

# Green Recommender Systems: Optimizing Dataset Size for Energy-Efficient Algorithm Performance

Ardalan Arabzadeh<sup>1</sup>, Tobias Vente<sup>1</sup>, and Joeran Beel<sup>1,2</sup>

<sup>1</sup> University of Siegen, Department of Electrical Engineering and Computer Science, Germany

<sup>2</sup> Recommender-Systems.com, Siegen, Germany

[joeran.beel@uni-siegen.de](mailto:joeran.beel@uni-siegen.de)

<https://www.recommender-systems.com>

**Abstract.** As recommender systems become increasingly prevalent, the environmental impact and energy efficiency of training these large-scale models have come under scrutiny. This paper investigates the potential for energy-efficient algorithm performance by optimizing dataset sizes through downsampling techniques. We conducted experiments on the MovieLens 100K, 1M, 10M and Amazon Toys and Games datasets, analyzing the performance of various recommender algorithms under different portions of dataset size. Our results indicate that while more training data generally leads to higher performance in algorithms, certain algorithms, such as FunkSVD and BiasedMF, particularly in cases involving more unbalanced and sparse dataset like Amazon Toys and Games, maintain high-quality recommendations with up to 50% reduction in training data, achieving nDCG@10 scores within  $\sim 13\%$  of their full dataset performance. These findings suggest that strategic dataset reduction can decrease computational and environmental costs without substantially compromising recommendation quality. This study advances sustainable and green recommender systems by providing actionable insights for reducing energy consumption while maintaining effectiveness.

**Keywords:** Sustainability, Green RecSys, Dataset Downsampling, Energy Efficiency, Environmental Impact.

**Reference:** Arabzadeh, Ardalán; Vente, Tobias; Beel, Joeran. 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)*.

## 1 Introduction

Advancements in recommender systems have enhanced user experience. However, these advancements came at a substantial computational and energy cost [30, 35]. Large datasets do not only increase operational expenses but also result in higher energy consumption and carbon emissions, contributing to a more significant environmental impact [1, 16, 30, 31, 35]. In extreme cases, energy consumption

between datasets differs by factor 1,444, such as between LastFM vs. Yelp with the DGCF algorithm [35].

Given the environmental and computational challenges associated with large datasets, it's important to question whether using the entire datasets is always necessary. For instance, datasets like MovieLens 10M are frequently employed in training recommender systems. But is it necessary to use 10 million instances, especially with simple baseline algorithms? Or would downsampling the dataset to e.g. 10% suffices, which, in turn, might save 90% of energy? Similarly, dataset size might be an important factor when choosing datasets for recommender-system experiments [10].

In our work, we investigate whether downsampling the dataset can lead to an acceptable trade-off between energy efficiency and the performance of recommender algorithms. We see our work in the context of "Green Recommender Systems" as defined by Beel et al. as follows [9].

*"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<sub>2</sub> emissions. [9]*

For our current work, we hypothesize that by downsampling recommender system datasets, we can save time, energy, and CO<sub>2</sub> emissions, while obtaining nearly the same performance as with the full datasets.

## 2 Related Work

The field of green recommender systems has only started evolving recently [30, 35]. We also recently proposed "e-fold cross-validation", an energy-efficient alternative to k-fold cross-validation [5, 11, 23]. Wegmeth et al. introduced EMERS, a tool to measure the electricity consumption of recommender system experiments [36]. Also, judging on the "accepted papers" list of the RecSoGood workshop, more related work is to be published very soon [26, 32].

While the "green" concept in recommender systems is new, other disciplines, like Automated Machine Learning, explore options to save energy for a longer time [2, 13, 14, 20, 27, 29, 33].

In the domain of recommender systems, several studies have explored the impact of dataset size on the efficiency of algorithmic performance, which aligns the key focus of this study. Notably, Bentzer and Thulin explore the trade-off between accuracy and computational efficiency in collaborative filtering algorithms under limited data conditions [12]. They found that IBCF algorithm performs better in terms of accuracy with smaller datasets compared to SVD algorithm, while SVD outperforms IBCF in terms of speed and scalability with larger datasets. Their study highlights the performance differences between these two algorithms

but does not address how other algorithms perform under similar constraints. This gap is relevant to our research, which seeks to evaluate a wider range of algorithms for optimizing both energy efficiency and performance.

Additionally, Jain and Jindal’s review emphasizes that strategic sampling and filtering can enhance recommendation efficiency by improving computational speed and accuracy [21]. However, their review lacks experimental validation of how these techniques impact algorithm performance with varying dataset sizes. Our study addresses this gap by empirically evaluating these effects on recommender systems. Judging based on the paper’s title and abstract, Spillo et al. appear to have conducted research similar to ours [31]. However, at the time of conducting our research and writing our manuscript, the work of Spillo et al. was not yet publicly available but only announced on the ACM Recommender Systems conference website as an accepted paper.

It is worth mentioning that most papers and experiments focusing on downsampling and data efficiency are predominantly conducted in domains like (automated) machine learning, AI and computer vision [2, 13, 14, 20, 27, 33], with more extensive research compared to the recommender systems domain. Moreover, studies within this broader field also corroborate the potential benefits of downsampling. Research by Zogaj et al. demonstrates that reducing dataset sizes can enhance both computational efficiency and predictive accuracy in genetic programming-based AutoML systems [37]. Their experiments show that downsampling large datasets can even in some cases result in better performance than using the full dataset, with shorter search times.

These studies underscore and evaluate the potential benefits of downsampling and its impact on model performance, but are not directly applicable to traditional recommender system algorithms, where such effects remain underexplored. Especially when thinking of automated recommender systems (AutoRecSys [3, 19, 34]), where large spaces of configurations must be searched, energy efficiency is fundamentally important.

### 3 Methodology

#### 3.1 Datasets & Preprocessing

We used four datasets for our experiment: **MovieLens 100K**, **MovieLens 1M**, **MovieLens 10M**, and **Amazon Toys and Games**. The MovieLens datasets feature relatively balanced ratings across a scale from 1 to 5. In contrast, the Amazon Toys and Games dataset exhibits a skewed distribution, with ~90% of ratings concentrated in the 4 and 5 ranges. The following preprocessing steps were applied to the datasets: removal of duplicate rows, averaging duplicate ratings, and applying 10-core pruning to retain users and items with at least 10 interactions.

The dataset details before and after preprocessing are in Table 1.

Table 1: Basic information of datasets before and after preprocessing

Dataset	Before Preprocessing					After Preprocessing				
	#Users	#Items	#Interactions	Avg. #Int. per user	Avg. #Int. per item	#Users	#Items	#Interactions	Avg. #Int. per user	Avg. #Int. per item
MovieLens 100K	943	1,682	100,000	106	59	943	1,152	97,953	103	85
MovieLens 1M	6,040	3,706	1,000,209	165	269	6,040	3,260	998,539	165	306
MovieLens 10M	69,878	10,677	10,000,054	143	936	69,878	9708	9,995,471	143	1029
Amazon Toys and Games	208,180	78,772	1,828,971	8	23	11,609	8,443	202,721	17	24

### 3.2 Data Splitting and Downsampling

We applied a User-Based Split [24], with 10% of each user’s interactions randomly selected for the test set, 10% for validation, and 80% for training. The validation set was used for hyperparameter tuning, maintaining a comparable size between the validation and test sets to account for the impact of the training-to-validation/test ratio on results, as highlighted in prior research [15]. The training set was downsampled to various proportions (10%, 20%, 30%, up to 100%) by randomly selecting different portions of each user’s interactions. This approach ensures consistency in user representation across all sets while varying the number of interactions in the training set.

### 3.3 Algorithms and Evaluation

We trained the following algorithms on the downsampled training sets using the LensKit [17] and RecPack [25] libraries:

Table 2: Information of algorithms used in our experiment

Algorithms	Bias	Popular	Random	UserKNN	ItemKNN	BiasedMF	FunksVD	Popularity	ItemKNN	SVD	NMF
Library	LensKit	LensKit	LensKit	LensKit	LensKit	LensKit	LensKit	RecPack	RecPack	RecPack	RecPack

Performance was evaluated using the nDCG@10 metric, ensuring that all libraries adhered to an identical standard calculation logic to facilitate a fair comparison of results across algorithms [28].

## 4 Results & Conclusion

Our research investigated the impact of downsampling on the efficiency of recommender system algorithms by analyzing performance metrics and dataset characteristics. Variations in user/item interaction densities and rating distributions, as discussed in subsection 3.1, impact algorithm performance. Preprocessing facilitated consistent evaluations across varying dataset sizes. Before presenting the experimental results (Figure 1), it is useful to estimate the potential environmental benefits of the downsampling strategy proposed in this work, specifically in terms of reducing carbon footprint and CO<sub>2</sub>e emissions, with a calculation example where the training set is downsampled to 50% of its original size.

Based on our observations and calculations, downsampling the training data to 50% reduces the runtime for training and evaluation phases to  $\sim 72\%$  of the runtime required for the full dataset, on average. Furthermore, the energy consumption for a single run of a recommender algorithm on one dataset is estimated at 0.51 kWh [35]. Assuming 10 hyperparameter configurations per algorithm and using the global average conversion factor of 481 gCO<sub>2</sub>e per kWh [18], and accounting for a potential increase by a factor of 40 to consider preliminary tasks such as algorithm prototyping, initial tests, debugging, and re-runs [35], we estimate the potential carbon equivalent emissions savings from downsampling the training set to 50% compared to the full set per algorithm per dataset as follows:

$$(100\% - 72\%) \times 0.51 \text{ kWh} \times 10 \times 481 \text{ gCO}_2\text{e/kWh} \times 40 \approx 27.4 \text{ KgCO}_2\text{e}.$$

This estimation roughly quantifies the reduction in CO<sub>2</sub>e emissions resulting from the training of a single algorithm on a single dataset, based solely on the reduction in runtime following downsampling. It assumes that the hardware used for the full dataset will also be employed for the downsampled dataset and that a nearly linear relationship exists between runtime, energy consumption, and carbon emissions, as supported by the ML CO<sub>2</sub> Impact calculator tool [22]. In the upcoming sections, we detail the principal observations derived from our results and delve into how they inform the objectives of our research. For simplicity in discussing the algorithms examined in this study, we have categorized the algorithms into two groups. This division reflects the observed similarities in performance and results within each group, with distinct behaviors compared to the other group as shown in Figure 1, facilitating clearer analysis of their comparative effectiveness. The Random algorithm serves as a baseline for comparison but is not included in the statistics of either group. Table 3 provides an overview of these categorizations.

Table 3: Categorization of Examined Algorithms

Group	Algorithms included
Group 1	UserKNN, SVD, ItemKNN (both LensKit and RecPack version), NMF
Group 2	Bias, Popularity, FunkSVD, BiasedMF, Popular

**Observations** Several key observations can be outlined from our analysis of recommender system algorithms across different datasets, each numbered for easy reference. (1) larger datasets consistently resulted in improved performance across all algorithms, with Group 1 algorithms benefiting significantly from increased data availability. (2) In examining the performance metrics, we observed that Group 1 algorithms displayed significant improvements when the dataset used for training exceeded  $\sim 30\%$  of the total data. Specifically, downsampling the MovieLens 100K dataset to  $\sim 50\%$  resulted in a  $\sim 50\%$  decrease in average nDCG@10 values of this group’s algorithms, while reducing to  $\sim 30\%$  led to a  $\sim 65\%$  decrease, highlighting a near-linear relationship between dataset size and performance. (3) Conversely, Group 2 algorithms demonstrated more gradual

performance improvements, with nDCG@10 values decreasing by  $\sim 23\%$  and  $\sim 29\%$  in average when the dataset was downsampled to  $\sim 50\%$  and  $\sim 30\%$ , respectively. (4) The sparse Amazon Toys and Games dataset particularly illustrated a more pronounced performance gap between these two groups of algorithms. When downsampling to  $\sim 50\%$  and  $\sim 30\%$ , Group 2 algorithms experienced only  $\sim 13\%$  and  $\sim 17\%$  average drops in performance, respectively, which is less severe compared to the denser MovieLens datasets.

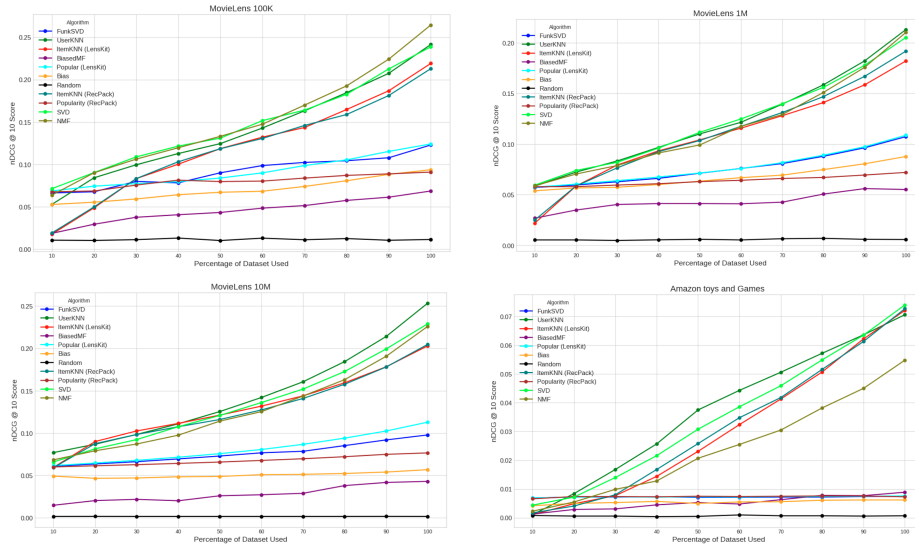


Fig. 1: nDCG@10 scores for each algorithm trained on varying portions of the datasets. The horizontal axis shows the percentage of the full training set, where 100% equals 80% of the total dataset, with other percentages relative to this.

**Interpretation** From these observations, it appears that the size and sparsity of datasets significantly influence the performance of recommender system algorithms. Observation (1) highlights that contrary to our expectations, larger data volumes, including those from extensive datasets like MovieLens 10M, generally lead to better algorithm performance. However, the extent of improvement depends on the specific algorithm and the characteristics of the dataset. Observations (2) and (3) highlight that Group 1 algorithms are highly dependent on larger datasets to perform optimally. In contrast, Group 2 algorithms maintain relatively stable performance even with reduced data, striking a balance between performance and computational efficiency. This observation is evident from the narrower gap in the nDCG@10 scores distribution box plot between 50% and 100% dataset utilization for algorithms in Group 2, compared to the larger gap seen in Group 1, as shown in Figure 2. The detailed analysis in observation

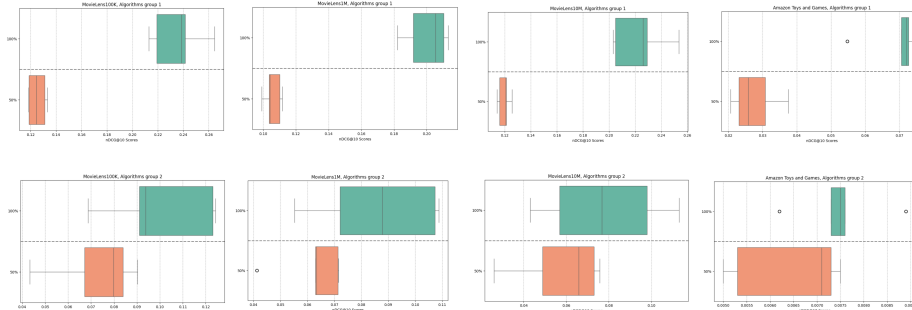


Fig. 2: Distribution of nDCG@10 Scores Across Four examined Datasets for Two Distinct Algorithm Groups at 50% and 100% Dataset Utilization.

(4) shows that in sparse environments, such as the Amazon Toys and Games dataset, downsampling effectively reduces computational demands with only minimal performance loss. This indicates that strategic downsampling can be a viable method especially in contexts where energy optimization is crucial without significantly compromising accuracy.

**Conclusion** This study underscores the potential for optimizing recommender systems through dataset size reduction. Although most algorithms demonstrate enhanced performance with larger training datasets, our analysis has pinpointed specific scenarios where the trade-off between energy efficiency and accuracy favors efficiency. In these cases, significant savings are achieved with minimal detriment to accuracy. Some algorithms consistently maintain high performance even with reduced data volumes, highlighting their potential for energy-efficient AI development.

Therefore, we answer our research question by affirming that it can be possible to identify an optimal trade-off between maintaining algorithmic performance and reducing dataset size. Specifically, our analysis shows that strategic downsampling may improve energy efficiency while maintaining performance comparable to the original dataset size, thereby supporting the optimization of AI systems and recommenders. However, more research is necessary to find out when exactly downsampling is a sensible approach, as sometimes, performance varies notably. We hope that in the long term, downsampling datasets becomes an accepted best-practice [7, 8], for the recommender-system community that helps contributing to green and sustainable recommender systems.

**Acknowledgment** This paper benefited from ChatGPT for grammar and wording improvements [6]. The code used for conducting the experiments is publicly available on GitHub [4].

## References

1. Al-Jarrah, O., Yoob, P., Muhaidat, S., Karagiannidis, G., Taha, K.: Efficient machine learning for big data: A review. *Big Data Research* **2**(3), 87–99 (2015). <https://doi.org/10.1016/j.bdr.2015.04.002>
2. Alzoubi, Y.I., Mishra, A.: Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production* **468**, 143090 (2024). <https://doi.org/https://doi.org/10.1016/j.jclepro.2024.143090>, <https://www.sciencedirect.com/science/article/pii/S0959652624025393>
3. Anand, R., Beel, J.: Auto-surprise: An automated recommender-system (autorecsys) library with tree of parzens estimator (tpe) optimization. In: 14th ACM Conference on Recommender Systems (RecSys). pp. 1–4 (2020), <https://arxiv.org/abs/2008.13532>
4. Arabzadeh, A.: Green recsys github repository (2024), <https://github.com/Ardalan224/RecSoGood2024/>
5. Baumgart, M., Wegmeth, L., Vente, T., Beel, J.: 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)
6. Beel, J.: Our use of ai-tools for writing research papers (2024), <https://isg.beel.org/blog/2024/08/19/our-use-of-ai-tools-for-writing-research-papers/>, in: Intelligent Systems Group, Blog
7. Beel, J.: A call for evidence-based best-practices for recommender systems evaluations. In: Bauer, C., Said, A., Zangerle, E. (eds.) Report from Dagstuhl Seminar 24211: Evaluation Perspectives of Recommender Systems: Driving Research and Education (2024). <https://doi.org/10.31219/osf.io/djuac>, [https://isg.beel.org/pubs/2024\\_Call\\_for\\_Evidence\\_Based\\_RecSys\\_Evaluation\\_\\_Pre\\_Print\\_.pdf](https://isg.beel.org/pubs/2024_Call_for_Evidence_Based_RecSys_Evaluation__Pre_Print_.pdf)
8. Beel, J., Jannach, D., Said, A., Shani, G., Vente, T., Wegmeth, L.: Best-practices for offline evaluations of recommender systems. In: Bauer, C., Said, A., Zangerle, E. (eds.) Report from Dagstuhl Seminar 24211 – Evaluation Perspectives of Recommender Systems: Driving Research and Education (2024)
9. Beel, J., Said, A., Vente, T., Wegmeth, L.: Green recommender systems – a call for attention. *Recommender-Systems.com Blog* (2024). <https://doi.org/10.31219/osf.io/5ru2g>, [https://isg.beel.org/pubs/2024\\_Green\\_Recommender\\_Systems-A\\_Call\\_for\\_Attention.pdf](https://isg.beel.org/pubs/2024_Green_Recommender_Systems-A_Call_for_Attention.pdf)
10. Beel, J., Wegmeth, L., Michiels, L., Schulz, S.: Informed dataset-selection with algorithm performance spaces. In: 18th ACM Conference on Recommender Systems (ACM RecSys) (2024)
11. Beel, J., Wegmeth, L., Vente, T.: 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>, <https://osf.io/preprints/osf/exw3j>
12. Bentzer, C., Thulin, H.: Recommender systems using limited dataset sizes (June 8 2023), degree Project in Computer Science and Engineering, First cycle, 15 credits, KTH Royal Institute of Technology
13. Castellanos-Nieves, D., García-Forte, L.: Strategies of automated machine learning for energy sustainability in green artificial intelligence. *Applied Sciences* (2076-3417) **14**(14) (2024)



14. Castellanos-Nieves, D., García-Forte, L.: Improving automated machine-learning systems through green ai. *Applied Sciences* **13**(20) (2023). <https://doi.org/10.3390/app132011583>, <https://www.mdpi.com/2076-3417/13/20/11583>
15. Cañamares, R., Castells, P., Moffat, A.: Offline evaluation options for recommender systems. *Information Retrieval Journal* **23**(4), 387–410 (2020). <https://doi.org/10.1007/s10791-020-09371-3>
16. Chen, C., Zhang, P., Zhang, H., Dai, J., Yi, Y., Zhang, H., Zhang, Y.: Deep learning on computational-resource-limited platforms: A survey. *Advances in Artificial Intelligence* (2020). <https://doi.org/10.1155/2020/8454327>, first published: 01 March 2020
17. Ekstrand, M.: Lenskit for python: Next-generation software for recommender systems experiments. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. pp. 2999–3006. Virtual Event, Ireland (2020). <https://doi.org/10.1145/3340531.3412778>
18. Ember: Carbon intensity of electricity generation – ember and energy institute (2024), <https://ourworldindata.org/grapher/carbon-intensity-electricity>, yearly Electricity Data by Ember; Statistical Review of World Energy by Energy Institute. Dataset processed by Our World in Data
19. Gupta, S., Beel, J.: Auto-caserec: Automatically selecting and optimizing recommendation-systems algorithms. *OSF Preprints DOI:10.31219/osf.io/4znmd*, (2020). <https://doi.org/10.31219/osf.io/4znmd>
20. Hennig, L., Tornede, T., Lindauer, M.: Towards leveraging automl for sustainable deep learning: A multi-objective hpo approach on deep shift neural networks. In: *arXiv* (2024), <https://arxiv.org/abs/2404.01965>
21. Jain, K., Jindal, R.: Sampling and noise filtering methods for recommender systems: A literature review. *Engineering Applications of Artificial Intelligence* **122**, 106129 (2023)
22. Lacoste, A., Luccioni, A., Schmidt, V., Dandres, T.: Quantifying the carbon emissions of machine learning. *arXiv preprint* (2019), <https://mlco2.github.io/impact/>, arXiv:1910.09700 [cs.CY]
23. Mahlich, C., Vente, T., Beel, J.: From theory to practice: Implementing and evaluating e-fold cross-validation. In: *International Conference on Artificial Intelligence and Machine Learning Research (CAIMLR)* (2024), <https://isg.beel.org/blog/2024/09/16/e-fold-cross-validation/>
24. Meng, Z., McCreadie, R., Macdonald, C., Ounis, I.: Exploring data splitting strategies for the evaluation of recommendation models. In: *Proceedings of RecSys '20: The 14th ACM Recommender Systems Conference (RecSys '20)*. p. 8. ACM, New York, NY, USA (2020). <https://doi.org/10.1145/1122445.1122456>
25. Michiels, L., Verachtert, R., Goethals, B.: Recpack: An(other) experimentation toolkit for top-n recommendation using implicit feedback data. In: *Proceedings of the 16th ACM Conference on Recommender Systems*. p. 648–651. RecSys '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3523227.3551472>, <https://doi.org/10.1145/3523227.3551472>
26. Plaza, A., Gil, J., Parra Santander, D.: 14 kg of co2: Analyzing the carbon footprint and performance of session-based recommendation algorithms. In: *RecSoGood Workshop* (2024)
27. Santos, S.O.S., Skiarski, A., García-Núñez, D., Lazzarini, V., De Andrade Moral, R., Galvan, E., Ottoni, A.L.C., Nepomuceno, E.: Green machine learning: Analysing the energy efficiency of machine learning models. In: *2024 35th Irish Signals and Systems Conference (ISSC)*. pp. 1–6 (2024). <https://doi.org/10.1109/ISSC61953.2024.10603302>

28. Schmidt, M., Prinz, T., Nitschke, J.: Evaluating the performance-deviation of itemknn in recbole and lenskit. arXiv preprint (2024), <https://arxiv.org/abs/2407.13531>, arXiv:2407.13531v1 [cs.LG]
29. Schwartz, R., Dodge, J., Smith, N.A., Etzioni, O.: Green ai. *Commun. ACM* **63**(12), 54–63 (Nov 2020). <https://doi.org/10.1145/3381831>, <https://doi.org/10.1145/3381831>
30. Spillo, G., De Filippo, A., Milano, M., Musto, C., Semeraro, G.: Towards sustainability-aware recommender systems: Analyzing the trade-off between algorithms performance and carbon footprint. In: *Proceedings of the ACM Conference*. p. 7. Singapore, Singapore (2023). <https://doi.org/10.1145/3604915.3608840>
31. Spillo, G., De Filippo, A., Musto, C., Milano, M., Semeraro, G.: Towards green recommender systems: Investigating the impact of data reduction on carbon footprint and algorithm performances. In: *18th ACM Conference on Recommender Systems* (2024)
32. Spillo, G., Valerio, A.G., Franchini, F., De Filippo, A., Musto, C., Milano, M., Semeraro, G.: Recsys carbonator: Predicting carbon footprint of recommendation system models. In: *RecSoGood Workshop* (2024)
33. Tornede, T., Tornede, A., Hanselle, J., Mohr, F., Wever, M., Hüllermeier, E.: Towards green automated machine learning: Status quo and future directions. *arXiv / Journal of Artificial Intelligence Research* **77**, 427–457 (2021 / 2023)
34. Vente, T., Ekstrand, M., Beel, J.: Introducing lenskit-auto, an experimental automated recommender system (autorecsys) toolkit. In: *Proceedings of the 17th ACM Conference on Recommender Systems*. pp. 1212–1216 (2023), <https://dl.acm.org/doi/10.1145/3604915.3610656>
35. Vente, T., Wegmeth, L., Said, A., Beel, J.: From clicks to carbon: The environmental toll of recommender systems. In: *Proceedings of the 18th ACM Conference on Recommender Systems*. p. 580–590. *RecSys '24*, Association for Computing Machinery, New York, NY, USA (2024). <https://doi.org/10.1145/3640457.3688074>, <https://arxiv.org/abs/2408.08203>
36. Wegmeth, L., Vente, T., Said, A., Beel, J.: 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), <https://arxiv.org/pdf/2409.15060>
37. Zogaj, F., Cambronero, J., Rinard, M., Cito, J.: Doing more with less: Characterizing dataset downsampling for automl. *Proceedings of the VLDB Endowment (PVLDB)* **14**(11), 2059–2072 (2021). <https://doi.org/10.14778/3476249.3476262>