

Green Recommender Systems – A Call for Attention

Joeran Beel

joeran.beel@uni-siegen.de
University of Siegen & Recommender-Systems.com
Siegen, Germany

Tobias Vente

tobias.vente@uni-siegen.de
Intelligent Systems Group
University of Siegen, Germany

Alan Said

alan@gu.se
University of Gothenburg
Gothenburg, Sweden

Lukas Wegmeth

lukas.wegmeth@uni-siegen.de
Intelligent Systems Group
University of Siegen, Germany

ACM Reference Format:

Joeran Beel, Alan Said, Tobias Vente, and Lukas Wegmeth. 2024. Green Recommender Systems – A Call for Attention. Recommender-Systems.com Blog. <https://recommender-systems.com/news/2024/10/07/green-recommender-systems-a-call-for-attention/>

Recommender systems are a cornerstone of many industries, from e-commerce to streaming services. Their role in shaping user choices is significant, driving customer engagement and revenue. However, the computational demands of recommender systems have surged – by a factor of 42 over the past decade [16]. As recommender systems become ubiquitous on digital platforms, the environmental and financial costs of their operation escalate. For example, creating a single recommender-system research paper emits an average of 3,297 kilograms of CO₂ [16], with large-scale production systems contributing even more to carbon emissions. This presents a new challenge to the community: sustainability.

Some researchers, including ourselves, have explored Green Recommender Systems, or sustainable recommender systems [2, 3, 8, 10, 11, 13, 16], and presented practical tools to measure the energy consumption of recommender system experiments [17]. However, much more work is needed to fully assess and reduce the environmental impact of recommender systems. Therefore, with this call for attention, we urge the recommender system community to thoroughly explore the topic of Green Recommender Systems as a framework for addressing sustainability.

We define Green Recommender Systems as follows

“Green Recommender Systems” are recommender systems designed to minimize their environmental impact throughout their life cycle – from research and design to implementation and operation. Green Recommender Systems typically aim to match the performance of traditional systems but may also accept trade-offs in accuracy or other metrics to prioritize sustainability. Minimizing environmental impact typically but not necessarily means minimizing energy consumption and CO₂ emissions.

Green Recommender Systems’ principles are not tied to specific algorithms or techniques¹. We also point out that we do not consider recommender systems that recommend eco-friendly items

as “green” if the systems themselves are not designed to minimize their environmental impact.

Examples of “green” recommender systems include the following. Training a deep learning-based recommender system on an Apple M1 chip, rather than an NVIDIA RTX 3090, may be considered “green”: For some algorithms, the M1 consumes up to 90% less energy than the NVIDIA GPU, while achieving the same performance, albeit with longer training times [16]. Similarly, down-sampling large datasets may be “green”, as not all recommender algorithms require large datasets for effective training or evaluation [2]. Also, alternatives to *k-fold* cross-validation, such as *e-fold* cross-validation [3, 5, 10], may be “green” since fewer folds are needed while maintaining comparable generalizability.

The importance of Green Recommender Systems lies in their potential to reduce the environmental footprint of large-scale recommendation engines and experiments. Recent estimates suggest machine learning systems, especially deep learning models, require vast computational power. Recommender systems, often deployed at massive scales, exacerbate this issue. These systems run continuously, analyzing user interactions, updating models, and delivering personalized recommendations in real time. The energy costs of these processes are often invisible but accumulate rapidly [14, 16].

Why should the community care about this issue? The push for sustainability in artificial intelligence and machine learning is rapidly gaining momentum [1, 6, 7, 9, 12, 15], but sustainable recommender systems, or Green Recommender Systems, remain significantly underrepresented in this discourse. Part of the reason is the pervasive focus on accuracy and user satisfaction metrics, which, while essential, overlook the growing environmental cost. Moreover, the current trend of ever-increasing model complexity and size directly conflicts with sustainability goals. Green Recommender Systems offer a necessary counterbalance by introducing a design philosophy that accounts for resource limitations and encourages the development of more efficient models.

Researchers and practitioners in recommender systems should consider energy and resource efficiency as part of their evaluation criteria. This shift could foster the development of more sustainable algorithms, architectures, and deployment strategies. The concept of “performance” should expand to include energy consumption alongside traditional metrics like accuracy and click-through rates. While some argue that energy-efficient models may sacrifice recommendation quality, the goal of Green Recommender Systems is to strike a balance between sustainability and performance or even maintain the same performance with fewer resources.

¹While we speak mainly of machine learning algorithms, Green Recommender Systems also apply to non-machine-learning algorithms.

Incorporating green principles into recommender systems offers practical benefits. Efficient models can reduce operational costs, lower server loads, and enhance user experience by delivering faster recommendations. As energy costs rise and concerns about the environmental impact of technology grow, companies have a financial incentive to invest in sustainable system designs. Green Recommender Systems can future-proof recommendation technologies against these emerging challenges.

The recommender systems community must recognize the critical importance of sustainability in its research agenda. By focusing on Green Recommender Systems, the field can contribute meaningfully to broader global efforts in AI for sustainability. It is no longer sufficient to pursue only accuracy and user satisfaction. Instead, the community must fully embrace the responsibility to develop systems that are both effective and sustainable.

The time has come for the recommender systems community to take action. Green Recommender Systems, or Green RecSys, are a practical, necessary response to the increasing computational demands of modern AI. The community must pay more attention to the environmental impact of their systems and work toward a future where sustainability is central to recommendation system design. This will not only benefit the planet but also ensure the long-term viability of the technology, while enhancing business value by reducing operational costs and aligning with the demand for eco-conscious solutions.²

REFERENCES

- [1] Yehia Ibrahim Alzoubi and Alok Mishra. 2024. Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production* 468 (2024), 143090. <https://doi.org/10.1016/j.jclepro.2024.143090>
- [2] Ardalan Arabzadeh, Tobias Vente, and Joeran Beel. 2024. Green Recommender Systems: Optimizing Dataset Size for Energy-Efficient Algorithm Performance. In *International Workshop on Recommender Systems for Sustainability and Social Good (RecSoGood) at the 18th ACM Conference on Recommender Systems (ACM RecSys) (2024-01-01)*.
- [3] Moritz Baumgart, Lukas Wegmeth, Tobias Vente, and Joeran Beel. 2024. e-Fold Cross-Validation for Recommender-System Evaluation. In *International Workshop on Recommender Systems for Sustainability and Social Good (RecSoGood) at the 18th ACM Conference on Recommender Systems (ACM RecSys) (2024-01-01)*.
- [4] Joeran Beel. 2024. Our use of AI-tools for writing research papers. In *Intelligent Systems Group, Blog*. <https://isg.beel.org/blog/2024/08/19/our-use-of-ai-tools-for-writing-research-papers/>
- [5] Joeran Beel, Lukas Wegmeth, and Tobias Vente. 2024. E-fold Cross-validation: A Computing and Energy-efficient Alternative to K-fold Cross-validation with Adaptive Folds [Proposal]. *OSF Preprints* (2024). <https://doi.org/10.31219/osf.io/exw3j>
- [6] Dagoberto Castellanos-Nieves and Luis García-Forte. 2024. Strategies of Automated Machine Learning for Energy Sustainability in Green Artificial Intelligence. *Applied Sciences (2076-3417)* 14, 14 (2024).
- [7] Dagoberto Castellanos-Nieves and Luis García-Forte. 2023. Improving Automated Machine-Learning Systems through Green AI. *Applied Sciences* 13, 20 (2023). <https://doi.org/10.3390/app132011583>
- [8] Alexander Felfernig, Manfred Wundara, Thi Ngoc Trang Tran, Seda Polat-Erdeniz, Sebastian Lubos, Merfat El Mansi, Damian Garber, and Viet-Man Le. 2023. Recommender systems for sustainability: overview and research issues. *Frontiers in big Data* 6 (2023), 1284511.
- [9] Leona Hennig, Tanja Tornede, and Marius Lindauer. 2024. Towards Leveraging AutoML for Sustainable Deep Learning: A Multi-Objective HPO Approach on Deep Shift Neural Networks. arXiv:2404.01965 [cs.LG] <https://arxiv.org/abs/2404.01965>
- [10] Christopher Mahlich, Tobias Vente, and Joeran Beel. 2024. From Theory to Practice: Implementing and Evaluating e-Fold Cross-Validation. In *International Conference on Artificial Intelligence and Machine Learning Research (CAIMLR) (2024-01-01)*. <https://isg.beel.org/blog/2024/09/16/e-fold-cross-validation/>
- [11] Pavel Merinov. 2023. Sustainability-oriented Recommender Systems. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (Limassol, Cyprus) (UMAP '23)*. Association for Computing Machinery, New York, NY, USA, 296–300. <https://doi.org/10.1145/3565472.3595617>
- [12] Samara O. S. Santos, Agustina Skiarski, Daniel García-Núñez, Victor Lazzarini, Rafael De Andrade Moral, Edgar Galvan, André L. C. Ottoni, and Erivelton Nepomuceno. 2024. Green Machine Learning: Analysing the Energy Efficiency of Machine Learning Models. In *2024 35th Irish Signals and Systems Conference (ISSC)*. 1–6. <https://doi.org/10.1109/ISSC61953.2024.10603302>
- [13] Giuseppe Spillo, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. [n. d.]. Towards Green Recommender Systems: Investigating the Impact of Data Reduction on Carbon Footprint and Algorithm Performances. In *18th ACM Conference on Recommender Systems (2024)*.
- [14] Giuseppe Spillo, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. 2023. Towards Sustainability-aware Recommender Systems: Analyzing the Trade-off Between Algorithms Performance and Carbon Footprint. In *Proceedings of the 17th ACM Conference on Recommender Systems (Singapore, Singapore) (RecSys '23)*. Association for Computing Machinery, New York, NY, USA, 856–862. <https://doi.org/10.1145/3604915.3608840>
- [15] Tanja Tornede, Alexander Tornede, Jonas Hanselle, Felix Mohr, Marcel Wever, and Eyke Hüllermeier. 2023. Towards green automated machine learning: Status quo and future directions. *Journal of Artificial Intelligence Research* 77 (2023), 427–457.
- [16] Tobias Vente, Lukas Wegmeth, Alan Said, and Joeran Beel. 2024. From Clicks to Carbon: The Environmental Toll of Recommender Systems. In *Proceedings of the 18th ACM Conference on Recommender Systems (ACM RecSys) (2024-01-01)*. <https://arxiv.org/pdf/2408.08203>
- [17] Lukas Wegmeth, Tobias Vente, Alan Said, and Joeran Beel. 2024. EMERS: Energy Meter for Recommender Systems. In *International Workshop on Recommender Systems for Sustainability and Social Good (RecSoGood) at the 18th ACM Conference on Recommender Systems (ACM RecSys) (2024-01-01)*. <https://arxiv.org/pdf/2409.15060>

²We used ChatGPT to improve the writing of this manuscript [4].