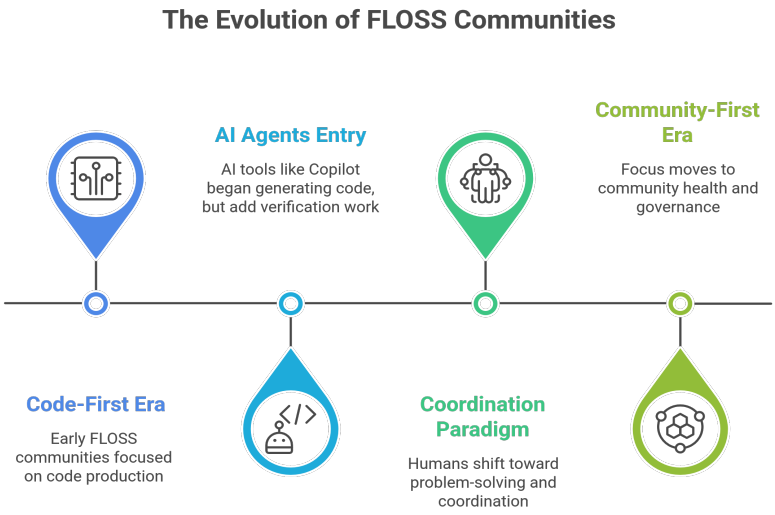

Accordingly, the project investigates a transition from direct human-to-human collaboration toward collaboration increasingly mediated by AI agents. This is not simply a matter of better tools; it implies a qualitative shift in the organization of contribution, coordination, and decision-making.

A useful illustration is the case of a software fork. In current FOSS practice, a contributor who wants to adapt an existing system typically studies the codebase, implements changes, tests them, and eventually proposes integration through pull requests and discussion with maintainers. In a more agent-mediated setting, by contrast, contributors may delegate much of this operational work to AI agents that inspect repositories, generate adaptations, compare implementations, test alternatives, and prepare synchronization proposals across forks. Human contributors would then intervene less at the level of direct implementation and more at the level of supervision, arbitration, and governance (Zheng et al., 2023; Vaithilingam et al., 2022; Nijkamp et al., 2023).

This shift matters because the introduction of agents into collaborative workflows does not only affect efficiency. It also changes how responsibility is distributed, how decisions are justified, how trust is established, and how collective oversight can be maintained. In this sense, the emergence of AI-mediated collaboration raises a broader socio-technical problem: how to preserve intelligibility, accountability, and sustainability when collective work is increasingly performed through computational intermediaries.

3. Research Gap

Existing research does not yet provide adequate models for understanding this transformation at the level of the ecosystem. Work on AI in software engineering has largely focused on productivity, code generation, and developer assistance, often at the level of individual users or small teams. Research on FOSS and online communities, by contrast, has developed strong accounts of coordination, governance, and collective sustainability, but without fully addressing the role of AI agents as active mediators of contribution.

As a result, an important gap remains: we still lack empirical and conceptual frameworks for analyzing how agent-mediated collaboration reshapes the social organization of decentralized software communities. In particular, current literature does not sufficiently explain how AI mediation affects role distribution, interaction patterns, review practices, authority relations, and the long-term resilience of collaborative ecosystems.

This gap is especially significant because early evidence suggests an ambivalent dynamic. Generative AI tools may improve productivity, but they can also increase technical debt and impose additional verification burdens on maintainers

(Yeverechyahu et al., 2025; Dhanore, 2025). More fundamentally, current systems still lack what can be described as social intelligence, understood here as the capacity to produce actions that are not only technically correct but also intelligible and appropriately situated within a collaborative process by taking into account shared context, intentions, and norms. This limitation is reinforced by recent benchmarks such as CooperBench, which show that coding agents still struggle to coordinate effectively with humans and with one another despite strong technical capabilities (Khatua et al., 2026).

The central research problem of this thesis is therefore not simply whether AI can support software development, but how AI-mediated collaboration reconfigures coordination and governance in FOSS ecosystems, and under what conditions these transformations can remain transparent, accountable, and sustainable.

This leads to the following central research question:

How does AI-mediated collaboration reconfigure coordination, governance, and sustainability in large-scale FOSS ecosystems, and which human-centered design principles are needed to preserve meaningful human oversight and community resilience?

4. Methodology

The methodology of this project starts from a simple idea: FOSS ecosystems leave rich traces of collaboration over time. These traces do not only document code production, but also how contributors coordinate tasks, review changes, negotiate decisions, and distribute responsibility across the community. Because these interactions are recorded in repositories, issues, pull requests, discussions, and review histories, FOSS provides a strong empirical basis for studying how large-scale software communities organize collaborative work.

The project uses these traces to reconstruct patterns of human-centered coordination in decentralized software ecosystems. This provides a baseline for analyzing how such patterns may change when AI agents begin to mediate contributions, reviews, and workflow decisions. Here, computational modeling is not meant to replace human judgment, but to make collaboration dynamics visible and analyzable at ecosystem scale.

Methodologically, the project combines large-scale empirical analysis, computational modeling, and governance-oriented interpretation. The research proceeds in five steps: data collection, construction of a collaboration knowledge graph, pattern discovery, analysis of AI-mediated collaboration, and translation of findings into governance principles for human-centered FOSS ecosystems.

4.1 Data Collection

We will collect large-scale data from software archives such as GitHub, Software Heritage, and other open development platforms. These datasets will include commits, pull requests, issue discussions, review threads, and interaction histories, allowing us to reconstruct collaborative dynamics over time.

4.2 Collaboration Knowledge Graph (CoKG)

We will build a Collaboration Knowledge Graph (CoKG) representing the structure of FOSS development (Zhong et al., 2023; Guo et al., 2020; Guo et al., 2024).

Nodes represent contributors, artifacts, issues, and discussions, while edges capture relationships such as authorship, review interactions, dependencies, and workflow transitions.

The graph integrates both social and technical layers of collaboration. It will be enriched using existing ontologies (Smith et al., 2011; Nundlall & Nagowah, 2022), as well as NLP techniques, including LLM-based classification of discussion content and Natural Language Inference (NLI) to extract relations such as dependency, resolution, and conflict between tasks and discussions.

4.3 Pattern Discovery

To identify coordination structures, contributor roles, workflows, and bottlenecks, we will apply graph-based mining techniques such as:

- gSpan algorithm (Yan et Han, 2002)
- FSG algorithm (Kuramochi et al., 2001)
- FPGraphMiner (Vijayalakshmi et al., 2011)

These approaches will be combined with Graph Neural Networks (GNNs) to capture global structural dependencies and improve reasoning over large collaborative graphs (Xu et al., 2025).

4.4 Analysis of AI-Mediated Collaboration

We will analyze how AI agents influence collaborative processes through empirical studies, case analyses, and simulations of AI-assisted development workflows.

We will study how mentorship structures evolve, how peer review processes change, and how consensus formation is affected when AI agents mediate contributions.

Evaluation will leverage CooperBench (Khatua et al., 2026), as well as datasets such as Allard et al., 2019.

4.5 Governance and Ecosystem Implications

Finally, we will translate findings into governance principles for AI-mediated FOSS ecosystems. This includes defining roles, coordination structures, and transparency mechanisms aligned with initiatives such as the Linux Foundation's Responsible Generative AI framework.

The objective is to ensure that:

- AI contributions remain interpretable to human maintainers
- coordination failures can be detected early
- and community health and sustainability can be preserved under AI-augmented development pressure

5. Expected Contributions and Societal Relevance

Scientifically, the project will contribute to understanding an emerging form of collaboration in which collective work is increasingly mediated by intelligent systems. It will bridge research on FOSS governance, large-scale digital collaboration, and human-AI interaction by focusing on the community level rather than only on individual use.

In particular, the thesis introduces a community-first perspective, where the primary unit of analysis is not the individual developer or the code artifact, but the evolving structure of collaboration within FOSS ecosystems. AI is therefore studied not as a tool, but as an active participant in reshaping community dynamics.

More specifically, the project proposes computational models such as collaboration knowledge graphs enriched with social and technical relations, and analytical methods that make both human and AI contributions legible at the ecosystem level. These tools aim to better understand coordination, governance, and collective health in hybrid human-AI systems.

Societally, the project addresses a significant issue for the future of the digital commons. FOSS infrastructures support critical public and private systems, yet their sustainability depends on forms of coordination, trust, and participation. If AI agents become deeply integrated into these ecosystems, it will be essential to ensure that automation does not erode transparency, accountability, inclusiveness, or maintainability.

By identifying how governance and collaboration must evolve in these contexts, the project aims to support more sustainable, legible, and human-centered forms of digital production.

6. PEPR eNSEMBLE relevance

This PhD project bridges the scientific objectives of PC4 CONGRATS and PC3 MATCHING by studying how AI agents reshape collaboration, governance, and sustainability in Free/Open Source Software (FOSS) communities—a case of large-scale, hybrid human-AI knowledge production. While CONGRATS focuses on the socio-technical dynamics of online communities and the design of tools to support their viability, MATCHING investigates adaptive human-AI cooperation, authority, and long-term interaction. This project contributes to both by:

- For CONGRATS: Providing an empirical, data-rich case study of FOSS communities as epistemic collectives, where AI agents mediate coordination, review, and decision-making. The project aligns with WP 4.2 (theorizing collaboration) and WP 4.3 (community-centric tools) by analyzing how AI integration affects community health, efficiency, and governance. Deliverables—such as Collaboration Knowledge Graphs and governance principles—will directly inform CONGRATS' efforts to design scalable tools for open collaboration.

- For MATCHING: Offering a real-world testbed for studying human-AI teaming in hybrid groups, where AI agents handle operational tasks (e.g., code generation, review automation) while humans shift toward supervision and arbitration. This addresses Axis 1 (group collaboration with AI) and Axis 3 (long-term co-adaptation), with a focus on preserving human agency, transparency, and skill evolution in AI-augmented workflows.

We suggest a Joint funding by PC4 CONGRATS and PC3 MATCHING as this PhD thesis presents an opportunity to merge the complementary strengths of socio-technical analysis of online communities (CONGRATS) and the study of human-AI interaction (MATCHING).

7. Research Environment

The project brings together complementary expertise in complex networks, software repository mining, and the governance of Free/Libre and Open Source Software (FLOSS) communities.

IMT Atlantique (Lab-STICC, Complex Networks) contributes expertise in network analysis, graph machine learning, and organizational data science, under the supervision of Prof. Cécile Bothorel.

Universidad Rey Juan Carlos (SoftDev) brings internationally recognized expertise in mining software repositories and empirical research on FOSS ecosystems, under the supervision of Prof. Gregorio Robles.

The project also benefits from the broader research environment of the ANR Data2Laws initiative and its partners, including OW2 and the Eclipse Foundation. These connections provide access to real FLOSS communities, contributors' expertise, and potential opportunities for interviews, empirical validation, and experimentation.

The team builds on strong previous achievements, including publications in leading venues such as ICSE, MSR, CSCW, and The Web Conference, as well as contributions to influential FOSS tools and to methods, algorithms, and software for complex network analysis; s; https://gitlab.imt-atlantique.fr/networks for network analysis (roles, detection of communities, anomaly detection in temporal graphs).

Bibliography

Allard, T., Alvino, P., Shing, L., Wollaber, A., & Yuen, J. (2019). A dataset to facilitate automated workflow analysis. PLoS ONE, 14(2), e0211486. https://doi.org/10.1371/journal.pone.0211486

CHAOSS (2026, March 11). https://chaoss.community : KPIs for good health of open-source communities.

Dabbish, L., Stuart, C., Tsay, J., & Herbsleb, J. (2012). Social coding in GitHub: Transparency and collaboration in an open software repository. In Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (pp. 1277–1286). https://doi.org/10.1145/2145204.2145396

Dhanore, Y. M. (2025). The Impact of Generative AI Tools in Open-Source Software Development. Master's Degree Project in IT Project Management, Spring term 2025, Supervisor: Binish Tanveer, Department of Computer and Systems Sciences, Stockholm University. https://www.diva-portal.org/smash/get/diva2:1979979/FULLTEXT01.pdf

Guo, D., Ren, S., Lu, S., Feng, Z., Tang, D., Liu, S., ... & Zhou, M. (2020). Graphcodebert: Pre-training code representations with data flow. arXiv preprint arXiv:2009.08366

Guo, Taicheng and Chen, Xiuying and Wang, Yaqi and Chang, Ruidi and Pei, Shichao and Chawla,Nitesh V. and Wiest, Olaf and Zhang, Xiangliang, Large Language Model Based Multi-agents: A Survey of Progress and Challenges, IJCAI24 (2024). https://doi.org/10.24963/ijcai.2024/890

Kalliamvakou, E., Gousios, G., Singer, L., & Damian, D. (2014). The promises and perils of mining GitHub. In Proceedings of the 11th Working Conference on Mining Software Repositories (pp. 92–101). https://doi.org/10.1145/2597073.2597074

Khatua, A., Zhu, H., Tran, P., Prabhudesai, A., Sadrieh, F., Lieberwirth, J. K., ... & Yang, D. (2026). CooperBench: Why Coding Agents Cannot be Your Teammates Yet. arXiv. https://arxiv.org/abs/2601.13295

Kuramochi, M, Karypis, G. Frequent subgraph discovery. In: Proc. 1st IEEE Int. Conf. on Data Mining, San Jose, USA, 29 November - 2 December, 2001, 2001:313–320.

Nijkamp, E., Pang, B., et al. (2022). CodeGen: An Open Large Language Model for Code with Multi-Turn Program Synthesis. arXiv. https://arxiv.org/abs/2203.13474

Nundlall C, Nagowah SD. Task allocation and coordination process in distributed agile software development: an ontology based approach. Inf Technol Manag. 2022;23(3):167-192. https://doi.org/10.1007/s10799-022-00365-9.

Smith, B. L., Tamma, V., & Wooldridge, M. (2011). An Ontology for Coordination. Applied Artificial Intelligence, 25(3), 235–265. https://doi.org/10.1080/08839514.2011.553376

Vaithilingam, P., Zhang, T., & Glassman, E. L. (2022). Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (Article 332, pp. 1–7). Association for Computing Machinery. https://dl.acm.org/doi/10.1145/3491101.3519665

Vijayalakshmi, R, Nadarajan, R, Roddick, JF, Thilaga, M, Nirmal, P. FP-GraphMiner-A
Fast Frequent Pattern Mining Algorithm for Network Graphs. Journal of Graph Algorithms and Applications, 2011, 15(6):753–76.

Xu W. et al., "Subgraph-Aware Joint Modeling for Graph Classification via Efficient Frequent Pattern Mining and Fusion," in IEEE Access, vol. 13, pp. 151894-151906, 2025,
doi:10.1109/ACCESS.2025.3602411

Yan, X, Han, J. g. Span: Graph-Based Substructure Pattern Mining. In: Proc. 2002 IEEE Intern. Conf. Data Mining, Maebashi City, Japan, 9-12 December, 2002:721–724)
Yeverechyahu, D., Mayya, R., & Oestreicher-Singer, G. (2024). The impact of large language models on open-source innovation: Evidence from GitHub Copilot. SSRN. https://ssrn.com/abstract=4684662

Zheng, Z., et al. (2023). A Survey of Large Language Models for Code: Evolution, Benchmarking, and Future Trends. arXiv. https://arxiv.org/abs/2311.10372

Zhong Lingfeng, Jia Wu, Qian Li, Hao Peng, and Xindong Wu. 2023. A Comprehensive Survey on Automatic Knowledge Graph Construction. ACM Comput. Surv. 56, 4, Article 94 (November 2023), 62 pages. https://doi.org/10.1145/3618295