

# Do Developers Adopt Green Architectural Tactics for ML-Enabled Systems? A Mining Software Repository Study

Vincenzo De Martino  
Software Engineering (SeSa) Lab  
University of Salerno, Italy  
vdemartino@unisa.it

Silverio Martínez-Fernández  
Universitat Politècnica de Catalunya  
Spain  
silverio.martinez@upc.edu

Fabio Palomba  
Software Engineering (SeSa) Lab  
University of Salerno, Italy  
fpalomba@unisa.it

**Abstract**—As machine learning (ML) and artificial intelligence (AI) technologies become more widespread, concerns about their environmental impact are increasing due to the resource-intensive nature of training and inference processes. Green AI advocates for reducing computational demands while still maintaining accuracy. Although various strategies for creating sustainable ML systems have been identified, their real-world implementation is still underexplored. This paper addresses this gap by studying 168 open-source ML projects on GitHub. It employs a novel large language model (LLM)-based mining mechanism to identify and analyze green strategies. The findings reveal the adoption of established tactics that offer significant environmental benefits. This provides practical insights for developers and paves the way for future automation of sustainable practices in ML systems.

**Index Terms**—Green AI; Machine Learning-Enabled Systems; Software Sustainability; Empirical Software Engineering.

## I. INTRODUCTION

Systems incorporating AI or ML, often referred to as AI-enabled or ML-enabled systems, are becoming increasingly common [1]. The growing use of these systems raises concerns about their sustainability, particularly in terms of energy consumption and CO<sub>2</sub> emissions, which have become increasingly pressing [2], [3]. While the focus has traditionally been on maximizing accuracy, there is increasing recognition of the need to balance this with environmental impact, leading to the concept of Green AI [4]. Green AI advocates for reducing computational demands while maintaining accuracy, emphasizing a responsible approach to AI development and deployment [5]. In recent years, the software engineering (SE) community has made considerable efforts in promoting software sustainability, with specific attention now being paid to ML-enabled systems [6], [7], [8]. This shift has identified various green tactics to reduce the environmental footprint of ML models. [6], [9]. For instance, Järvenpää et al. [9] cataloged 30 green architectural tactics for ML-enabled systems through a literature review and validated by experts. However, the actual relevance of these tactics is not yet established.

▲ *There is still a gap of knowledge on the extent to which green tactics are actually implemented in real-world contexts.*

Our objective seeks to determine the extent to which the green tactics cataloged by Järvenpää et al. [9] are adopted in the development of ML-enabled systems. Understanding this level of adoption is essential for bridging the gap between state-of-the-art research and the practices employed in real-world projects. Hence, we ask:

**Q RQ:** *To what extent are green architectural tactics for ML-enabled systems adopted in software projects?*

This paper addresses gaps in previous studies by conducting a mining software repository study to understand the adoption of green tactics in ML systems. Building on the catalog by Järvenpää et al. [9], the study analyzes 168 open-source ML projects on GITHUB using a novel LLM-based mechanism to identify green tactics at the code level.

## II. RELATED WORK

Interest in sustainable and Green AI is growing, with the SE and AI communities publishing numerous studies to identify and catalog green practices that reduce energy consumption throughout the ML pipeline [6], [9], [10], [11]. However, most existing research focuses on identifying these practices rather than evaluating their adoption in real-world projects.

Salehi et al. [12] conducted a systematic review to classify data-centric approaches in Green AI, discussing practical applications, challenges, and future directions. Xu et al. [13] focused on Green Deep Learning, mapping approaches to reduce computational and energy costs in deep learning models while maintaining performance, using a taxonomy of lifecycle stages and artifacts. Building on this, Verdecchia et al. [6] expanded the scope to encompass the entire domain of Green AI, identifying a broader set of characteristics. Järvenpää et al. [9] later synthesized and validated these green practices with input from ML experts, classifying them into clusters, expanding on existing tactics, and identifying new ones.

To sum up, much work has been done to advance knowledge on what approaches can improve the fields of sustainable and green SE. However, to the best of our knowledge, our study is the first to investigate the actual adoption of green architectural

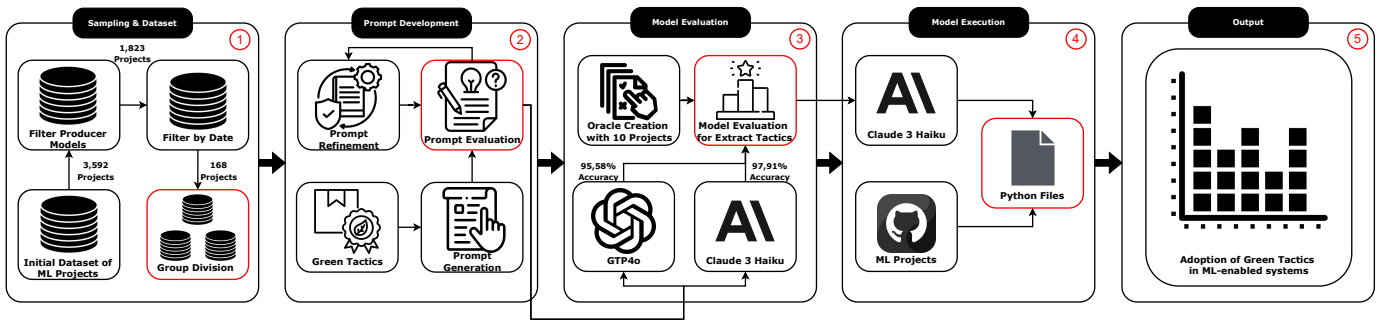


Fig. 1: Overview of the research process.

tactics in real software projects, marking the primary scientific contribution of our work. We specifically study the set of tactics defined in the literature, developing and evaluating an LLM-based tool to assess the adoption and frequency of these tactics in real-world projects.

### III. RESEARCH METHOD

This section discusses the methodology of our work; the main steps are overviewed in Figure 1. To report the results of our exploratory mining study, we adhere to the standards for repository mining as outlined in the “*Repository Mining*” and “*General Standard*” categories of the *ACM/SIGSOFT Empirical Standards*.<sup>1</sup>

#### A. Sampling and Dataset Description

We initially identified a candidate set of 3,592 ML projects coming from the combination of two popular datasets. The first dataset, proposed by Gonzalez et al. [14] and later revised by Rzig et al. [15], contains 3,020 ML projects. The second dataset, NICHE [16], includes 572 ML projects. The datasets include a variety of ML projects, from frameworks and tools to those creating or integrating ML models. However, not all were suitable for studying sustainable tactics in ML model lifecycles. Projects unrelated to ML model training were filtered out, narrowing the original 3,592 projects to a relevant subset aligned with the study’s goals. Projects building and training ML models were identified by analyzing APIs for ML library imports and training method usage [17], using references such as the well-know AI & ML Libraries list [18]. This step refined the dataset to 1,823 projects. Using the GitHub API [19], only projects updated in 2024 were included, following prior research methodologies [20]. This further reduced the dataset to 168 projects, ensuring a focus on active projects.

The final set comprises 168 projects. Table I summarizes the characteristics of these projects, including metrics such as age, number of stars, and number of contributors. Additionally, we use the LOC to classify projects into three size groups: small ( $LOC \leq 10,000$ ), medium ( $10,000 < LOC \leq 100,000$ ), and large ( $LOC > 100,000$ ). We use these groups to make comparisons in our analyses, which are in line with the methodology used by previous studies [21], [20].

<sup>1</sup><https://github.com/acmsigsoft/EmpiricalStandards>.

#### B. Selection of Architectural Green Tactics

After selecting the ML projects for our study, we focused on identifying green architectural tactics using the comprehensive catalog by Järvenpää et al. [9], which build upon the work of Verdecchia et al. [6]. This catalog, the most extensive resource available, organizes 30 green tactics into six clusters addressing various aspects of energy efficiency and sustainability across the ML pipeline. Unlike other studies focusing on domain-specific or isolated strategies [9], [10], [11], it provides a holistic framework for analyzing green practices in ML systems. From these candidate tactics, our study focuses on those pertaining to three specific clusters: ‘*Algorithm design*’, ‘*Model optimization*’, and ‘*Model training*’, encompassing a total of 15 tactics available in our appendix [22]. These particular tactics are operational in nature, as they involve code-level and model-level adjustments that can be implemented using existing libraries or developed manually, depending on the developer’s expertise. This flexibility ensures that the tactics are accessible to any development team, regardless of their resources, and can be effectively analyzed using the content available in GITHUB repositories. While other tactics from the original study are undoubtedly valuable, they present significant challenges in terms of measurability.

#### C. Extraction of Architectural Green Tactics

Extracting architectural green tactics from software repositories may be challenging, as it requires discerning whether developers have implemented specific practices to support sustainability efforts.

More specifically, we mined the latest snapshot of the repositories associated with the selected software projects, extracting all *Python* files, the files that may contain source code related to ML model creation, training, and deployment. To detect architectural green tactics, we applied LLMs directly to the extracted *Python* files. LLMs are highly effective at interpreting the source code, enabling them to identify subtle patterns and context-specific practices within these files [23]. The advanced architecture and extensive training of LLMs enable them to capture complex code relationships and subtle patterns, making them effective for identifying context-specific green tactics in ML projects. By understanding source code context, e.g., recognizing optimized algorithms or model prun-

TABLE I: Average (AVG) and Median (Med) of each group of projects.

Group	#Projects	Age		Stars		Size		Contributors		Forks		Commit	
		Avg	Med	Avg	Med	Avg	Med	Avg	Med	Avg	Med	Avg	Med
Small	60	6,3	5,67	1.101,68	255	3.466,75	2.629,50	9,35	4,5	244,57	65,5	532,63	291,5
Medium	59	6,67	6,63	2.543,78	669	46.739,27	42.446	15,08	13	412,85	148	1.114,46	691
Large	49	6,98	6,6	6.703,76	563	435.995,41	248.857	19,8	21	1.681,33	154	4.257,47	1.506
Total	168	6,63	6,29	3.242,07	519,5	144.817,84	38.744,5	14,41	12	722,72	122,5	1.823,38	674

ing for energy efficiency, LLMs can detect green tactics even without explicit documentation. We chose LLMs over other ML models for several reasons. Their architecture and diverse training enhance their ability to interpret code [24], and their capacity to process large volumes of code suits our study’s scale. Moreover, LLMs can execute tasks from instructions without additional training data [25], which is crucial given the lack of a labeled dataset for detecting green architectural tactics in software repositories.

To this aim, we devised an LLM-based mechanism aiming at detecting the application of green architectural tactics within software code. Our approach employs prompt engineering, where carefully crafted input prompts guide the LLM to generate outputs that accurately reflect the presence of tactics [26], [27]. The prompt was structured into five components to achieve the study objective. First, the ‘*Objective*’ defines the goal, instructing the LLM to analyze source code for sustainable practices, focusing on green architectural tactics for ML. Next, the ‘*Task Description*’ outlines the task to detect 15 predefined sustainable tactics. The ‘*Additional Instructions*’ guide the LLM on how to analyze the code and produce the output. In the ‘*Reporting Instructions*’, the LLM is directed to present findings clearly, including relevant code snippets, and explicitly note the absence of tactics when applicable. Lastly, the ‘*Examples*’ section provides templates to ensure the results are formatted consistently across scenarios. We tested the prompt with GPT-4o (OpenAI) [28], and Claude 3 Haiku (Anthropic) [29], selected for their accuracy and ability to interpret complex unstructured data like source code [30]. The prompt was refined iteratively through collaborative sessions by the first two authors, using tests on a project subset to improve instructions and examples. The prompt’s development process and iterations are detailed in the appendix [22].

#### D. Model Validation

To evaluate the feasibility of using LLMs to identify green architectural practices, we conducted an initial test with GPT-4o and Claude 3 Haiku on 10 projects (86 Python files) selected through random sampling [31], [32]. The first author, leveraging expertise in sustainable SE, manually created an oracle to classify files based on the presence of sustainable tactics using a well-established catalog [9]. This oracle served as the reference for assessing model accuracy. GPT-4o achieved 95.58% accuracy, offering precise but minimal contextual explanations. Claude 3 Haiku performed slightly better with 97.91% accuracy, providing clear natural language

explanations that enhanced understanding of its classifications. Despite occasionally omitting code snippets, Claude 3 Haiku’s insights and cost efficiency (\$0.25/\$1.25 per million tokens vs. GPT-4o \$5/\$15) made it more suitable for analyzing larger datasets. Ultimately, we selected Claude 3 Haiku for its higher accuracy and lower cost. To address its occasional omission of specific code snippets, we implemented a mechanism to track the repository and file where a tactic was detected, ensuring reliable localization even without explicit code output.

#### E. Data Analysis

To ease the data analysis process, we developed a Python script [22] that systematically traverses the files in each repository and records key information. For each project, it generates a `csv` file, with each row corresponding to an analyzed file. The `csv` file captures (1) the file name and (2) the output generated by Claude 3 Haiku. In particular, we implemented a script that processed Haiku’s responses by segmenting the first part of the output—specifically the classification of existing tactics—into separate columns, with each column representing one of the tactics considered in the study. After collecting all responses, the data was aggregated into a single file that tracked the occurrence of each tactic. The script calculated the total frequency of each tactic across all projects and used this information to address the **RQ**, providing an overview of how green architectural tactics are adopted in the considered ML-enabled systems.

#### F. Threats to validity

This section highlights potential factors influencing the study’s conclusions and generalizability [33]. Errors from poorly designed prompts or misinterpreted tactic definitions were mitigated through iterative refinement and the use of well-established definitions from the literature [9]. Risks included randomness in LLM outputs and incomplete tactic extraction, which were addressed by refining prompts, conducting tests, and tracking repositories for verification. The study focuses on Python due to its prominence in ML, ensuring generalizability through dataset quality by using well-established sources like NICHE[16] and the dataset by Gonzalez et al.[14], with preprocessing to exclude outdated or irrelevant projects. Sustainable tactics were selected based on Järvenpää et al.[9] catalog for their suitability for code-level analysis. Default LLM parameters were used to ensure consistency in evaluation. While based on a sizable sample, the findings cannot be fully generalized to all ML projects. A

comparative evaluation of GPT-4o [28] and Claude 3 Haiku [29], using a manually constructed oracle, was conducted to identify the reliable model for accuracy and tactic discovery.

#### IV. ANALYSIS OF THE RESULTS

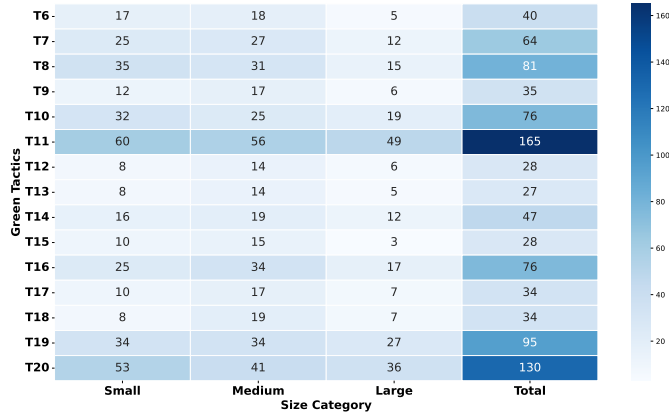


Fig. 2: Frequency of Green Tactics in ML Projects.

Figure 2 shows the adoption frequency of green tactics across 168 ML projects, categorized by size: ‘Small’, ‘Medium’, ‘Large’, and the full set. ‘Medium’ projects showed the highest overall adoption of tactics, with ‘Small’ projects favoring tactics like ‘Decrease model complexity’, and ‘Use dynamic parameter adaptation’. ‘Large’ projects demonstrated the lowest adoption rates, likely due to their complexity and a focus on stability and performance, which can hinder the integration of sustainable tactics [34]. Across all projects, the most widely adopted tactic was ‘Use of built-in library functions’ 165 (98.2%), reflecting the reliance on optimized libraries and frameworks. ‘Design for memory constraints’ 130 (77.4%) was also prevalent, emphasizing efficiency in memory usage. Other frequently adopted tactics included ‘Use of checkpoints during training’ 95 (56.5%), ‘Reducing model complexity’ 81 (48.2%), ‘Use dynamic parameter adaptation’ 76 (45.2%) and ‘Consider transfer learning’ 76 (45.2%). In contrast, tactics like ‘Enhance model sparsity’ 47 (28%) and ‘Choose a lightweight algorithm alternative’ 40 (23.8%) were less common. The least adopted tactics were ‘Consider graph substitution’, ‘Set energy consumption as a model constraint’, and ‘Consider energy aware pruning’ 27-28 (16.1%-16.7%). These low adoption rates may stem from their complexity or a lack of awareness among developers. This variation underscores the need for targeted education and tooling to facilitate the adoption of less-used but potentially impactful green architectural tactics.

🔗 **Answer to RQ.** Green tactics proposed in the literature are widely adopted in ML projects. However, some tactics are less adopted, suggesting the need to raise awareness and give more support to the tools to facilitate their wider integration.

#### V. DISCUSSION AND IMPLICATIONS

Our study highlights several implications for the SE community and society to reduce the energy footprint. When

comparing the adoption of green tactics identified in previous studies, particularly the catalog of 30 tactics by Järvenpää et al. [9], we analyzed 168 open-source ML projects on GitHub using LLMs. Our analysis confirms that practitioners adopt many of these tactics. Specifically, tactics like ‘Use of built-in library functions’ and ‘Design for Memory Constraints’ have high adoption rates, suggesting that these tactics are well-supported and effectively integrated into current development processes. Conversely, other tactics such as ‘Set Energy Consumption as a Model Constraint’ and ‘Consider Graph Substitution’ exhibit lower adoption rates in our study. Although our analysis is based on open-source projects, the results have broader implications for different ML contexts. Many tactics are applicable to proprietary and enterprise-level systems, but some, such as energy-aware optimizations, may require adaptations for specific domains such as resource-constrained environments. Green tactics also hold potential for improving accessibility and affordability by reducing hardware and energy costs, but their adoption must be balanced against potential trade-offs in performance and usability, particularly in critical applications such as healthcare and autonomous systems. Adoption barriers, including organizational resistance, lack of awareness, and perceived costs, are particularly evident in larger projects where coordination and competing priorities complicate sustainability efforts. Addressing these barriers will require better education, improved tooling, and incentives for adopting green tactics. Engaging directly with developers through interviews or analyzing project repositories could uncover additional insights into these challenges and inform strategies for overcoming them.

#### VI. CONCLUSION AND FUTURE WORK

This paper examines the adoption of green tactics in ML-enabled systems through a study of 168 open-source ML projects and introduces a novel LLM-based mechanism for identifying green tactics in software repositories. Our contributions include an empirical analysis of green tactic adoption, a novel LLM tool to identify tactics, and a public replication package with study materials [22]. Future work will explore the impact of green tactics on sustainability through longitudinal studies and the identification of new green tactics. It will also examine socio-technical factors, such as knowledge gaps and implementation challenges while promoting lesser-used tactics with enhanced documentation, tools, and case studies. These steps will help integrate green tactics more effectively into ML development, contributing to the creation of more sustainable systems.

#### ACKNOWLEDGMENT

This work has been partially supported by the European Union - NextGenerationEU through the Italian Ministry of University and Research, Projects PRIN 2022 “QualAI: Continuous Quality Improvement of AI-based Systems” (grant n. 2022B3BP5S, CUP: H53D23003510006), and by the GAISSA research project (ref. TED2021-130923B-I00; MCIN/AEI/10.13039/501100011033).

## REFERENCES

- [1] S. Martínez-Fernández, J. Bogner, X. Franch, M. Oriol, J. Siebert, A. Trendowicz, A. M. Vollmer, and S. Wagner, "Software engineering for ai-based systems: a survey," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 31, no. 2, pp. 1–59, 2022.
- [2] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres, "Quantifying the carbon emissions of machine learning," *arXiv preprint arXiv:1910.09700*, 2019.
- [3] F. Durán, S. Martínez-Fernández, M. Martinez, and P. Lago, "Identifying architectural design decisions for achieving green ml serving," in *Proceedings of the IEEE/ACM 3rd International Conference on AI Engineering-Software Engineering for AI*, 2024, pp. 18–23.
- [4] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green ai," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.
- [5] A. Van Wynsberghe, "Sustainable ai: Ai for sustainability and the sustainability of ai," *AI and Ethics*, vol. 1, no. 3, pp. 213–218, 2021.
- [6] R. Verdecchia, J. Sallou, and L. Cruz, "A systematic review of green ai," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 13, no. 4, p. e1507, 2023.
- [7] D. A. Tamburri, "Sustainable mlps: Trends and challenges," in *2020 22nd international symposium on symbolic and numeric algorithms for scientific computing (SYNASC)*. IEEE, 2020, pp. 17–23.
- [8] S. Martínez-Fernández, X. Franch, and F. Durán, "Towards green ai-based software systems: an architecture-centric approach (gaissa)," *arXiv preprint arXiv:2307.09964*, 2023.
- [9] H. Järvenpää, P. Lago, J. Bogner, G. Lewis, H. Muccini, and I. Ozkaya, "A synthesis of green architectural tactics for ml-enabled systems," in *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Society*, 2024, pp. 130–141.
- [10] H. K. Natarajan, D. de Paula, C. Dremel, and F. Uebernickel, "A theoretical review on ai affordances for sustainability," 2022.
- [11] S. Mittal, "A survey on optimized implementation of deep learning models on the nvidia jetson platform," *Journal of Systems Architecture*, vol. 97, pp. 428–442, 2019.
- [12] S. Salehi and A. Schmeink, "Data-centric green artificial intelligence: A survey," *IEEE Transactions on Artificial Intelligence*, 2023.
- [13] J. Xu, W. Zhou, Z. Fu, H. Zhou, and L. Li, "A survey on green deep learning," *arXiv preprint arXiv:2111.05193*, 2021.
- [14] D. Gonzalez, T. Zimmermann, and N. Nagappan, "The state of the ml-universe: 10 years of artificial intelligence & machine learning software development on github," in *Proceedings of the 17th International conference on mining software repositories*, 2020, pp. 431–442.
- [15] D. E. Rzig, F. Hassan, C. Bansal, and N. Nagappan, "Characterizing the usage of ci tools in ml projects," in *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, 2022, pp. 69–79.
- [16] R. Widayarsi, Z. Yang, F. Thung, S. Q. Sim, F. Wee, C. Lok, J. Phan, H. Qi, C. Tan, Q. Tay *et al.*, "Niche: A curated dataset of engineered machine learning projects in python," *arXiv preprint arXiv:2303.06286*, 2023.
- [17] V. De Martino, G. Recupito, G. Giordano, F. Ferrucci, D. Di Nucci, and F. Palomba, "Into the ml-universe: An improved classification and characterization of machine-learning projects," *Available at SSRN 4960056*.
- [18] J. Misiti, "Awesome machine learning: A curated list of awesome machine learning frameworks, libraries and software," <https://github.com/josephmisiti/awesome-machine-learning>, 2015, accessed: 2024-08-31.
- [19] "GitHub REST API documentation - GitHub Docs — docs.github.com," <https://docs.github.com/en/rest>, 2022.
- [20] J. H. Bernardo, D. A. Da Costa, S. Q. d. Medeiros, and U. Kulesza, "How do machine learning projects use continuous integration practices? an empirical study on github actions," in *Proceedings of the 21st International Conference on Mining Software Repositories*, 2024, pp. 665–676.
- [21] W. Felidré, L. Furtado, D. A. Da Costa, B. Cartaxo, and G. Pinto, "Continuous integration theater," in *2019 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. IEEE, 2019, pp. 1–10.
- [22] V. D. Martino, S. Martínez-Fernández, and F. Palomba, "Do Developers Adopt Green Architectural Tactics for ML-Enabled Systems? A Mining Software Repository Study," 9 2024. [Online]. Available: [https://figshare.com/articles/dataset/Do\\_Developers\\_Adopt\\_Green\\_Architectural\\_Tactics\\_for\\_ML-Enabled\\_Systems\\_A\\_Mining\\_Software\\_Repository\\_Study/27101503](https://figshare.com/articles/dataset/Do_Developers_Adopt_Green_Architectural_Tactics_for_ML-Enabled_Systems_A_Mining_Software_Repository_Study/27101503)
- [23] C. V. Lopes, V. I. Klotzman, I. Ma, and I. Ahmed, "Commit messages in the age of large language models," *arXiv preprint arXiv:2401.17622*, 2024.
- [24] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, and H. Wang, "Large language models for software engineering: A systematic literature review," *arXiv preprint arXiv:2308.10620*, 2023.
- [25] B. Zhang, B. Haddow, and A. Birch, "Prompting large language model for machine translation: A case study," in *International Conference on Machine Learning*. PMLR, 2023, pp. 41 092–41 110.
- [26] Y. Zhou, A. I. Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba, "Large language models are human-level prompt engineers," *arXiv preprint arXiv:2211.01910*, 2022.
- [27] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [28] "Chatgpt api keys," <https://openai.com/index/openai-api/>, 2024.
- [29] "Getting started - anthropic," <https://docs.anthropic.com/en/api/getting-started>, 2024.
- [30] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang *et al.*, "A survey on evaluation of large language models," *ACM Transactions on Intelligent Systems and Technology*, vol. 15, no. 3, pp. 1–45, 2024.
- [31] S. Baltes and P. Ralph, "Sampling in software engineering research: A critical review and guidelines," *Empirical Software Engineering*, vol. 27, no. 4, p. 94, 2022.
- [32] S. K. Thompson, *Sampling*. John Wiley & Sons, 2012, vol. 755.
- [33] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in software engineering*. Springer Science & Business Media, 2012.
- [34] J. M. Gonzalez-Barahona, G. Robles, M. Michlmayr, J. J. Amor, and D. M. German, "Macro-level software evolution: a case study of a large software compilation," *Empirical Software Engineering*, vol. 14, pp. 262–285, 2009.