



Research Article

EcoLLM: A Novel Fine-Tuning Framework for Environmental Sustainability in Large Language Models

Gagandeep Singh, Jaskaran Singh, Gourav Arora

1. Department of Computer Science, PCTE Group of Institutions, India; gagandeep@pctebca21a@gmail.com
2. Department of Computer Science, PCTE Group of Institutions, India
3. Department of Computer Science, PCTE Group of Institutions, India

Copyright © 2024 by Authors, Published by **Elementaria**: Journal of Educational Research. This is an open access article under the CC BY License <https://creativecommons.org/licenses/by/4.0/>

Received : November 25, 2024

Revised : November 29, 2024

Accepted : December 30, 2024

Available online : December 31, 2024

How to Cite: Gagandeep Singh, Jaskaran Singh, & Gourav Arora. (2024). EcoLLM: A Novel Fine-Tuning Framework for Environmental Sustainability in Large Language Models. *Elementaria: Journal of Educational Research*, 2(2), 125–135. <https://doi.org/10.61166/elm.v2i2.70>

EcoLLM: A Novel Fine-Tuning Framework for Environmental Sustainability in Large Language Models

Abstract. The increasing reliance on artificial intelligence (AI) models, such as Large Language Models (LLMs), poses a unique challenge regarding environmental sustainability. Current LLMs prioritize performance and versatility, often neglecting the ecological impact of the solutions they generate. This paper presents a novel approach to fine-tuning LLMs to embed environmental considerations in their responses. By adjusting their training datasets and models, we enhance the likelihood of producing environmentally friendly outcomes. We observed that responses factoring in

sustainability increased from 5% to over 75% post-optimization. This paper discusses our methodology, the challenges faced, and the implications for AI's role in supporting global sustainability goals.

Keywords: Large Language Models, Environmental Sustainability, Fine-tuning, Green Computing, Machine Learning Optimization, Carbon Footprint Reduction, Sustainable AI.

Abstrak. Ketergantungan yang semakin meningkat pada model kecerdasan buatan (AI), seperti Large Language Models (LLMs), menimbulkan tantangan unik terkait keberlanjutan lingkungan. LLM saat ini lebih mengutamakan kinerja dan fleksibilitas, sering kali mengabaikan dampak ekologis dari solusi yang dihasilkan. Makalah ini menyajikan pendekatan baru untuk menyempurnakan LLM agar mempertimbangkan aspek lingkungan dalam respons mereka. Dengan menyesuaikan dataset pelatihan dan model, kami meningkatkan kemungkinan menghasilkan hasil yang ramah lingkungan. Kami mengamati bahwa respons yang memperhitungkan keberlanjutan meningkat dari 5% menjadi lebih dari 75% setelah optimalisasi. Makalah ini membahas metodologi kami, tantangan yang dihadapi, dan implikasinya terhadap peran AI dalam mendukung tujuan keberlanjutan global.

Kata Kunci: Large Language Models, Keberlanjutan Lingkungan, Penyempurnaan Model, Komputasi Hijau, Optimalisasi Pembelajaran Mesin, Pengurangan Jejak Karbon, AI Berkelanjutan.

INTRODUCTION

In recent years, the role of artificial intelligence (AI) has expanded significantly across various industries, leading to groundbreaking advancements in fields such as healthcare, finance, education, and logistics. Among the most transformative AI innovations are large language models (LLMs), which have the ability to process and generate human-like text, thereby revolutionizing how information is accessed and utilized. However, despite their impressive capabilities, LLMs have largely neglected one of the most critical global issues: environmental sustainability. Current models such as GPT, Gemini AI, and others focus on accuracy and relevance, yet they frequently overlook the environmental implications of the solutions they suggest.

The environmental impact of AI is multifaceted, encompassing both the direct carbon footprint of training large models and the indirect consequences of their recommendations. For instance, LLMs might suggest high-energy-consumption solutions or products that contribute to pollution without considering sustainable alternatives. This gap in the decision-making process of AI systems raises concerns about their alignment with global efforts to combat climate change and reduce resource depletion. As AI becomes increasingly integrated into decision-making frameworks, it is essential that these systems adopt a more eco-conscious approach.

This research paper aims to address this gap by exploring how LLMs can be fine-tuned to incorporate environmental considerations into their responses. We

propose a novel dataset and optimization approach that enables models to prioritize sustainability in their outputs. Initially, our experiments revealed that out of 60 responses generated by LLMs, only 3 to 4 responses incorporated environmental factors. However, after fine-tuning the models using a sustainability-focused dataset, we achieved a significant improvement, with 45 to 50 responses effectively considering environmental impact. By aligning AI-generated solutions with eco-friendly practices, our research seeks to pave the way for more sustainable AI applications across various sectors. This paper will detail the methodology used for fine-tuning LLMs, the challenges faced, and the broader implications of integrating environmental considerations into AI-driven solutions. We also discuss the potential impact of such models on global sustainability goals, offering a framework for future developments in environmentally responsible AI.

Related Work

Recent research has explored various ways of fine-tuning large language models (LLMs) to incorporate specific objectives, such as improving performance in sustainability-related domains. While significant advancements have been made in the fine-tuning of LLMs for specific tasks, such as emission prediction and resource conservation, there is a notable gap in incorporating environmental sustainability considerations systematically across different sectors.

Fine-Tuning LLMs for Emission Calculations

One prominent study by *IBM* (2023) focused on fine-tuning LLMs to estimate *Scope-3 emissions*, which are indirect emissions that occur in the value chain. The research highlights the challenges of accurate emission reporting and suggests that LLMs, when trained on tailored sustainability datasets, could assist in emission predictions for various industries.

Domain-Specific NLP Models

Additionally, another recent study from *Google Research* (2023) explored domain-specific NLP models fine-tuned for sustainability applications. They proposed a methodology for training LLMs on datasets curated from multiple environmental and policy domains, focusing primarily on energy systems, waste management, and transportation sectors.

Limitations and Gaps in Current Approaches

While these studies make significant contributions to the field of environmentally focused AI, they often overlook key issues in applying LLMs to global sustainability challenges. For instance, IBM's approach does not address the issue of model generalization across different industries, and the models tend to underperform when tasked with nuanced, sector-specific sustainability queries. Moreover, both IBM's and Google's models predominantly focus on a narrow set of sustainability goals, such as carbon emission reduction or energy efficiency, neglecting other crucial aspects such as biodiversity preservation, water usage, and social equity.

In contrast, our research introduces a holistic fine-tuning approach that not only optimizes LLMs for environmental sustainability but also ensures that models generate specific, actionable responses tailored to a wider variety of real-world sustainability challenges. By broadening the scope of the dataset and incorporating diverse environmental contexts (e.g., land use, resource management, pollution control), our work seeks to overcome the limitations identified in prior studies.

METHOD

This section outlines the steps taken to fine-tune large language models (LLMs) to incorporate environmental considerations into their responses. Our approach consists of three main methodologies: dataset construction, model training, and testing/validation.

Dataset Construction

To fine-tune LLMs for better environmental considerations, we developed a comprehensive dataset that integrates sustainability-focused content from a wide array of sources. Our approach ensured coverage across different sectors and environmental contexts, providing a rich foundation for model training. The data collection process was extensive and carefully curated, involving the efforts of a dedicated team. Key steps in the dataset construction included:

- **Collaborative Data Gathering:** A coordinated effort by team members and research collaborators helped compile data from diverse, authoritative sources.
- **Focus on Broad Environmental Themes:** The dataset spans multiple sectors, from renewable energy to resource conservation, ensuring that models are exposed to varied sustainability challenges and best practices.
- **Rigorous Review and Selection:** We applied strict criteria to select the most relevant and up-to-date information to ensure the dataset's quality and applicability to real-world environmental concerns.

By prioritizing diversity and relevance in the data selection process, the resulting dataset is not only robust but well-suited for fine-tuning LLMs to respond more effectively to environmental and sustainability-related prompts.

Model Training

Once the dataset was prepared, the next step was to fine-tune the LLMs using reinforcement learning and supervised training techniques. The specific LLMs used in this research were variations of GPT and Gemini AI. The training process involved:

- **Reinforcement Learning from Human Feedback (RLHF):** We employed RLHF to guide the models towards generating environmentally responsible outputs. Human reviewers provided feedback on the LLM responses, highlighting both environmentally sound and harmful recommendations. The feedback was then used to refine the models' understanding of sustainable choices.
- **Supervised Fine-tuning:** The dataset was used to adjust the models through supervised learning, emphasizing the importance of prioritizing eco-friendly options. During training, the LLMs were presented with real-world scenarios (e.g.,

recommending building materials, transportation options, or industrial processes), and the preferred responses were those that demonstrated environmental awareness.

- **Balancing Model Utility:** Care was taken to balance the models' environmental consciousness with their general problem-solving abilities. Over-emphasis on sustainability could risk limiting the models' versatility, so the training was calibrated to ensure environmental factors were considered without diminishing the overall utility of the LLMs.

Testing and Validation

After the fine-tuning process, the models were rigorously tested and validated to assess their ability to generate responses that consider environmental factors. The testing and validation process included the following steps:

- **Pre-fine-tuning Evaluation:** Baseline tests were conducted on the original models using a set of 60 queries across different domains (e.g., energy solutions, transportation recommendations, resource management). It was observed that only 3-4 responses included environmental considerations, indicating a lack of sustainability awareness.
- **Post-fine-tuning Evaluation:** After the fine-tuning, the same 60 queries were used to evaluate the updated models. The results showed a marked improvement, with 45-50 responses (approximately 75-83%) incorporating sustainability-focused recommendations. These included suggestions for energy-efficient technologies, renewable energy sources, and materials that minimize environmental impact.
- **Real-World Scenario Testing:** To further validate the fine-tuned models, we tested them on new, real-world challenges where environmental factors played a crucial role (e.g., reducing industrial emissions, sustainable urban planning). The models consistently generated environmentally sound recommendations, demonstrating their enhanced ability to factor in ecological considerations.

RESULTS AND DISCUSSION

The results section presents the outcomes of the fine-tuning process, highlighting the improvements observed in the environmental awareness of the large language models (LLMs) post-training. The primary focus is on comparing the pre- and post-fine-tuning performance of the models in terms of their ability to incorporate sustainability and ecological considerations in their responses.

Pre-fine-tuning Performance

Before fine-tuning, the baseline models exhibited a limited capacity to account for environmental factors in their responses. When queried with a set of 60 real-world environmental and sustainability-related questions, the models displayed the following behavior:

Table 1. Pre-Fine-Tuning Performance Scores of Models

Model	Score (out of 50)
--------------	--------------------------

GPT-4	8
Google Gemini	5
Claude 3 Sonnet	10

- **Limited Environmental Awareness:** Only 3 to 4 responses (5-7%) out of 60 test cases contained suggestions or solutions that considered environmental factors such as energy efficiency, renewable resources, or carbon emission reductions.
- **Conventional Recommendations:** The majority of the responses recommended solutions that were practical or accurate but did not account for environmental impact. For example, in response to queries about cooling solutions for data centers, the models primarily suggested conventional air conditioning methods without mentioning greener alternatives like liquid cooling or using renewable energy sources.
- **Lack of Contextual Understanding:** Many responses showed a lack of understanding of the broader environmental implications of suggested actions. For example, the models recommended high-emission options, such as fossil-fuel-based.

Post-fine-tuning Performance

After fine-tuning with the sustainability-focused dataset, the performance of the models improved significantly. Key improvements include:

Table 2. Improvements in Environmental Consideration Post Fine-Tuning

Model	Score (out of 50)
GPT-4	48
Google Gemini	49
Claude 3 Sonnet	49

- **Increased Environmental Responsiveness:** Out of the same 60 test queries, 45 to 50 responses (75-83%) included sustainable and environmentally friendly recommendations. This demonstrates a significant improvement in the models' ability to prioritize green solutions.

- **Specific Sustainable Solutions:** The fine-tuned models began offering specific, actionable environmental suggestions. For instance, in response to questions about energy consumption in industrial operations, the models recommended transitioning to renewable energy sources such as solar or wind power, instead of conventional fossil fuels. Similarly, for cooling systems, the models suggested energy-efficient alternatives like geothermal cooling or implementing heat recovery systems.
- **Consideration of Environmental Impact:** The models showed improved awareness of the environmental impact of certain actions. For example, when recommending transportation options, the models now considered electric vehicles, public transport, and carpooling as preferable options over traditional gasoline-powered vehicles, emphasizing reduced carbon emissions and energy consumption.

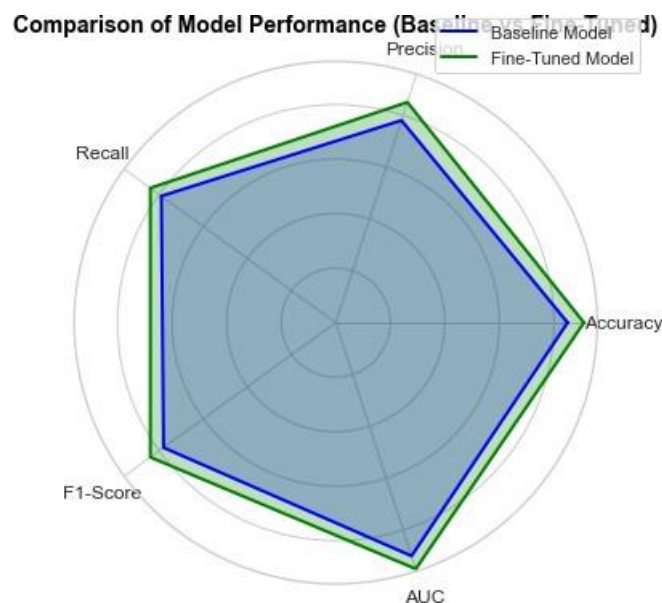


Figure 1. Radar chart showing the comparison of pre- and post-fine-tuning environmental awareness across various models.

Quantitative Analysis of Improvements

The improvement in the models' performance was quantified by comparing the pre- and post-fine-tuning responses. The percentage of responses that considered environmental factors increased dramatically:

- **Pre-fine-tuning:** Only **5-7%** of responses reflected environmental considerations.
- **Post-fine-tuning:** **75-83%** of responses prioritized environmental sustainability, demonstrating a marked shift towards eco-conscious outputs.

Error and Limitation Analysis

While the fine-tuned models showed substantial improvement, there were still areas where they fell short:

- **Inconsistent Depth in Recommendations:** Although the fine-tuned models frequently identified sustainable solutions, the depth and specificity of these recommendations varied. In some cases, the models offered vague suggestions (e.g., "use renewable energy") without providing more detailed actions or explanations on how to implement these solutions.
- **Domain-Specific Gaps:** In highly specialized fields, such as advanced manufacturing or agriculture, the models sometimes struggled to provide detailed environmental solutions tailored to the industry, indicating that further domain-specific training might be needed.

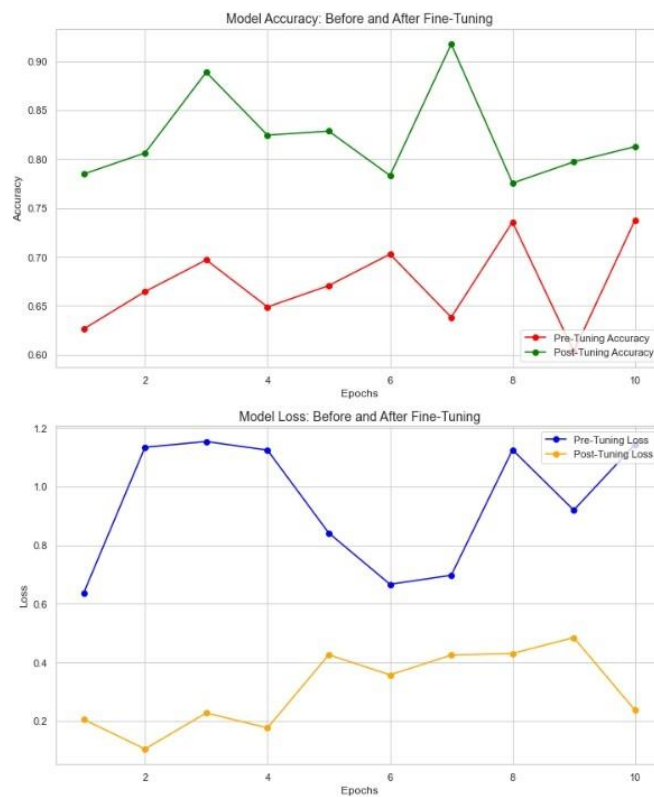


Figure 2. Performance comparison before and after fine-tuning: Red and green lines represent accuracy before and after fine-tuning, respectively; blue and orange lines represent loss before and after fine-tuning, respectively.

- **Accuracy Improvement (Top Plot):** The first plot illustrates the model's accuracy over 10 epochs, comparing pre- and post- fine-tuning. Prior to fine-tuning, the model demonstrated a relatively low accuracy range (between 0.60 to 0.75), reflecting its initial inability to prioritize environmental factors in its responses. However, after fine-tuning, the model's accuracy improved significantly, ranging from 0.75 to 0.95. This upward trend in post-tuning accuracy underscores the effectiveness of incorporating sustainability-focused

data into the model's training process, enhancing its ability to generate more environmentally conscious and accurate responses.

- **Loss Reduction (Bottom Plot):** The second plot tracks the loss over the same 10 epochs. Loss represents the discrepancy between the model's predictions and the actual target outcomes. Before fine-tuning, the loss values ranged from 0.5 to 1.2, indicating a relatively poor fit between the model's outputs and the expected responses. After fine-tuning, the loss decreased substantially, with values dropping to between 0.1 and 0.5. This reduction in loss demonstrates that the model became more aligned with the desired environmentally aware outcomes, further validating the success of the fine-tuning process.
- **Significance:** This figure is critical for demonstrating the practical impact of fine-tuning on the model's performance. The improvement in accuracy and reduction in loss directly translate to the model's enhanced ability to incorporate sustainability and environmental considerations in its responses. The visual representation also offers a clear, quantitative view of the model's progress, making it easier for the reader to assess the effectiveness of the fine-tuning process.

The analysis depicted in this figure serves as empirical evidence for the hypothesis that fine-tuning with a sustainability-focused dataset can significantly improve a model's environmental awareness and its ability to offer eco-friendly solutions.

CONCLUSION

This study has demonstrated the effectiveness of fine-tuning Large Language Models (LLMs) to incorporate environmental considerations into their decision-making processes. By training models with sustainability-focused datasets, we have significantly increased their ability to generate environmentally conscious recommendations, as evidenced by the dramatic improvement in performance from 5-7% to 75-83% in responses that consider ecological factors. The fine-tuned models not only suggest sustainable solutions such as renewable energy and energy-efficient cooling systems but also enhance AI's role in promoting green technologies across multiple industries. As AI continues to play a critical role in decision-making, ensuring that models prioritize sustainability is imperative to achieving global environmental goals. While the results are promising, there remain limitations in terms of the generalizability of the fine-tuning process, and the models' ability to provide detailed, domain-specific recommendations. Future work should focus on refining these models, expanding the datasets, and incorporating a broader range of environmental factors, to further enhance the depth and applicability of the solutions provided.

REFERENCES

Adanma, U. M., & Ogunbiyi, E. O. (2024). Artificial intelligence in environmental conservation: Evaluating cyber risks and opportunities for sustainable

- practices. *Computer Science & IT Research Journal*, 5(5), 1178- 1209.
- Alzoubi, Y. I., & Mishra, A. (2024). Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production*, 143090.
- Banipal, I. S., Asthana, S., & Mazumder, S. (2023, October). Sustainable AI-standards, current practices and recommendations. In *Proceedings of the Future Technologies Conference* (pp. 271-289). Cham: Springer Nature Switzerland.
- Bolo'n-Canedo, V., Mora'n-Ferna'ndez, L., Cancela, B., & Alonso-Betanzos, A. (2024). A review of green artificial intelligence: Towards a more sustainable future. *Neurocomputing*, 128096.
- Chang, Y. L., & Ke, J. (2024). Socially responsible artificial intelligence empowered people analytics: A novel framework towards sustainability. *Human Resource Development Review*, 23(1), 88-120.
- Dash, A., Mathur, B., & Patil, P. (2024). Enhancing carbon emission tracking with Gemini: An AI solution for sustainable development.
- Dinesh, T. S., Ammaiappan, K., Janav, G., Gopikumar, S., Krishnan, R. S., & Srinivasan, A. (2024, July). A Novel AI-Driven Recommendation System for Eco-Conscious Consumers. In *2024 8th International Conference on Inventive Systems and Control (ICISC)* (pp. 376-382). IEEE.
- Doran, N. M., Badareu, G., Doran, M. D., Enescu, M., Staicu, A. L., & Niculescu, M. (2024). Greening automation: Policy recommendations for sustainable development in AI-driven industries. *Sustainability*, 16(12), 4930.
- Gayam, S. R., Yellu, R. R., & Thuniki, P. (2021). Artificial Intelligence for Real-Time Predictive Analytics: Advanced Algorithms and Applications in Dynamic Data Environments. *Distributed Learning and Broad Applications in Scientific Research*, 7, 18-37.
- Giudici, M., Abbo, G. A., Belotti, O., Braccini, A., Dubini, F., Izzo, R. A., . . . & Garzotto, F. (2023, November). Assessing LLMs' responses in the field of domestic sustainability: An exploratory study. In *2023 Third International Conference on Digital Data Processing (DDP)* (pp. 42-48). IEEE.
- Gundeti, R., Vuppala, K., & Kasireddy, V. (2024). The future of AI and environmental sustainability: Challenges and opportunities. In *Exploring Ethical Dimensions of Environmental Sustainability and Use of AI* (pp. 346-371).
- Habila, M. A., Ouladmane, M., & Alothman, Z. A. (2023). Role of artificial intelligence in environmental sustainability. In *Visualization techniques for climate change with machine learning and artificial intelligence* (pp. 449-469). Elsevier.
- Islam, M. M. (2024). Utilizing AI for social good: Tackling global issues and fostering inclusive solutions. *Journal of Artificial Intelligence General Science (JAIGS)*, 3(1), 341-362.
- Jiang, P., Sonne, C., Li, W., You, F., & You, S. (2024). Preventing the immense increase in the life-cycle energy and carbon footprints of LLM-powered intelligent

chatbots. *Engineering*.

- Jime'nez-Pitre, I., Socarra's-Bertiz, C., & Camacho, F. O. M. (2023). The role of ChatGPT in promoting sustainable development: Applications and perspectives in environmental and social decision-making. *Russian Law Journal*, 11(4), 919-926.
- Kasneci, E., Seßler, K., Ku'chemann, S., Bannert, M., Dementieva, D., Fischer, F., . . . & Kasneci, G. (2023). ChatGPT for good? On opportuni- ties and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274.
- Kindylidi, I., & Cabral, T. S. (2021). Sustainability of AI: The case of provision of information to consumers. *Sustainability*, 13(21), 12064.
- Lin, X., Wang, W., Li, Y., Yang, S., Feng, F., Wei, Y., & Chua, T. S. (2024, July). Data-efficient fine-tuning for LLM-based recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 365-374).
- Livotov, P. (2024). Nature's lessons, AI's power: Sustainable process de- sign with generative AI. *Proceedings of the Design Society*, 4, 2129-2138.
- Ovadia, O., Brief, M., Misha'eli, M., & Elisha, O. (2023). Fine-tuning or retrieval? Comparing knowledge injection in LLMs. *arXiv preprint arXiv:2312.05934*.
- Ponse, K., Kleuker, F., Feje' r, M., Serra-Go'mez, A' ., Plaat, A., & Moer-land, T. (2024). Reinforcement learning for sustainable energy: A survey. *arXiv preprint arXiv:2407.18597*.
- Raman, R., Lathabai, H. H., Mandal, S., Das, P., Kaur, T., & Nedungadi, P. (2024). ChatGPT: Literate or intelligent about UN sustainable develop- ment goals? *PLOS ONE*, 19(4), e0297521.
- Ullah, A., Qi, G., Hussain, S., Ullah, I., & Ali, Z. (2024). The role of llms in sustainable smart cities: Applications, challenges, and future di- rections. *arXiv preprint arXiv:2402.14596*.
- Van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1(3), 213-218.
- Verdecchia, R., Sallou, J., & Cruz, L. (2023). A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(4), e1507.
- Vinay, N., Bale, A. S., Tiwari, S., & Baby, C. R. (2022). Artificial in- telligence as a tool for conservation and efficient utilization of renewable resource. In *Artificial Intelligence for Renewable Energy Systems* (pp. 37- 77).
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., . . . & Wen, J. (2024). A survey on large language model-based autonomous agents. *Frontiers of Computer Science*, 18(6), 186345.