GreenChat
Group 12

Simon Biennier Jasper Heijne Paul Lindhorst Huib Sprangers

April 4, 2025

Abstract

As artificial intelligence (AI) becomes increasingly integrated into daily activities, its environmental impact often remains unnoticed. Large language models (LLMs) like ChatGPT consume significant computational resources, leading to high energy usage, carbon emissions, and water consumption. While users are growing, there is limited transparency regarding the resource demands of these systems, leaving end users unaware of their digital footprint.

This paper introduces **GreenChat**, a browser extension designed to provide real-time feedback on the environmental impact of ChatGPT conversations. By estimating energy consumption, carbon emissions and water usage, the tool aims to increase user awareness and encourage more sustainable AI usage. We present the design and implementation of the extension, the methodology behind its estimations, and a case study assessing its effectiveness in influencing user behaviour. Our findings highlight the challenges of accurately quantifying AI's environmental impact and explore the broader implications of integrating sustainability metrics into user-facing applications.

A Introduction

The increasing reliance on AI in everyday tasks has brought significant advancements in efficiency and accessibility. However, the environmental cost of these AI-driven services often goes unnoticed. This becomes particularly important given that AI is likely to grow more and more.

With the increasing use of LLMs, such as ChatGPT, there has been a corresponding increase in electricity consumption along with carbon emissions [14]. This is directly related to the intensive computations and energy requirements performed by Graphics Processing Units (GPUs), necessary for AI model operations, as well as the energy requirements of data centers to host these models. Not to mention the costs of training these models in the first place, which have been likened to the emission of five cars over their full lifespan, resulting in approximately 300 tons of CO_2 [13].

Despite the increasing attention given to sustainable computing, most users remain unaware of the energy footprint associated with AI interactions. This is emphasised by the difficulty in measuring these consequences [3]. Unlike traditional computing tasks, AI model usage occurs in cloud-based infrastructures, making its energy consumption less tangible. Without transparency, users lack the necessary information to make environmentally conscious decisions about their AI usage.

Moreover, with ChatGPT handling over 300 million weekly active users and 1 billion user messages every day, its energy consumption is projected to reach approximately 226.8 GWh per year, which could "run the entire country of Finland or Belgium for a day" [8, 5, 11].

What's more, the answers to many questions asked to LLMs can easily be found using a simple preferred search engine query. This is a more energy-efficient method than asking LLMs: when you submit a query to ChatGPT, it passes your input through the model's parameters to identify patterns and generate a response based on its training data. This is less energy efficient than making a query to a domain name server, as search engines do. As stated by Hoffman: "Each time you ask ChatGPT a question, it uses about 0.0029 kilowatt-hours of electricity. This is nearly ten times more than the energy needed for a typical Google search, which consumes about 0.0003 kilowatt-hours per query, according to The Electric Power Research Institute (EPRI)." [4, 2]

To address these issues, we have developed **GreenChat**, a browser extension that provides real-time feedback on the environmental impact of ChatGPT conversations. By estimating energy consumption, carbon emissions and water usage, the tool aims to raise awareness and encourage users to reflect on their AI consumption. This paper presents the development of our extension, the methodology used to estimate energy impact, and a case study evaluating its effectiveness in raising awareness and influencing behaviour. We also discuss the challenges of accurate estimation and the broader implications of making AI sustainability metrics accessible to users.

A.1 Related work

Various other projects have been carried out to measure the emissions made by ChatGPT or make people more aware of the environmental impact they might have. One among these is Tilburg AI's CodeCarbon project [15], which allows a user to track their emission with a Python package, though it only tracks the emissions of programs run while CodeCarbon is running itself and needs some experience with coding to use. We believe our project can improve upon this by directly implementing an extension within ChatGPT to make be more user-friendly.

Another project that aims to bring awareness to the environmental impact of large language models would be Maria the emissions reduction expert [1], a GPT made to answer any questions people might have related to the environmental impact of LLMs. This is a great tool to help raise awareness, but we see a downside of using an AI model to teach AI sustainability, as this causes undue emissions by the emissions reduction AI, and might hallucinate to give untruthful answers.

In addition to projects such as these, raising awareness about the environmental impacts of AI appears to be done mainly through blog posts and articles [6, 17]. Which is less intuitive for the public than creating an extension directly integrated with large language models. Our project aims to accurately convey the impact from using LLMs on the environment to its users. We feel our work is valuable for providing a tool that builds off previous work, to help educate the LLM users intuitively on the carbon emissions and water usage needed to power ChatGPT.

B Our solution

We aim to spread awareness among ChatGPT users and inform them of the high energy usage of LLMs. We achieve this by creating a browser extension, showing the users how much energy ChatGPT consumed in the current "chat" (thread of multiple prompts and responses). We create this extension for Chrome, since it is the most widely used browser worldwide [12]. The energy consumption information is seamlessly integrated with ChatGPT's user interface. The interface of a ChatGPT-chat con-

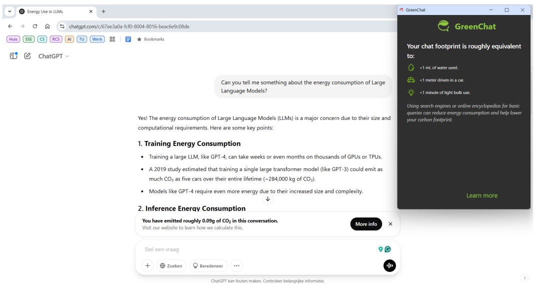

Figure 1: The interface of GreenChat, the ChatGPT extension.

sists of a query field at the bottom of the screen. The queries and responses fill the rest of the screen. The extension presents the current chat's carbon emission in a banner just above the query field for conversations.

The extension provides data on the energy usage of the chat in a way that is intuitively understandable for users. This means that besides the computed carbon emission, we report a comparison to more intuitive measurements (e.g. driving a car for a certain distance or having a light bulb turned on for some time). This information is not provided in the banner during the chat to avoid cluttering the user interface, but can be requested by clicking "More info". This information will then be provided in a popup. In this popup, we also advise users on greener options to retrieve the answers to their queries.

See Figure 1 for the interface of the GreenChat extension. At the bottom of the screen, above the query field, the extension shows the banner with carbon emissions. This banner can be closed with the close button. A small circle with the extension's logo will then appear at the bottom right of the screen to reopen the banner. When the "More info" button in the banner is pressed, the popup at the top right appears. The popup can be closed by pressing the cross button, in case the user is currently not interested in their energy usage. The "Learn more" field in the popup will direct the user to the extension's website, where we elaborate on our carbon calculations and provide more information about the extension.

C Implementation

This section explains how we implemented our extension. The code is publicly available on GitHub[1].

[1] https://github.com/simonbiennier/greenchat

Then, we elaborate on our calculations for the estimation of ChatGPT's carbon and water usage.

C.1 Code

The browser extension was developed using ReactJS [7], enabling us to inject custom HTML elements directly into the ChatGPT page to display real-time carbon emission data depending on the number of tokens. Additionally, we included a button that launches a popup window with more information, as described in Section B.

C.2 Carbon calculations

Due to the closed source nature of OpenAI's policy regarding ChatGPT-4 and its derivatives, it is not viable for anyone outside of the company to accurately predict how much carbon their tool emits. However, we can make educated guesses, considering available information.

Firstly, we need to find out how many watt-hours are used by a query, depending on the number of tokens. In a study conducted by Alex De Vries [16], this amount would be around 3 watt-hours for a typical query. However, Josh You writing for Epoch AI [18] claims that this is an overly pessimistic number, citing that while longer queries would result in a similar energy cost of about 2.4 watt-hours, a truly typical ChatGPT query would be far shorter than presumed by De Vries, and only result in an energy cost of around 0.3 watt-hours, which would be about on par with a simple Google search according to De Vries.

The main cause for the discrepancy between these two sources comes from several assumptions they make: a disagreement on what would be considered a typical query's token length, differing assumptions on which type of GPU was used to train and run ChatGPT-4 and differing assumptions on the parameter size of the model. Since none of these factors have been officially revealed by OpenAI, it is not possible to objectively say which estimation is more precise.

The second factor determining the emission of carbon is the carbon intensity of the places where OpenAI keeps its data centers. Once again, this is not publicly disclosed, though it is known that these data centers are located within the United States, which has a national average of about 0.429 $kg\ CO_2e/kWh$ [9]. This means that the emissions would lie between around $0.429 \cdot 0.3 = 0.1287$ $g\ CO_2e$ and about $0.429 \cdot 3 = 1.287$ $g\ CO_2e$. The average carbon intensity is calculated by taking the average carbon intensity of each state, meaning our estimations are rough, since OpenAI's data centers are unlikely to be evenly spaced throughout the country.

Lastly, we elaborate on the token count: per which number of tokens should we rule this estimate? The optimistic estimate by Mr.You [18] calculates the given number of 0.3 watt-hours using only the assumption of 500 tokens of output, considering the input tokens for such a query to be negligible. Meanwhile, the research by De Vries [16] pessimistically presumes an input of 4000 tokens with 2000 tokens as output. This means that the carbon emission range will be calculated with the range of $0.1287/500 = 0.2574$ $mg\ CO_2e$ per token and $1.287/6000 = 0.2145$ $mg\ CO_2e$ per token. For our extension, we then take the average of these two numbers.

C.3 Water Usage

Much like with the carbon usage, the lack of clarity from OpenAI regarding their internal statistics, the efficiency of their program and any information regarding their data centers, we cannot claim to objectively measure the water usage of ChatGPT. In work done by The Green Grid[10] the average water usage efficiency (WUE) of the United States is stated to be around 1.8 L/kWh, which means that in accordance with the earlier estimates given for the number of watt-hours used by ChatGPT, we assume that the amount of water used lies between $0.3 * 1.8/500 = 0.000108mL$ and $3 * 1.8/6000 = 0.0009mL$ per token.

Like with the carbon emissions, these numbers are averaged for our extension to derive the final value per token calculation. As with the carbon intensity, the WUE of the United States is calculated by taking the average across the United States, and is therefore a rough estimate for the true amount of water used.

D Evaluation

We conduct a case study to evaluate whether our proposed solution will create awareness among ChatGPT users. We interview multiple participants and let them interact with the extension. Based on these results, we can evaluate if the extension achieves its purpose of spreading awareness and what possible changes are required to improve the extension.

D.1 Setup

Prior to the case study, we ask the prospective participant for their consent to collect data on their knowledge and opinion on software to improve AI sustainability. When asking for consent, we will describe the steps listed in this section to inform participants of what they are participating in. We also inform participants that the data obtained is only used for academic purposes and will be kept private at all times.

To start the case study we ask our participants questions to poll their current knowledge on the energy usage of LLMs. In this phase they will not have any information from us or from the extension, so we obtain a baseline measurement of their knowledge. Questions we will ask them in this phase are:

- What method do you think is the most energy efficient for translating one sentence to another language?
    1. Use a website to translate (i.e. Google Translate).
    2. Ask a specialized small LLM to translate.
    3. Ask ChatGPT to translate.
- Given that a light bulb consumes approximately 60 watt-hours, how much energy do you think 10 prompts to ChatGPT cost?
    - Answer in decimal value of watt-hours.
- How much water do you think a chat of 10 prompts to ChatGPT costs?
    - Answer in decimal value of liters.

In the next phase, we explain that ChatGPT and LLMs in general are large consumers of energy. We explain that the idea of this extension is to raise awareness among ChatGPT users about the energy they are using with the chat. We follow this with the following questions about the participant's expectations of such an extension:

- Would you be interested in an extension for ChatGPT to make you more aware of your energy consumption?
    - Answer with yes/no and optionally a brief explanation.
- Do you think your behaviour with ChatGPT would change when you are confronted with your energy usage?
    - Answer with yes/no and optionally a brief explanation.

Next, we will let the participants interact with the extension. We will not assist them in this process to see how user-friendly the interface is. During the participant's interaction with the extension, we will monitor them to see what things work intuitively and what parts of the extension are not immediately clear.

After the participant has explored the extension, we show parts of the extension they might have missed in their exploration. Hereby we make sure that the participant has a complete understanding of all functionalities of the extension.

Finally, we end the case study by asking the participants for their opinion on the extension. We do this with the following questions:

- Did you learn new information on AI energy consumption from the extension?
    - Answer with yes/no and optionally a brief explanation.
- Would you change your ChatGPT behaviour based on the information supplied by the extension?
    - Answer with yes/no and optionally a brief explanation.
- Would you use the extension yourself?
    - Answer with yes/no and optionally a brief explanation.
- What is missing in the extension?
    - Answer in the form of a brief explanation of the missing elements.

D.2 Participants

Due to the limited time available, we opt to ask our close friends and family for the extension's evaluation. We select this group as we can get in contact with them in a short amount of time, and they do not require compensation for participating. For the distribution of age groups, see Table 1.

A large portion (66,7%) of the participants has a technical background. This means they are following or have completed a technical study. One of

Age Range	Number of Participants
18 - 25	5
26 - 35	1
36 - 45	0
46 - 55	1
56 - 65	2

Table 1: The age groups of the participants in the case study.

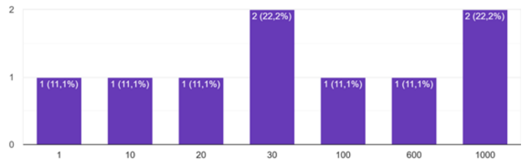

Figure 2: Participants' prediction for watt-hour usage of ChatGPT.

these participants is also employed in a computer science-related field. This is because the pool of participants largely consist of fellow students from the TU Delft. Participants may therefore have more knowledge on AI and energy usage than the average ChatGPT user.

D.3 Results

From the pre-knowledge questions from the case study, we see that 33,3% of participants think that asking an LLM to translate is more energy efficient than using a translation website (i.e. Google Translate). The estimation by participants for the amount of watt-hours ChatGPT uses and the amount of water used, we see very scattered results, see Figure 2 for the results of the watt-hour estimations. The water consumption estimations ranged from 0.01L to 50L of water for 10 ChatGPT prompts.

When asked about possible interest in this extension, a majority of the participants indicated that they would consider using the extension. Some interesting points are mentioned by participants. One indicates that they would only consider using the extension if it is easy to install, and another indicates that the extension should not be obstructive when using ChatGPT. This indicates that accessibility and usability are important for this extension. Most participants also say they would change their behaviour if they were to be confronted with the energy consumption.

After interacting with the extension, all participants indicate that their estimations were off. Participants say their behaviour might slightly change because of the extension. Two participants mentioned that they would mostly change their behaviour for smaller unimportant prompts. Someone also explained why they would not change their behaviour by comparing the small energy usage of ChatGPT to other things they do in their daily life that pollute more.

Overall, participants are mostly positive as the extension is not too obstructive and can be removed when it is no longer wanted. Most participants are interested in the information that the extension supplies and are therefore willing to use it.

For the future of the extension, we received feedback on the user interface that we could incorporate in a future patch of the extension. Improvements mentioned were centering the "learn more" button and matching the popup interface with the ChatGPT and banner interfaces. Since these are minor adjustments that do not add greatly to the cause of raising awareness on AI energy consumption, we choose to exclude these changes from the initial launch of the extension.

E Discussion

The findings from our case study highlight a significant gap in public awareness regarding the energy consumption of LLMs such as ChatGPT. Estimations for energy and water consumption are very scattered, indicating that this information is not widely accessible. This lack of transparency is primarily caused by the secrecy that big AI companies have about the energy consumption of their models. Extensions to raise awareness are a step toward greater openness and accountability.

Although our estimations are not perfect due to the lack of publicly available data, this does not undermine their value. To achieve our goal of raising awareness, rough estimations of emissions are sufficient. These rough estimations can cause ChatGPT users to consider the options they have for their queries and make green behavioural decisions.

For such extensions to be effective, they must be non-intrusive and free from advertisements. ChatGPT and other LLMs are very useful, so the intent of the extension is not to prohibit users from interacting with these LLMs. The case study indicates that people are willing to learn about the energy consumption of LLMs. Learning about the consequences of using ChatGPT does not fix the pollution made by LLMs, but it is a good first step to raise awareness of the problem. Our findings suggest that users may adjust their behaviour, particularly by reducing low-priority or unnecessary queries. Although such small changes may not drastically cut overall energy consumption, they contribute to a movement toward more sustainable AI usage.

That being said, the effectiveness of the extension relies on user adoption. The extension will only spread awareness among people who took the effort to install the extension. This raises a paradox where people have to be somewhat informed on LLM energy consumption to be incentivised to install the extension which aims to make them informed on LLM energy consumption. Therefore, the users who would benefit most from this extension, who are completely unaware of ChatGPT's energy consumption, are unlikely to find and install this extension. To solve this, a direct integration into LLM interfaces displaying awareness information as presented by this extension, without asking users to manually install an extension, would greatly improve the impact on awareness.

Beyond individual action, systemic change is necessary to reduce the impact of LLMs on the environment. The biggest challenge is the lack of transparency from large AI companies regarding their energy consumption. Without this data, it is difficult for both users and policymakers to grasp the scale of LLM emissions. Ideally, large AI companies should openly report their energy usage, allowing for more accurate estimations and informed discussions. However, given that such transparency could harm their reputation, it is unlikely that they will make this information public voluntarily. Therefore, government regulations for openness about AI energy consumption may be an essential step toward increasing public awareness and creating accountability within the industry.

While this extension does not directly solve the issue of AI-related pollution, it plays an important role in raising awareness. This is an important first step to decrease the pollution by large AI companies. Increased awareness can drive further discussions, policy changes, and the development of greener AI initiatives. By acknowledging the environmental impact of LLMs, we can take the necessary steps toward more sustainable AI practices.

E.1 Threats to validity

The biggest ambiguity of the extension is the computation of emitted carbon. In Section C.2, we explain how we came to the current estimation in the extension. This estimation is based on two research papers that make an educated estimate of ChatGPT's energy usage. However, OpenAI has not released the information required to make a precise estimation of ChatGPT's energy consumption. Therefore, we cannot be confident that the provided estimations are precise. However, as the goal of the extension is to create awareness among ChatGPT users, an exact estimation is not a hard requirement. Awareness can also follow from a rough estimation, as we provided.

Besides the estimation of carbon usage we also estimate how much water ChatGPT uses. As OpenAI is also not transparent on this information, as described in C.3, this is also a rough estimate. This estimation is based on the average water usage efficiency, which is also an estimation.

Another complication in predicting the emission of an LLM is the energy used during the training phase of the model. For our energy estimation, we used the energy that the model consumes to stay running. However, a lot of energy is also used in the initial training step of the model. We do not use this energy consumption in our estimation, as training is only done once and will therefore not differ based on the number of requests to ChatGPT. Therefore, we deem this energy consumption irrelevant for the emission of a chat, but awareness of energy consumption of training LLMs is also important.

To compare the emission of ChatGPT to intuitive measurements we have estimated the emission of driving a car and using a light bulb. In contrary to the emission data of large AI companies, this data is available publicly. However, there are many types of cars and light bulbs, each with different energy usage. We use a general mean to estimate these values, as we feel this will provide enough information to raise awareness in an intuitive manner.

Due to time constraints, the evaluation case study is conducted with participants from our personal network, including family and friends. This demographic consists of individuals who are inherently supportive and inclined to assist, which may influence the outcomes of the study. This group of people is likely more supportive and positive about the created extension, possibly causing skewed results.

The demographic for the case study is also not representative of the target audience of the extension. The number of participants is small and they are from similar societal backgrounds as we are. The target audience of the extension is all users of ChatGPT, which is a worldwide, diverse audience. Therefore, it is unlikely that the small group of like-minded individuals is representative of the complete target audience.

F Conclusion

To sum up, we developed an innovative browser extension to increase awareness of carbon emissions associated with ChatGPT use. The extension intuitively shows real-time carbon footprint data with relatable comparisons, and informs users that energy-efficient alternatives for simple queries exist. According to our preliminary assessment, most participants were previously unaware of these emissions but became more mindful after using the tool. Ultimately, this extension promotes more sustainable digital practices among ChatGPT users.

F.1 Future work

To further increase the precision of our carbon emissions tracking, it would be essential to acquire more official and precise energy usage data directly from OpenAI. Furthermore, we intend to offer users the option to redirect simple queries to more energy-efficient alternatives, such as standard Google searches or smaller, specialized LLMs, to reduce unnecessary energy consumption. In the future, these recommendations could be tailored to the queries made by the user. Lastly, we could add gamification components with long-term impact tracking, which could better motivate users to reduce their carbon footprint over longer periods of time.

As previously stated, the evaluation of our tool was very limited due to time constraints, so conducting a more intensive evaluation in the future with a larger and diverse user group would further validate the effectiveness of GreenChat.

References

[1] Lana Aljuaid. *Maria the emissions reduction expert.* 2025. URL: https://chatgpt-plugin.org/custom-gpt/maria-the-emissions-reduction-expert/.

[2] Electric Power Research Institute (EPRI). *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption.* Brochure. Electric Power Research Institute (EPRI), May 28, 2024, p. 35. URL: https://www.epri.com/research/products/000000003002028905.

[3] Ahmad Faiz et al. *LLMCarbon: Modeling the end-to-end Carbon Footprint of Large Language Models.* 2024. arXiv: 2309.14393 [cs.CL]. URL: https://arxiv.org/abs/2309.14393.

[4] Paul Hoffman. "AI's Power Demand: Calculating ChatGPT's Electricity Consumption for Handling Over 365 Billion User Queries Every Year". In: (Jan. 30, 2025). URL: https://www.bestbrokers.com/forex-brokers/ais-power-demand-calculating-chatgpts-electricity-consumption-for-handling-over-78-billion-user-queries-every-year/?utm_source=chatgpt.com.

[5] Dara Kerr. "AI Brings Soaring Emissions for Google and Microsoft, a Major Contributor to Climate Change". In: *NPR* (July 2024). URL: https://www.npr.org/2024/07/12/g-s1-9545/ai-brings-soaring-emissions-for-google-and-microsoft-a-major-contributor-to-climate-change?utm_source=chatgpt.com.

[6] Sophie McLean. *The Environmental Impact of ChatGPT: A Call for Sustainable Practices In AI Development.* 2023. URL: https://earth.org/environmental-impact-chatgpt/.

[7] Meta Platforms, Inc. *React.* URL: https://react.dev/.

[8] OpenAI Newsroom. *OpenAI Usage Statistics Tweet.* Twitter. Dec. 2024. URL: https://x.com/OpenAINewsroom/status/1864373399218475440.

[9] David Patterson et al. *Carbon Emissions and Large Neural Network Training.* 2021. arXiv: 2104.10350 [cs.LG]. URL: https://arxiv.org/abs/2104.10350.

[10] M. Patterson et al. *Water usage effectiveness (WUE™): a green grid data center sustainability metric*. 2011. URL: https://airatwork.com/wp-content/uploads/The-Green-Grid-White-Paper-35-WUE-Usage-Guidelines.pdf.

[11] RWDigital. “How Much Energy Do Google Search and ChatGPT Use?” In: *RWDigital Blog* (Oct. 2024). URL: https://www.rwdigital.ca/blog/how-much-energy-do-google-search-and-chatgpt-use/?utm_source=chatgpt.com.

[12] StatCounter. *Browser market share worldwide*. Feb. 2025. URL: https://gs.statcounter.com/browser-market-share#monthly-202309-202410-bar.

[13] Emma Strubell, Ananya Ganesh, and Andrew McCallum. “Energy and Policy Considerations for Deep Learning in NLP”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Ed. by Anna Korhonen, David Traum, and Lluís Màrquez. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 3645–3650. DOI: 10.18653/v1/P19-1355.

[14] Neil C. Thompson et al. “Deep Learning’s Diminishing Returns: The Cost of Improvement is Becoming Unsustainable”. In: *IEEE Spectrum* 58.10 (2021), pp. 50–55. DOI: 10.1109/MSPEC.2021.9563954.

[15] Matthijs ten Tije. *How Much CO2 Does ChatGPT Generate? Monitor Your Emissions inside the OpenAI API*. 2024. URL: https://tilburg.ai/2024/09/co2-chatgpt-emissions/.

[16] Alex de Vries. “The growing energy footprint of artificial intelligence”. In: *Joule* 7.10 (2023), pp. 2191–2194. ISSN: 2542-4351. DOI: https://doi.org/10.1016/j.joule.2023.09.004.

[17] *What is the CO2 emission per ChatGPT query?* 2024. URL: https://smartly.ai/blog/the-carbon-footprint-of-chatgpt-how-much-co2-does-a-query-generate.

[18] Josh You. *How much energy does ChatGPT use?* Feb. 7, 2025. URL: https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use.