
BENEFITS AND LIMITATIONS OF USING GENAI FOR POLITICAL EDUCATION AND MUNICIPAL ELECTIONS

Raphael Fischer
Lamarr Institute for ML and AI
TU Dortmund University
Dortmund, Germany
raphael.fischer@udo.edu

Youssef Abdelrahim
TU Dortmund University
Dortmund, Germany
youssef.abdelrahim@udo.edu

Katharina Poitz
Lamarr Institute for ML and AI
TU Dortmund University
Dortmund, Germany
katharina.poitz@udo.edu

ABSTRACT

Generative artificial intelligence (GenAI) presents both challenges and opportunities across all areas of education. Facing the municipal elections in North Rhine-Westphalia, the *Young AI Leaders* in Dortmund asked themselves—could GenAI be used to make political programs more accessible, in order to facilitate political education? To explore respective potentials and limitations, we therefore performed an experimental study that combines different GenAI approaches. Language models were used to automatically translate and analyze the contents of each program, deriving five potential visual appearance changes to the city of Dortmund. Based on each analysis, we then generated images with diffusion models and published all results as an interactive webpage. All GenAI models were locally deployed on a Dortmund-based computing cluster, allowing us to also investigate environmental impacts. This manuscript explores the project in full depth, discussing technical details and critically reflecting on the results. As part of the global Young AI Leaders Community, our work promotes the Sustainable Development Goal *Quality Education* (SDG 4) by transparently discussing the pros and cons of using GenAI for education and political agendas.

Keywords Generative AI · AI for Education · AI for Politics · Large Language Models · AI for Good · Trustworthy AI

1 Introduction

Modern generative artificial intelligence (GenAI) services are available to anyone with internet access, and thus play a critical role for all areas of education [1]. In the context of political activities, large language models (LLMs) can for example be utilized to build interactive voting advice applications [2] or to visualize political agendas [3]. While this allows for interactive and customizable learning about politics, GenAI availability however also raises ethical questions with regard to misinformation, bias, and misuse [4]. With regard to politics, LLMs could thus potentially be used for “generative propaganda” [5] and moreover have been analyzed to not be neutrally informative in itself [6, 7]. While these works already highlight potentials and limitations, additional empirical studies are needed to understand how AI can truly benefit political and ideological education.

To further investigate the capabilities of GenAI for advancing political education, the Dortmund-based *Young AI Leaders*¹ have decided to launch the *Dortmund-Wahl-KI* (Dortmund-Election-AI) project. Our goal is to analyze and visualize local election programs with the help of GenAI, making the agendas more accessible in order to stimulate voter engagement and discourse. For this project, we developed a custom AI pipeline and published all results in the form of an interactive webpage²—see Figure 1 for an overview. Our work is among the first projects of the worldwide *Young AI Leaders* initiative and promotes the mission of *AI for Good*³, who empower the global young to lead the AI revolution with positive change in mind. In particular, the insights from this project shed light on the role of AI in the context of politics and education, thus relating to *Quality Education* as the fourth Sustainable Development Goal (SDG

¹<https://youngaileaders-dortmund.de/>

²<https://dortmund-wahl-ki.de/>

³<https://aiforgood.itu.int/>




Figure 1: Website showcasing the *Dortmund-Wahl-KI* results.

4) [8]. The manuscript at hand acts as a formal documentation of our project and is structured into four parts—we start with reviewing the literature related to *Dortmund-Wahl-KI*, then explore the details of our methodology, after that present the published results, and lastly self-critically discuss the outcome of our project. Our work follows open science practices, as we provide supplementary materials such as the experiment code and results in our public *GitHub* repository⁴. Moreover, we transparently report on the resource consumption of our AI pipeline, thus promoting the spirit of sustainable reporting [9, 10].

2 Related Work

We first discuss existing literature that discusses how AI can be applied for political education and then discuss respective limitations and risks.

AI can be applied in many different educational areas [1], including but not limited to the context of ideologies and politics. Opportunities are for example seen for conforming to social advancement trends, meeting personalized learning demands, strengthening education efficiency, and facilitating the access to educational resources [11]. To give a specific first example, AI methods can be leveraged to analyze student behavior for providing personalized and interest-oriented learning content [12]. AI can also be utilized through virtual reality or augmented reality, allowing for a brand new way of interactive learning [13]. Similarly, GenAI models can be used to illustrate hypothetical or complex textual information, for example visualizing potential climate disasters [14] or political election programs [3], resulting in higher accessibility. In the context of elections, conversational LLMs are moreover leveraged for offering interactive voting advice applications [2, 6]. Such initiatives result in an alternative representation of political agendas that is more accessible and inclusive than plain text documents, thus connecting to SDG 4 (*Quality Education*) [8].

With the educational benefits in mind, it is however also important to acknowledge associated risks. The potential misuse of AI and LLMs in particular raises ethical concerns [4], as they for example enable “generative propaganda” [5] as already actively adopted by German far-right actors [15]. In addition, GenAI is known to occasionally output factually wrong information, necessitating to establish “hallucination awareness” for students and teachers [16]. AI can also mirror and “even amplify existing biases” [17], for example due to unfair training data or biased algorithmic design choices. Studies have for example evidenced that *ChatGPT*, arguable the most prominent GenAI tool, exhibits biases relating to gender, ethnicity, religion, age, education, and social class [18, 19]. Research suggests that people tend to be more likely to accept and rely on AI when the bias aligns with user opinions [20]. For the political domain, biases of AI are unfortunately more difficult to identify and eliminate [18], thus potentially having a strong influence on people [21]. Several studies have revealed that GenAI models commonly lean toward democratic and left-wing positions [6, 22, 23, 24, 18]. Going into specifics, distinctions can be seen for different models: Answering politically charged prompts, *ChatGPT* and *Claude* were found to exhibit a liberal bias, while *Perplexity* takes more conservative stances and *Gemini* has a centrist view [7]. Another study suggests that *Gemini* might actually be more left-wing and liberal-oriented than *ChatGPT* [24], and moreover showed that the choice of query language can impact the resulting biases. Naturally, understanding such biases is crucial when using GenAI, however a recent study also showcased that “judgmental wording has a stronger impact on truthfulness predictions than political leaning” in the context of automated fact checking [25].

⁴<https://github.com/youngaileadersdortmund/dortmund-wahl-ki>




Figure 2: Schematic visualization of the GenAI pipeline behind *Dortmund-Wahl-KI*.

3 Methodology

To further investigate the capabilities of GenAI for political education, we decided to perform an experimental study in the context of the 2025 municipal elections in Dortmund. Our goal is to promote SDG 4 *Quality Education* by raising awareness about benefits and limitations of using AI for making political content more accessible. Similar to the *bundestagswahl.ai* project [3], we used generative LLMs for analyzing and visualizing political agendas for the Dortmund election. However, due to the aforementioned risks associated with GenAI, it was our goal to implement this project as transparently as possible. As such, instead of relying on black-box AI services like *ChatGPT*, we implemented a specialized AI pipeline that combines different open-weights models and moreover allows for local (i.e., on-premise) deployment. The different steps of the resulting analysis pipeline are summarized in Figure 2 and will be explained in the following.

We started with collecting the election programs of different parties from their respective websites. As a secondary data source, we also analyzed the individual party responses for the *Kommunalomat* (local election compass)⁵, obtained from the *Jugendring (Youth Council) Dortmund*⁶. In the first step of our pipeline, we enabled downstream processing with non-German language models by translating the primary data into English via *TowerInstruct-13B-v0.1*⁷. This model builds on an open-weights multilingual LLM, which was specialized for different tasks related to translation [26]. In our implementation, we split the text into chunks of ten sentences and fed them to the LLM with the following prompt:

Translate the following sentences from German into English: [sentences from the program]

To reduce the complexity of the translated programs, we next compressed them with the help of *BART-Large-CNN*⁸, a fine-tuned denoising autoencoder for text summarization [27]. These program summaries were then analyzed for potential visual changes of the city appearance. The analysis was performed with the *Qwen3-30B-A3B* reasoning model⁹, which is a distilled variant of a powerful foundation model [28]. It uses a mixture-of-experts architecture for efficiency and also supports chain-of-thought inference, which results in more accurate and transparent outputs. For analyzing the program summaries, we used the following prompt:

Identify five important visual aspects of a city appearance that would be affected or impacted by this political program. Describe each aspect in an informative and concise way, with 3 to 6 words. Return these five visual descriptions as a comma-separated list.
[program summary]

As the model provides visual descriptions and chain-of-thought output in English, we once more used *TowerInstruct-13B-v0.1* to translate the reasoning output back to German, in order to better address the local audience. Based on the English descriptions, we lastly used *FLUX.1 [schnell]*¹³ for generating images that visualize the respective changes. The base rectified flow transformer was trained using latent adversarial diffusion distillation and then optimized for fast and accurate image generation [29]. In our pipeline, we used it with the following prompt, generating five image variants per query:

Dortmund city, with additional [visual descriptions]

All AI experiments were performed on the computing cluster of the local *Lamarr Institute*¹⁴, using a single NVIDIA A100 graphics processing unit (GPU). As already displayed in Figure 1, we presented the obtained results to the general

⁵<https://dortmund.waehlt.nrw/>

⁶<https://www.jrdo.de/>

⁷<https://huggingface.co/Unbabel/TowerInstruct-13B-v0.1>

⁸<https://huggingface.co/facebook/bart-large-cnn>

⁹<https://huggingface.co/Qwen/Qwen3-30B-A3B>

¹³<https://huggingface.co/black-forest-labs/FLUX.1-schnell>

¹⁴<https://lamarr-institute.org/>

Table 1: GenAI results based on party programs

Christian Democratic Union ¹⁰	Social Democratic Party ¹¹	Alliance 90/The Greens ¹²
“Revitalized City Center, Lively Public Squares, Green Infrastructure, Safe Streets, Expanded Green Spaces”	“Affordable housing developments, Modern school buildings, Expanded bike lanes, Barrier-free playgrounds, Green infrastructure expansion”	“Increased urban greenery and tree coverage, expanded bike and pedestrian pathways, solar panel installations on buildings, de-paved and permeable surfaces, barrier-free street designs”
		

public via an interactive webpage, powered by Vite¹⁵, React¹⁶, JavaScript, and TailwindCSS¹⁷. The main page features a sliding element that showcases the analysis results for individual parties, as well as a grid of all images that provide detailed information upon user interaction (hover or click). In addition, the website features pages for general information about the project and details on the technical realization. We also placed a disclaimer on every page, clarifying the experimental nature of this project and pointing toward risks and biases of GenAI. We offer German and English languages as well as desktop and mobile support, and deploy the website via GitHub Pages¹⁸. To foster open science practices, we make the complete GenAI pipeline implementation, website code, and experimental results publicly available at <https://github.com/youngaileadersdortmund/dortmund-wahl-ki>.

4 Results

In total, we visualized eight party programs and corresponding election compass responses, resulting in $8 \cdot 2 \cdot 5 = 80$ images. This section is limited to only discussing some selected results, however the rest can be explored online.

The GenAI outputs for three exemplary party programs are displayed in Table 1, featuring the descriptions of potential visual changes obtained from AI reasoning as well as the respectively generated images (two variants each). Comparing

¹⁵<https://vite.dev/>

¹⁶<https://react.dev/>

¹⁷<https://tailwindcss.com/>

¹⁸<https://pages.github.com/>

Table 2: GenAI results for three exemplary parties, based on election compass responses

Christian Democratic Union	Social Democratic Party	Alliance 90/The Greens
“Clean school exteriors, expanded public libraries, increased bike parking facilities, safer cycle paths with markings, more public shade trees”	“Public Learning Spaces, Safer Cycle Paths, Bike Parking Spaces, Tree-lined Public Spaces, Solar Panel Installations”	“Clean and well-maintained school areas, enhanced urban infrastructure, increased greenery and tree planting, improved public hygiene facilities, expanded pedestrian and cycling infrastructure ”
		
		

the textual results, we see a strong overlap—all programs were analyzed to advocate a green transformation, two of them mention expanded bike lanes, and they all seem to promote a modern and lively city. Investigating the images below, some of these aspects seem to be well-represented, for example displaying urban greenery (all parties), a lively city center (CDU), and bike lanes (SPD and the Greens). However, several other aspects cannot be found in the images—solar panels and modern buildings are missing, even though they would be easy to visualize, while characteristics such as safe streets, affordable housing, and permeable surfaces are arguably harder to incorporate in images. In general, the images also seem to be quite similar in terms of daytime (bright daylight) and perspective (especially when comparing the variants per party).

In addition to the election programs, we used the responses of the parties for the local election compass (*Kommunalomat*) as an alternative data source for our experiments. These results are displayed in Table 2, showcasing that the respectively analyzed visual points and resulting images are even more similar among the three selected parties. For all of them, the AI analysis highlights clean school areas and cycling infrastructure, while more intricate details are hardly displayed in the images (libraries for CDU, solar panels for SPD, hygiene facilities for the Greens). In these experiments, the nature of the primary data likely is a major issue—as the questions are the same for all parties, only *approved* or *rejected*, the AI analysis is likely biased toward the question wording itself, not the party response. Recalling the reasoning prompt given in our Methodology, it is also interesting that both the program and election compass descriptions for the Greens is significantly longer than instructed, also in comparison to the other parties. Another aspect of random behavior can be seen in the capitalization, which is not unified across the descriptions (despite using the same base prompt).

Promoting the spirit of sustainable reporting in AI [9, 10], we also estimated the energy consumption and CO₂ emissions of running our AI pipeline. Respective numbers were assessed with the help of *CodeCarbon* [30], as listed in Table 3.

Table 3: GenAI processing steps with corresponding duration, energy use, and emissions.

Processing Step	Duration	Energy Use	Power Consumption	Emissions
Translate (DE → EN)	132.06 min	0.71 kWh	322.58 W	0.27 kg CO ₂ eq
Summarize	5.04 min	0.02 kWh	238.10 W	0.01 kgCO ₂ eq
Reason	360.24 min	1.41 kWh	234.84 W	0.54 kg CO ₂ eq
Translate (EN → DE)	18.96 min	0.10 kWh	316.46 W	0.04 kg CO ₂ eq
Generate Images	22.45 min	0.12 kWh	320.71 W	0.05 kg CO ₂ eq
Total	538.75 min	2.37 kWh	–	0.90 kg CO₂eq

The GenAI models behind the individual processing steps vary in their complexity and resulting power demand. The reasoning analysis had the highest environmental impact, mostly due to long processing times. With these results, note that (1) *CodeCarbon* is known to underestimate the actual energy demand of AI by 20–30% [31] and (2) these results only reflect the environmental impact originating from the energy demand of one analysis run across all parties—as such, the reported numbers do not entail overhead impacts of project development, AI model life cycle aspects (e.g., training) [32], or embodied impacts of the utilized hardware [33].

5 Discussion

We want to close our work with a self-critical discussion of our project, summarizing the results and linking them to established literature. As stated in our Related Work section, we are aware that LLMs and GenAI models exhibit biases [6, 22] and potentially do not correctly capture the political agendas [16]. At the same time, we face the fact that political actors already use GenAI for pushing their agendas [5, 15] and that young people might include AI models in their (self-)education [2, 3], despite the risks. To us, this highlights the need for additional empirical studies and motivated us to conduct our project. However, we strived for mitigating risks and maximizing transparency by following open science practices, using a reasoning model for the analysis, and prominently featuring a disclaimer on general problems of GenAI on the webpage.

Reflecting on our results, we believe that there are clear benefits but also strong limitations of using GenAI for political education. Our project was positively received and sparked many interesting discussions around the municipal elections and political usage of AI, which otherwise would not have taken place. On this personal scope, the feedback demonstrated how experimental projects can stimulate political education and discourse. As a big benefit, GenAI capabilities for automated summarization and visualization can facilitate the exchange and learning about political agendas without being required to read extensive documents. However, we also found limitations because (1) many aspects were not correctly represented in the images, (2) the visualizations generally look very similar, and (3) we found a strong gap when comparing the analysis outcomes of the programs and election compass answers. A critical finding thus is the limited ability of our AI pipeline to capture subtle nuances and distinct priorities among different political agendas. Instead, the results follow general trends that makes it hard to find differences among the actors. Interestingly, the results of the closely related *bundestagswahl.ai* project appear very different across all parties [3]—unfortunately, due to the black-box nature of their utilized AI service (*ChatGPT*), it is hard to understand why the results are so different. One root cause for the diverging project outcomes might be that election programs are more similar on a municipal level than on the federal level, resulting in more similar AI visualizations. Another option could be that *ChatGPT* more explicitly focuses on (and potentially even exaggerates) the characteristics of the individual parties—it is unknown how much it actually pays attention to the given program, or just repeats biases from online reputation. In short, our findings demonstrate that GenAI can serve political education, but it is extremely important to do so in a transparent, trustworthy, and responsible manner.

From these conclusions arise several opportunities for future research work. On the technical site, one could aim at obtaining less similar results by developing an AI pipeline that focuses on the differences among political agendas, instead of conducting the analysis and visualization of each primary data source individually. The significant limitations of the image generation model used in our work warrants further investigation and trying to find solutions that more accurately visualize the given aspects. In addition to improving our proposed methodology with regard to model selection and prompting, it could also applied to other political contexts and electoral analyses. Interdisciplinary research could focus on evaluating human feedback on projects like ours in order to better characterize and potentially

alleviate the limitations of GenAI for political education. For the future, we also plan to collect and characterize the various opportunities for interdisciplinary work in the field via expert consulting.

6 Conclusion

The availability of GenAI necessitates to carefully investigate modeling capabilities in the context of education and politics. Our project not only provided an overview of work in the field, but also is the first to analyze and visualize municipal election programs with LLMs and diffusion models. The results highlight the potential for making political agendas more accessible, however also showcase the limitations for grasping intricate details and differences among parties. With our work, we hope to make a helpful contribution to research on using AI for politics and education, while also promoting the sustainable use of AI, for good.

References

- [1] Fati Tahiru. AI in Education: A Systematic Literature Review. *Journal of Cases on Information Technology (JCIT)*, 23(1):1–20, 2021. ISSN 1548-7717. doi: 10.4018/JCIT.2021010101.
- [2] Martin Schiele, Yannick Gittmann, Stefan Ilchmann, Ante Gojsalić, Dominik Jurinčić, and Phylis Klempt. Voting advice applications: Implementation of rag-supported llms. July 2024. doi: 10.36227/techrxiv.172115156.64500701/v1.
- [3] Max Mundhenke. bundestagswahl.ai. <https://bundestagswahl.ai/>, 2025. Last Accessed: 2025/10/10.
- [4] Sara Saylam, Nilgun Duman, Yetkin Yildirim, and Kseniya Satseviv. Empowering education with ai: Addressing ethical concerns. *London Journal of Social Sciences*, (6):39–48, Sep. 2023. doi: 10.31039/ljss.2023.6.103.
- [5] Madeleine I. G. Daepp, Alejandro Cuevas, Robert Osazuwa Ness, Vickie Yu-Ping Wang, Bharat Kumar Nayak, et al. Generative propaganda. 2025. doi: 10.48550/arXiv.2509.19147.
- [6] Ina Dormuth, Sven Franke, Marlies Hafer, Tim Katzke, Alexander Marx, et al. A cautionary tale about "neutrally" informative ai tools ahead of the 2025 federal elections in germany. 2025. doi: 10.48550/arXiv.2502.15568.
- [7] Tavishi Choudhary. Political bias in large language models: A comparative analysis of chatgpt-4, perplexity, google gemini, and claude. *IEEE Access*, 13:11341–11379, 2025. doi: 10.1109/ACCESS.2024.3523764.
- [8] Ellen Boeren. Understanding Sustainable Development Goal (SDG) 4 on “quality education” from micro, meso and macro perspectives. *International Review of Education*, 65(2):277–294, April 2019. ISSN 1573-0638. doi: 10.1007/s11159-019-09772-7.
- [9] Raphael Fischer, Thomas Liebig, and Katharina Morik. Towards more sustainable and trustworthy reporting in machine learning. *Data Mining and Knowledge Discovery*, 2024. ISSN 1573-756X. doi: 10.1007/s10618-024-01020-3.
- [10] Raphael Fischer. *Advancing the Sustainability of Machine Learning and Artificial Intelligence via Labeling and Meta-Learning*. PhD thesis, TU Dortmund University, 2025.
- [11] Hongmin Fan and Xiaomin Wang. The role of artificial intelligence in ideological and political education: Opportunities and challenges. In *Proceedings of the International Conference on Information Education and Artificial Intelligence (ICIEAI)*, page 815–820, 2024. ISBN 9798400716157. doi: 10.1145/3660043.3660188.
- [12] Huang Aiwa and Tao Hongwei. Model construction and application of artificial intelligence in improving the efficiency of ideological and political education. In *International Conference on Artificial Intelligence, Deep Learning and Neural Networks (AIDLNN)*, pages 126–131, 2024. doi: 10.1109/AIDLNN65358.2024.00028.
- [13] Jiang Shan. Opportunities and challenges of ideological and political education in the era of artificial intelligence. *International Journal of Social Sciences in Universities*, 25, 2024. URL <http://www.acadpubl.com/ijssu.aspx>.
- [14] Victor Schmidt, Alexandra Luccioni, Mélisande Teng, Tianyu Zhang, Alexia Reynaud, et al. ClimateGAN: Raising climate change awareness by generating images of floods. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=EZN0b_uNpJk.
- [15] Anna Hiller and Pablo Maristany de las Casas. Generative AI and the german far right: Narratives, tactics and digital strategies, 2025. URL <https://www.isdglobal.org/wp-content/uploads/2025/02/The-use-of-generative-AI-by-the-German-Far-Right.pdf>. Last Accessed: 2025/10/10.
- [16] Bogdan-Iulian Ciubotaru. The hallucination problem in generative artificial intelligence: Accuracy and trust in digital learning. In *20th International Conference on Virtual Learning (ICVL)*, pages 35–45, 04 2025. doi: 10.58503/icvl-v20y202503.

- [17] Emilio Ferrara. Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. *Sci*, 6(1), 2024. ISSN 2413-4155. doi: 10.3390/sci6010003.
- [18] Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. More human than human: measuring ChatGPT political bias. 198(1):3–23. ISSN 1573-7101. doi: 10.1007/s11127-023-01097-2.
- [19] Yao Qu and Jue Wang. Performance and biases of large language models in public opinion simulation. 11(1): 1095. ISSN 2662-9992. doi: 10.1057/s41599-024-03609-x.
- [20] Uwe Messer. How do people react to political bias in generative artificial intelligence (AI)? 3:100108. ISSN 2949-8821. doi: 10.1016/j.chbah.2024.100108.
- [21] Uwe Peters. Algorithmic political bias in artificial intelligence systems. 35(2):25. ISSN 2210-5441. doi: 10.1007/s13347-022-00512-8.
- [22] Jérôme Rutinowski, Sven Franke, Jan Endendyk, Ina Dormuth, Moritz Roidl, and Markus Pauly. The self-perception and political biases of chatgpt. *Human Behavior and Emerging Technologies*, 2024(1):7115633, 2024. URL <https://onlinelibrary.wiley.com/doi/abs/10.1155/2024/7115633>.
- [23] David Rozado. The political biases of ChatGPT. *Social Sciences*, 12(3), 2023. ISSN 2076-0760. doi: 10.3390/socsci12030148.
- [24] Dogus Yuksel, Mehmet Cem Catalbas, and Bora Oc. Language-dependent political bias in ai: A study of chatgpt and gemini. 2025. doi: 10.48550/arXiv.2504.06436.
- [25] Charlott Jakob, David Harbecke, Patrick Parschan, Pia Wenzel Neves, and Vera Schmitt. Polbix: Detecting llms' political bias in fact-checking through x-phemisms. 2025. doi: 10.48550/arXiv.2509.15335.
- [26] Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, et al. Tower: An open multilingual large language model for translation-related tasks. 2024. doi: 10.48550/arXiv.2402.17733.
- [27] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, et al. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. 2019. doi: 10.48550/arXiv.1910.13461.
- [28] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, et al. Qwen3 technical report. 2025. doi: 10.48550/arXiv.2505.09388.
- [29] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, et al. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space. 2025. doi: 10.48550/arXiv.2506.15742.
- [30] Benoit Courty, Victor Schmidt, Sasha Luccioni, Goyal-Kamal, Marion Coutarel, et al. mlco2/codecarbon: v2.4.1. 2024. doi: 10.5281/zenodo.11171501.
- [31] Raphael Fischer. Ground-truthing AI energy consumption: Validating CodeCarbon against external measurements. 2025. doi: 10.48550/arXiv.2509.22092.
- [32] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, et al. Sustainable AI: Environmental implications, challenges and opportunities. 2022. doi: 10.48550/arXiv.2111.00364.
- [33] Sophia Falk, David Ekchajzer, Thibault Pirson, Etienne Lees-Perasso, Augustin Wattiez, et al. More than carbon: Cradle-to-grave environmental impacts of GenAI training on the Nvidia A100 GPU. 2025. doi: 10.48550/arXiv.2509.00093.