Green Recommender Systems: A Call for Attention

Joeran Beel
University of Siegen
Germany
joeran.beel@uni-siegen.de

Alan Said
University of Gothenburg
Sweden
alan@gu.se

Tobias Vente
University of Siegen
Germany

Lukas Wegmeth
University of Siegen
Germany

Abstract

The computational demands of recommender systems have drastically increased over the last decade, leading to higher energy consumption and carbon emissions. As recommender systems become more and more central to industries worldwide, their environmental impact is growing rapidly. However, this challenge also presents an opportunity. Designing recommender systems to minimize energy consumption and reduce environmental costs – this is what we define as *Green Recommender Systems* – can offer direct benefits not only for sustainability but also for companies and users. By optimizing resource use, these systems can reduce operational costs, improve efficiency, and even enhance user experience through faster recommendations. This paper calls for the recommender systems community to urgently integrate energy efficiency into their designs, considering it alongside traditional performance metrics. Addressing the environmental impact of recommender systems is essential, not just for the planet, but also for the long-term viability and cost-effectiveness of the technology.

1 Environmental Costs of Recommender Systems

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 [Vente et al., 2024]. 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 CO_2 [Vente et al., 2024], with large-scale production systems contributing even more to carbon emissions. This presents a new challenge to the community: sustainability.

2 Current Research on Green Recommender Systems

Some researchers, including ourselves, have explored Green Recommender Systems, or sustainable recommender systems [Purificato and Silvestri, 2024; Felfernig et al., 2023; Merinov, 2023; Spillo et al.; Vente et al., 2024; Arabzadeh et al., 2024; Baumgart et al., 2024; Spillo et al., 2024; Mahlich et al., 2024; Plaza et al., 2024], and presented practical tools to measure the energy consumption of recommender system experiments [Wegmeth et al., 2024]. 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.

3 Definition of 'Green Recommender Systems'

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 CO_2 emissions.*

Green Recommender Systems' principles are not tied to specific algorithms or techniques.[1] 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.

4 Examples of Green Recommender Systems

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 [Vente et al., 2024]. Similarly, down-sampling large datasets may be "green", as not all recommender algorithms require large datasets for effective training or evaluation [Arabzadeh et al., 2024; Spillo et al.]. Also, alternatives to *k-fold* cross-validation, such as *e-fold* cross-validation [Baumgart et al., 2024; Beel et al., 2024; Mahlich et al., 2024], may be considered "green" since fewer folds are needed while maintaining comparable generalizability.

5 The Importance of Green Recommender Systems

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

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

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 [Spillo et al., 2023; Vente et al., 2024].

Why should the community care about this issue? The push for sustainability in artificial intelligence and machine learning is rapidly gaining momentum [Santos et al., 2024; Alzoubi and Mishra, 2024; Castellanos-Nieves and García-Forte, 2024; Hennig et al., 2024; Tornede et al., 2023; Castellanos-Nieves and García-Forte, 2023], 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.

6 Call to Action

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.

Incorporating green principles into recommender systems offers practical benefits. Efficient models reduce operational costs and 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]

[2] We used ChatGPT to improve the writing of this manuscript [Beel, 2024].

References

Yehia Ibrahim Alzoubi and Alok Mishra. Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production*, 468:143090, 2024. ISSN 0959-6526. doi: https://doi.org/10.1016/j.jclepro.2024.143090. URL https://www.sciencedirect.com/science/article/pii/S0959652624025393.

Ardalan Arabzadeh, Tobias Vente, and Joeran Beel. 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.

Moritz Baumgart, Lukas Wegmeth, Tobias Vente, and Joeran Beel. 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.

Joeran Beel. Our use of ai-tools for writing research papers. In *Intelligent Systems Group, Blog*, 2024. URL https://isg.beel.org/blog/2024/08/19/our-use-of-ai-tools-for-writing-research-papers/.

Joeran Beel, Lukas Wegmeth, and Tobias Vente. E-fold cross-validation: A computing and energy-efficient alternative to k-fold cross-validation with adaptive folds [proposal]. *OSF Preprints*, 2024. doi: 10.31219/osf.io/exw3j. URL https://osf.io/preprints/osf/exw3jhttps://isg.beel.org/blog/2024/04/07/e-fold-cross-validation-a-computing-and-energy-efficient-alternative-to-k-fold-cross-validation-with-adaptive-folds-proposal/.

Dagoberto Castellanos-Nieves and Luis García-Forte. Strategies of automated machine learning for energy sustainability in green artificial intelligence. *Applied Sciences (2076-3417)*, 14(14), 2024.

Dagoberto Castellanos-Nieves and Luis García-Forte. Improving automated machine-learning systems through green ai. *Applied Sciences*, 13(20), 2023. ISSN 2076-3417. doi: 10.3390/app132011583. URL https://www.mdpi.com/2076-3417/13/20/11583.

Alexander Felfernig, Manfred Wundara, Thi Ngoc Trang Tran, Seda Polat-Erdeniz, Sebastian Lubos, Merfat El Mansi, Damian Garber, and Viet-Man Le. Recommender systems for sustainability: overview and research issues. *Frontiers in big Data*, 6:1284511, 2023.

Leona Hennig, Tanja Tornede, and Marius Lindauer. Towards leveraging automl for sustainable deep learning: A multi-objective hpo approach on deep shift neural networks. 2024. URL https://arxiv.org/abs/2404.01965.

Christopher Mahlich, Tobias Vente, and Joeran Beel. From theory to practice: Implementing and evaluating e-fold cross-validation. In *International Conference on Artificial Intelligence and Machine Learning Research (CAIMLR)*, 2024. URL https://isg.beel.org/blog/2024/09/16/e-fold-cross-validation/.

Pavel Merinov. Sustainability-oriented recommender systems. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*, UMAP '23, page 296–300, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399326. doi: 10.1145/3565472.3595617. URL https://doi.org/10.1145/3565472.3595617.

Alejandro Plaza, Juan Gil, and Denis Parra Santander. 14 kg of co2: Analyzing the carbon footprint and performance of session-based recommendation algorithms. In *RecSoGood Workshop*, 2024.

Antonio Purificato and Fabrizio Silvestri. Eco-aware graph neural networks for sustainable recommendations. *arXiv*, 2024. URL https://arxiv.org/abs/2410.09514.

Samara O. S. Santos, Agustina Skiarski, Daniel García-Núñez, Victor Lazzarini, Rafael De Andrade Moral, Edgar Galvan, André L. C. Ottoni, and Erivelton Nepomuceno. Green machine learning: Analysing the energy efficiency of machine learning models. In *2024 35th Irish Signals and Systems Conference (ISSC)*, pages 1–6, 2024. doi: 10.1109/ISSC61953.2024.10603302.

Giuseppe Spillo, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. Towards green recommender systems: Investigating the impact of data reduction on carbon footprint and algorithm performances. In *18th ACM Conference on Recommender Systems*.

Giuseppe Spillo, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. 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*, RecSys '23, page 856–862, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400702419. doi: 10.1145/3604915.3608840. URL https://doi.org/10.1145/3604915.3608840.

Giuseppe Spillo, Alberto Gaetano Valerio, Felice Franchini, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. Recsys carbonator: Predicting carbon footprint of recommendation system models. In *RecSoGood Workshop*, 2024.

Tanja Tornede, Alexander Tornede, Jonas Hanselle, Felix Mohr, Marcel Wever, and Eyke Hüllermeier. Towards green automated machine learning: Status quo and future directions. *Journal of Artificial Intelligence Research*, 77:427–457, 2023.

Tobias Vente, Lukas Wegmeth, Alan Said, and Joeran Beel. From clicks to carbon: The environmental toll of recommender systems. In *Proceedings of the 18th ACM Conference on Recommender Systems (ACM RecSys)*, 2024. URL https://arxiv.org/pdf/2408.08203.

Lukas Wegmeth, Tobias Vente, Alan Said, and Joeran Beel. 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. URL https://arxiv.org/pdf/2409.15060.