# The HCI GenAI CO2ST Calculator: Calculating and Offsetting the Carbon Footprint of Generative AI Use in Human-Computer Interaction Research

Nanna Inie
Center for Computing Education Research (CCER)
IT University of Copenhagen
Copenhagen, Denmark
nans@itu.dk

Jeanette Falk
Computer Science
Aalborg University
Copenhagen, Denmark
jfo@cs.aau.dk

Raghavendra Selvan
University of Copenhagen
Copenhagen, Denmark
raghav@di.ku.dk


**Figure 1: A touch display is integrated in the stem of a large cardboard tree. The glowing "leaves" of the tree are yellow and red to evoke associations with late summer in Aarhus. A person can input how they have used GenAI in research and the screen will show an estimate of the** $CO_2$e **that this research has cost. A small thermal printer will print two receipts of their result: one to hang on the tree, and one to bring home. The receipts hanging on the tree will constitute a growing body of evidence for the** $CO_2$**st of GenAI use for HCI research. The user can exchange their** $CO_2$e **result into the number of tree seeds that, if grown for the next 10 years, would be enough to offset their carbon footprint.**

## Abstract

Increased usage of generative AI (GenAI) in Human-Computer Interaction (HCI) research has caused a sustainability crisis in computing, due to the excessive power consumption of developing and running these models. Energy consumption causes a massive carbon footprint. The exact energy usage and and subsequent carbon emissions are difficult to estimate in HCI research because HCI researchers most often use cloud-based services where the hardware and its energy consumption are hidden from plain view. The HCI GenAI $CO_2$st Calculator is a tool designed specifically for the HCI research pipeline, to help researchers estimate the energy consumption and carbon footprint of using generative AI in their research, either a priori (allowing for mitigation strategies or experimental redesign) or post hoc (allowing for transparent documentation of carbon footprint in written reports of the research).


*AAR Adjunct 2025, Aarhus N, Denmark*

ACM ISBN 979-8-4007-1968-4/25/08
https://doi.org/10.1145/3737609.3747110

## CCS Concepts

• **Computing methodologies** → *Artificial intelligence*; • **General and reference** → **Estimation**; • **Hardware** → **Impact on the environment**.

## Keywords

Sustainable HCI, Generative AI, Carbon Footprint

**ACM Reference Format:**
Nanna Inie, Jeanette Falk, and Raghavendra Selvan. 2025. The HCI GenAI CO2ST Calculator: Calculating and Offsetting the Carbon Footprint of Generative AI Use in Human-Computer Interaction Research. In *Adjunct proceedings of the sixth decennial Aarhus conference: Computing X Crisis (AAR Adjunct 2025), August 18–22, 2025, Aarhus N, Denmark.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3737609.3747110

## 1 Introduction and motivation

The extensive use of generative artificial intelligence (GenAI) models in research worldwide has a significant carbon footprint. In recent years, the electricity use by Meta, Amazon, Microsoft, and Google — main providers of cloud compute services — has more than doubled, and the the electricity consumption by global data centers has increased by 20-40% [10, 19]. Irish data centers are on the path to derail the entire country's climate targets [1].

In this paper we present the *HCI GenAI $CO_2$st Calculator*: a calculator designed specifically for Human-Computer Interaction (HCI) researchers to estimate the operational carbon emissions due to electricity consumption of GenAI use in their research. When researchers explore, test, and prototype with GenAI, their use causes $CO_2$e consumption. Additionally, HCI researchers create a massive downstream carbon footprint of GenAI by integrating these technologies into more systems, processes, and end user interactions.

With this calculator, we hope to make two contributions to the HCI community: First, we wish to enable HCI researchers to be fully transparent of their own research and acknowledge their own climate impact. Second, we hope to evoke critical thinking about the importance and necessity of HCI research conducted with the use of GenAI.

## 2 Background — sustainable HCI and carbon tracking

In 2007, Eli Blevis coined the term Sustainable Interaction Design (SID) and argued that "sustainability can and should be a central focus of interaction design" [5]. This perspective includes the responsible audit of tools we use to conduct our research. The large-scale adoption of GenAI tools is more than likely to contribute to the replication of "our modern society's overconsumption habits of natural resources within the digital space" [26].

Research on sustainable AI and machine learning (ML) generally fall into two camps: AI and ML *for* sustainability and sustainability *of* AI and ML, see e.g. [14, 27]. While a growing number of publications are directed towards AI *for* the United Nations (UN) Sustainable Development Goals, there is little research addressing the, often hidden, environmental costs of AI [14]. These efforts are significantly higher in ML and AI communities than in HCI. For example, because a model's architecture can affect how much power it consumes [3, 8], different more energy-efficient approaches in the IT-infrastructure, data, modeling, training, deployment, and evaluation of ML models have been suggested — see Bartoldson et al. [4], Mehlin et al. [20]. However, the inference stage or *use* of these models happens at a far greater scale; as of February 2025, ChatGPT alone boasted more than 400 million weekly active users [23].

There are several carbon and energy tracking tools available for ML/AI methods — Wright and colleagues discuss pros and cons of seven of these [27]. None of them, however, focus on HCI research, and many of their metrics do not make sense in an HCI context (such as specifying hardware used for computation and the ML tasks performed). Similarly, mitigation strategies directed at the architecture and training of models are rarely relevant to researchers outside ML, who rely on off-the-shelf, *multi-purpose* models. This exacerbates the sustainability issue for HCI researchers, since multi-purpose, generative architectures (such as the GPT models) are *orders of magnitude* more expensive than task-specific systems [19]. The lack of transparency from large multi-purpose model providers (such as OpenAI, Microsoft, and Google) about critical data, such as model training and hosting, complicates the issue.

## 3 The HCI $CO_2$st Calculator

The HCI $CO_2$st Calculator is, at its core, a calculator through which the researcher can input how they used GenAI in their research, which model they used, how much they used it, and the calculator will give an estimate of the carbon footprint in kg$CO_2$e, this GenAI research use has cost.

The physical exhibition of the calculator (see Figure 1) is an interactive piece, integrating a touch display in the stem of a large cardboard tree which has balloons for leaves. The glowing "leaves" of the tree are yellow and red to evoke associations with late summer/fall in Aarhus. A person can input how they have used GenAI in research — we will suggest that they take departure in research conducted for Aarhus Decennial 2025 if they have used GenAI in work submitted to the conference — and the screen will show an estimate of the $CO_2$e that this research has cost. The user can then print their result and hang their paper receipt on the tree stem (like someone might carve their names into the bark of a physical tree). The receipts hanging on the stem will constitute a growing body of evidence of our lasting carbon footprint. After this, the user can exchange their resulting $CO_2$e number for a number of Danish native tree seeds that, if planted and grown for the next 10 years, will offset their carbon footprint reported in this research.

The tree that will house the calculator for the physical demo is chosen as an exhibition piece to manifest the otherwise abstract relationship between computing use and its direct impact on climate and nature. Trees are one of our main sources for reducing $CO_2$e in the atmosphere, and we hope the tree will inspire people to consider the value of natural resources consumed in the pipeline that feeds our technological ecosystem.

### 3.1 Front-end and design

Based on Inie et al. [12], we create a flow that begins with choosing which research phase the use was part of (Research planning, Prototyping & building, Evaluation & user studies, Data collection, Analysis & synthesis, Dissemination & communication, or AI model training or fine-tuning). Based on a user's selection of research phase, the input fields will change to reflect which factors needs to be input to obtain an estimate.

The factors we use to calculate a credible estimate of $CO_2$e consumption of model use are: ***model type***, ***usage numbers***, and ***input/output resolution*** (depending on the model type). If GenAI is integrated into a prototype or used as part of a user study, we need to know the ***number of test runs*** and ***number of interactions***

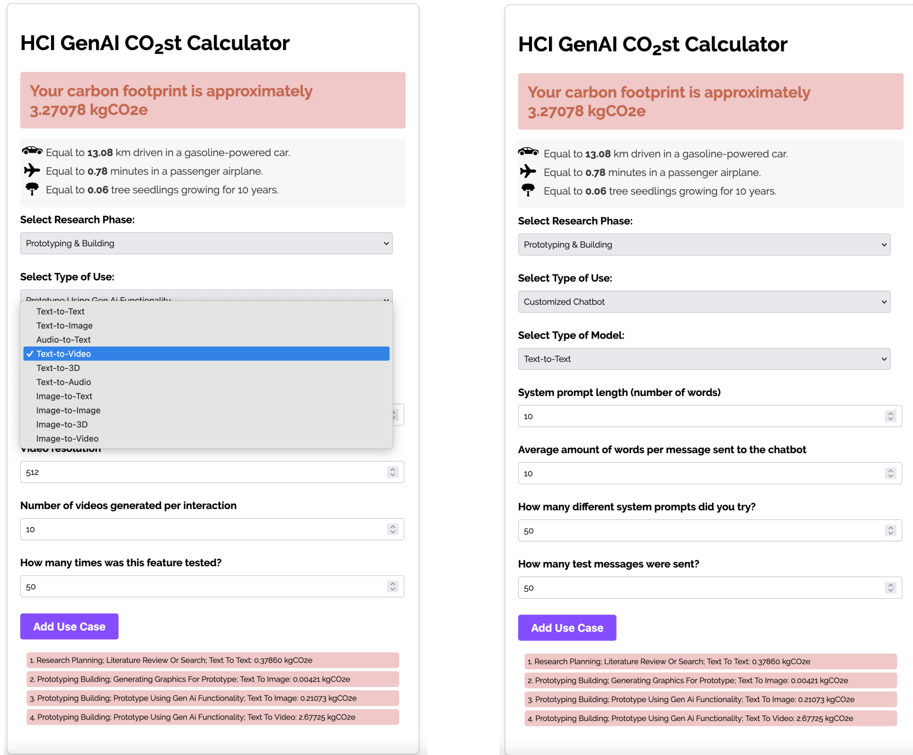


**(a) For the Prototyping using GenAI functionality use type, the user can choose between all model types.**

**(b) For the Customized chatbot type of use, the model type is "locked" to text-to-text.**

**Figure 2: Screenshots from the interface, showing how the input fields change when the user chooses different research phases and "stack" their use cases to account for the entire research pipeline.**

with the system). However, not all of these factors are relevant to all HCI research pipelines.

The goal is to translate the technical factors affecting the carbon footprint of GenAI use into an interface that makes it easy for an HCI researcher to audit their empirical research. The categorization imposed by the calculator encourages reflection about different GenAI uses that incur $CO_2$st which the researcher had not thought of, such as automatic transcription, automatic proofreading, or the generation of images for slides for a conference presentation.

Figure 2 shows two examples from the calculator. We see that the input fields are different when the *Type of use* is changed, mirroring the direct relevance to HCI research and simplifying the input. We attempt to limit the amount of choices, the user has to make to simplify the interaction as much as possible.

The result of the calculation is shown in a colored box on top of the page and updated when the user presses "Add use case". Use cases can be stacked because each research pipeline is likely to incur several GenAI uses, e.g., one for prototyping, and one for the subsequent user evaluation of a prototype, one for automatic transcription of audio data, and so forth. The result in kg$CO_2$e is translated into equivalent numbers: km. driven in a gasoline-powered car, number of minutes as a passenger on a commercial airplane, and number of tree seedlings grown for 10 years. These numbers are based on the EPA Greenhouse Gas Equivalencies Calculator. [1] We will hand out native tree seeds to participants for them to plant at home and thereby, over time, help "offset" their carbon footprint. We would like for this demo to not just be an informative gimmick, but to actually have a positive sustainability impact for the people who interact with it.

## 3.2 Back-end and algorithms

At a high-level, for each task we have measured the energy consumption for a single use (or prompt) denoted $E_p$ watt-hour (kWh). Using an in-house set-up comprising an NVIDIA-RTX3090 GPU, Intel-i7 processor with 32GB memory, we measured the energy consumption for various models using Carbontracker [2], which are reported in Table 1. The specific models shown in this table are used as proxies for the different model types (text-to-text, text-to-image, etc.) based on their popularity, ease-of-use, and availability (open-source). These choices provide useful approximations of the actual costs, which can vary between users due to differences in models and hardware used.

[1] https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator

**Table 1: Energy consumption per interaction for different model types.**

| Task | Model | $E_p$ (Wh) |
|---|---|---|
| Text-to-text | Llama-3.1-Instruct [9] | 0.004685 |
| Text-to-image | Stable-diffusion-XL [24] | 0.001301 |
| Audio-to-text | Whisper [25] | 0.006335 |
| Text-to-Video | AnimateDiff [17] | 0.021742 |
| Text-to-3D model | Shap-E [13] | 0.026320 |
| Text-to-Audio | MusicGen [7] | 0.011418 |
| Image-to-text | BLIP [15] | 0.003423 |
| Image-to-image | Instruct-Pix2Pix [6] | 0.000885 |
| Image-to-3D | One-2-3-45 [18] | 0.013010 |
| Video-to-text | XCLIP [21] | 0.001040 |
| Video-to-video | RIFE [11] | 0.026020 |
| Audio-to-audio | FreeVC [16] | 0.006335 |
| Image-to-video | SadTalker [28] | 0.026020 |

We have not included multi-modal models in the calculator, but instead reduced them to the most computationally heavy parameters (which results in a conservative estimate). For example, text+image-to-image becomes image-to-image or image+video-to-image becomes video-to-image, and so forth. This is done partially for simplicity of the interface, and partially to reduce our own carbon footprint when reproducing experiments.

We aggregate the usage information based on the user input into $N$, which is then used to estimate the overall energy consumption per use-case: $E = N \cdot E_p$ (kWh). This energy consumption is then converted to the carbon footprint using the global average carbon intensity of $CI = 0.481$ (kgCO2e/kWh) [22]. The final carbon footprint $C$ is estimated as: $C = CI \cdot E$ (kgCO2e).

## 4 Impact: awareness, transparency, and mitigation

When planning research with GenAI there is a range of trade-offs which the individual HCI researcher can make to reduce their carbon footprint. Many of these are opaque to a user of cloud-based models, as the factors which increase $CO_2$st are not clear or open. We intend for this system to have two practical impacts: First, to raise *awareness* of the carbon footprint caused by GenAI as it is typically used in HCI research, and second, enabling the HCI community to expect and increase *transparency* in reporting of research carbon footprint. The online version of the calculator (available at **www.hcico2st.com**) will enable HCI researchers to report the estimated carbon footprint of their research in a research paper's ethical statement. Hopefully, both awareness and transparency will lead to increased reflection upon researchers' own practices and potentially mitigation strategies for the planning of future experiments.

The calculator will show that the energy consumption grows almost linearly with the task load i.e., longer prompts or more images or images of higher resolution cost more in energy. It will show that the far most carbon intensive research use (aside from developing, training and fine-tuning new models) on average is large-scale open-ended generation of datasets, either for exploration or evaluation. Through the demo, we will also refer to the webpage of the calculator for concrete mitigation strategies for limiting one's carbon footprint when using GenAI in HCI research, such as enforcing user limits on prompting or making visible counters to show the user how many times they have prompted, refining prompting strategies (reducing the need for several attempts), and choosing task-specific rather than general-purpose models. We hope the physical exhibition will spark discussion and downstream use of this research tool.

## 5 Summary

This paper presents the HCI GenAI $CO_2$st Calculator, a system designed to help HCI researchers estimate the carbon footprint of using generative AI in their research. The interface is designed to represent typical HCI pipelines, and the calculations performed by the calculator are based on estimates derived from experiments run on our own hardware. The calculator is intended to be exhibited in a large cardboard tree with the receipts of different research pipelines hanging from the branches, and the possibility of offsetting one's concrete energy expenses by receiving tree seeds to plant after the conference. With this system, we hope to promote increased awareness and transparency in the HCI community about the climate impact of using GenAI in research.

## Acknowledgments

Funding acknowledgments: NI received funding from the VILLUM Foundation, grant 37176 (ATTiKA: Adaptive Tools for Technical Knowledge Acquisition). RS acknowledges funding received under Independent Research Fund Denmark (DFF) under grant 4307-00143B, and European Union's Horizon Europe Research and Innovation programme under grant agreements No. 101070284, No. 10107040 and No. 101189771.

## References

[1] Jillian Ambrose. 2024. Ireland's datacentres overtake electricity use of all urban homes combined — theguardian.com. https://www.theguardian.com/world/article/2024/jul/23/ireland-datacentres-overtake-electricity-use-of-all-homes-combined-figures-show. [Accessed 25-07-2024].

[2] Lasse F. Wolff Anthony, Benjamin Kanding, and Raghavendra Selvan. 2020. Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models. ICML Workshop on Challenges in Deploying and monitoring Machine Learning Systems. arXiv:2007.03051.

[3] Pedram Bakhtiarifard, Christian Igel, and Raghavendra Selvan. 2024. EC-NAS: Energy consumption aware tabular benchmarks for neural architecture search. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5660–5664.

[4] Brian R Bartoldson, Bhavya Kailkhura, and Davis Blalock. 2023. Compute-efficient deep learning: Algorithmic trends and opportunities. *Journal of Machine Learning Research* 24, 122 (2023), 1–77.

[5] Eli Blevis. 2007. Sustainable interaction design: invention & disposal, renewal & reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 503–512. doi:10.1145/1240624.1240705

[6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 18392–18402.

[7] Jade Copet, Felix Kreuk, Itai Gat, Tal Remez, David Kant, Gabriel Synnaeve, Yossi Adi, and Alexandre Défossez. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems* 36 (2024).

[8] Leon Derczynski. 2020. Power consumption variation over activation functions. *arXiv preprint arXiv:2006.07237* (2020).

[9] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).

[10] Ralph Hintemann and Simon Hinterholzer. 2022. Cloud computing drives the growth of the data center industry and its energy consumption. *Data centers* (2022).

[11] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. 2022. Real-time intermediate flow estimation for video frame interpolation. In *European Conference on Computer Vision*. Springer, 624–642.

[12] Nanna Inie, Jeanette Falk, and Raghavendra Selvan. 2025. How CO2STLY Is CHI? The Carbon Footprint of Generative AI in HCI Research and What We Should Do About It. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*.

[13] Heewoo Jun and Alex Nichol. 2023. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463* (2023).

[14] Sustainable AI Lab. 2024. About us - Bonn Sustainable AI Lab — sustainable-ai.eu. https://sustainable-ai.eu/about-us/. [Accessed 25-08-2024].

[15] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*. PMLR, 12888–12900.

[16] Jingyi Li, Weiping Tu, and Li Xiao. 2023. Freevc: Towards high-quality text-free one-shot voice conversion. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.

[17] Shanchuan Lin and Xiao Yang. 2024. Animatediff-lightning: Cross-model diffusion distillation. *arXiv preprint arXiv:2403.12706* (2024).

[18] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. 2024. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems* 36 (2024).

[19] Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power hungry processing: Watts driving the cost of AI deployment?. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 85–99.

[20] Vanessa Mehlin, Sigurd Schacht, and Carsten Lanquillon. 2023. Towards energy-efficient Deep Learning: An overview of energy-efficient approaches along the Deep Learning Lifecycle. arXiv:2303.01980 [cs.LG] https://arxiv.org/abs/2303.01980

[21] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. 2022. Expanding language-image pretrained models for general video recognition. In *European Conference on Computer Vision*. Springer, 1–18.

[22] Our World in Data. 2024. Carbon intensity of electricity generation. https://ourworldindata.org/grapher/carbon-intensity-electricity.

[23] Sarah Perez. 2025. ChatGPT doubled its weekly active users in under 6 months, thanks to new releases | TechCrunch — techcrunch.com. https://techcrunch.com/2025/03/06/chatgpt-doubled-its-weekly-active-users-in-under-6-months-thanks-to-new-releases/. [Accessed 06-05-2025].

[24] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=di52zR8xgf

[25] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*. PMLR, 28492–28518.

[26] Vanessa Utz and Steve DiPaola. 2023. Climate implications of diffusion-based generative visual ai systems and their mass adoption. In *Proceedings of the 14th International Conference on Computational Creativity*.

[27] Dustin Wright, Christian Igel, Gabrielle Samuel, and Raghavendra Selvan. 2023. Efficiency is Not Enough: A Critical Perspective of Environmentally Sustainable AI. arXiv:2309.02065 [cs.LG] https://arxiv.org/abs/2309.02065

[28] Wenxuan Zhang, Xiaodong Cun, Xuan Wang, Yong Zhang, Xi Shen, Yu Guo, Ying Shan, and Fei Wang. 2023. Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8652–8661.