



Thinking Geographically about AI Sustainability

Meilin Shi ¹, Kitty Currier ², Zilong Liu ¹, Krzysztof Janowicz^{1,2}, Nina Wiedemann ³,
Judith Verstegen ⁴, Grant McKenzie ⁵, Anita Graser ⁶, Rui Zhu ⁷, and Gengchen Mai ⁸

¹Department of Geography and Regional Research, University of Vienna, Austria

²Department of Geography, University of California, Santa Barbara, USA

³Institute of Cartography and Geoinformation, ETH Zurich, Switzerland

⁴Department of Human Geography and Spatial Planning, Utrecht University, the Netherlands

⁵Department of Geography, McGill University, Canada

⁶AIT Austrian Institute of Technology, Austria

⁷School of Geographical Sciences, University of Bristol, UK

⁸Department of Geography, University of Georgia, USA

Correspondence: Meilin Shi (meilin.shi@univie.ac.at)

Abstract. Driven by foundation models, recent progress in AI and machine learning has reached unprecedented complexity. For instance, the GPT-3 language model consists of 175 billion parameters and a training-data size of 570 GB. While it has achieved remarkable performance in generating text that is difficult to distinguish from human-authored content, a single training of the model is estimated to produce over 550 metric tons of CO_2 emissions. Likewise, we see advances in GeoAI research improving large-scale prediction tasks like satellite image classification and global climate modeling, to name but a couple. While these models have not yet reached comparable complexity and emissions levels, spatio-temporal models differ from language and image-generation models in several ways that make it necessary to (re)train them more often, with potentially large implications for sustainability. While recent work in the machine learning community has started calling for greener and more energy-efficient AI alongside improvements in model accuracy, this trend has not yet reached the GeoAI community at large. In this work, we bring this issue to not only the attention of the GeoAI community but also present ethical considerations from a geographic perspective that are missing from the broader, ongoing AI-sustainability discussion. To start this discussion, we propose a framework to evaluate models from several sustainability-related angles, including energy efficiency, carbon intensity, transparency, and social implications. We encourage future AI/GeoAI work to acknowledge its environmental impact as a step towards a more resource-conscious society. Similar to the current push for reproducibility, future publications should also report the energy/carbon costs of improvements over prior work.

Keywords. GeoAI Ethics, AI Sustainability, Carbon Cost, Transparency, Foundation Model

1 Introduction

The increasing availability of powerful hardware, such as GPUs and TPUs, has enabled the scaling of many computationally intensive tasks including training deep neural networks. With more capabilities and higher accuracy, recent machine learning models have seen an enormous increase in complexity compared to those developed a decade ago. For instance, large language models, such as GPT-3 (Brown et al., 2020) and PaLM (Wei et al., 2022), have hundreds of billions of parameters. The ever-increasing training and serving costs—in money, energy, and greenhouse gas (GHG) emissions—have raised concerns within the machine learning community. For example, the carbon cost of just training a BERT language model on GPUs without hyperparameter tuning is comparable to a trans-American flight (Strubell et al., 2019).

As machine learning becomes an invaluable tool in large-scale data analysis, we see a similar trend with GeoAI models getting increasingly computationally expensive, such as in the use of deep neural networks for weather forecasting (Pathak et al., 2022; Sønderby et al., 2020), satellite imagery classification (Jiao et al., 2018; Malof et al., 2017; Cong et al., 2022), traffic prediction (Cai et al., 2020; Yin et al., 2021), and so forth. Associated with these models are increasing financial costs, environmental costs, and social costs. It is interesting to note that models used to quantify sustainability issues or simulate the potential effects of measures to mitigate these issues (e.g., Heuvelink

et al., 2021; Mayfield et al., 2017) are, themselves, heavy energy consumers.

Calls to improve the energy efficiency and reduce the GHG cost of large computational models have become more frequent within the machine learning community. A variety of strategies have been put forward to quantify, report, and ultimately reduce the carbon footprint of modeling, discussed in later sections. We believe most are steps in the right direction and encourage their consideration within the GeoAI community. However, some proposed strategies entail ethical trade-offs regarding geographic disparities that have received little attention to date. In this work, we explain these trade-offs while encouraging the GeoAI community to prioritize sustainability alongside model performance. To start this discussion, we propose a framework (intended eventually to be refined into a numerical index) to evaluate models from several sustainability-related angles, including energy efficiency, carbon emissions, transparency, and social consequences. We hope that reporting on these angles will motivate our community to make resource-conscious decisions regarding model architecture design and evaluation.

The remainder of this paper is organized as follows. In Section 2, we briefly describe how attention to the energy and carbon costs of computing has grown, software tools intended to help reduce these costs, and how to take the full life cycle of hardware into account. We then propose a framework for model sustainability that incorporates a (geo)spatial and temporal perspective in Section 3. In Section 4, we discuss trade-offs between model improvement and sustainability, together with implications for geographic information. Finally, we summarize our findings and conclude with recommendations to the GeoAI community in Section 5.

2 Related Work

Growing awareness of climate change, and the role of greenhouse gas emissions driving it, have prompted many to consider the environmental consequences of computing. Early concerns focused on energy consumption of information and communication technologies in general (e.g., Hilty et al., 2009; Gelenbe and Caseau, 2015; Malmudin et al., 2013). Cryptocurrency mining has more recently popularized the topic in mainstream media (e.g., Gonzalez, 2022; Hinsdale, 2022; Schmidt and Powell, 2022). Concurrently, developers of AI and other computationally intensive modeling techniques in Earth system sciences (Loft, 2020; Fuhrer et al., 2018), computational biology (Lannelongue et al., 2021), precision medicine (Samuel and Lucassen, 2023), and other communities have increasingly questioned their role in climate change mitigation. Together, they have spawned initiatives such as Green AI

(Schwartz et al., 2020), Green Algorithms¹, and the Green Software Foundation².

Software tools have been developed for estimating, visualizing, and reporting operational energy use and carbon emissions associated with machine learning algorithms. For example, CodeCarbon (Lottick et al., 2019), Carbon Tracker (Anthony et al., 2020), Machine Learning Emissions Calculator Tool (Lacoste et al., 2019), EnergyVis (Shaikh et al., 2021), and experiment-impact-tracker (Henderson et al., 2020) analyze code, hardware, and in some cases the region where the code was executed to estimate the amount of energy used and carbon emissions generated. Often, such tools will predict how energy or carbon costs can be reduced by changing the hardware used or the geographic region where the code is executed. These tools are intended to simplify footprint estimation and reporting. At the same time, they promote resource-use efficiency as an evaluation standard within the machine learning community.

While the carbon costs of operating devices, or *operational emissions*, have grown somewhat in the past decade, they have been eclipsed by the carbon costs of producing devices, or *embodied emissions* (Gupta et al., 2022; Wu et al., 2022). Embodied emissions relate to the construction of hardware-manufacturing facilities, procuring raw materials, and the fabrication, assembly, packaging, and recycling of devices. Replacing or augmenting old hardware with newer, more powerful machines may increase operational efficiency but at the expense of a larger overall carbon footprint when considering operational and embodied emissions together (Wu et al., 2022; Lannelongue et al., 2021). Embodied emissions have been calculated to far outweigh operational emissions in contemporary computing (Gupta et al., 2022), yet embodied emissions are generally harder to estimate and have received less attention.

3 A Sustainability Framework for Model Evaluation

To promote resource awareness in GeoAI research, we propose a framework for evaluating model sustainability, eventually intended to become a multipart numerical index comprising a list of indicators. Just as authors of papers, submitted to AGILE and other conferences, are asked to provide their data and code to improve reproducibility, we hope that future papers will also report on sustainability indicators. We encourage sustainability to be part of the evaluation criteria rather than focusing solely on accuracy improvements over existing baselines, regardless of GHG costs. Fig. 1 provides an overview of the proposed components through three conceptual lenses, namely energy, social consequences, and transparency.

¹<https://www.green-algorithms.org/>

²<https://greensoftware.foundation/>



Figure 1. Lenses of the proposed sustainability framework.

Energy: (a) **Efficiency:** Where hardware specifications allow, the amount of energy consumed to train and re-train a model should be reported. If that is not possible, other indicators of energy consumption, like a combination of runtime and the hardware used, should be substituted. (b) **Carbon³ intensity:** When publicly available, the power-generation source should be disclosed, for example, "carbon-friendly" solar, hydro, or wind vs. "carbon-intensive" fossil fuels. If possible, an estimate of the carbon emissions generated as a result of the computations in the paper should be made available. Ideally, research would also report energy costs together with the carbon intensity of the data centers used so that they can be openly compared to alternative setups.

Social Consequences: For large-scale, computationally intensive models, include a statement evaluating potential risks and benefits to populations that may be affected. When outsourcing cloud computing jobs to carbon-friendly regions so as to reduce overall carbon emissions, we should also take into account the potential negative impact for the population in those regions. A risk-benefit assessment is adopted here to roughly estimate the social consequences.

Transparency: Building upon the reproducible research initiative of AGILE⁴, we envision a new submission guideline for transparency. In addition to data and software availability, one should include as much information as possible regarding the model setup, such as training time, sensitivity to hyperparameters, hardware used, spatio-temporal resolution (if applicable), and potential usage of pre-trained models.

Fig. 2 shows a hypothetical evaluation of a model according to the sustainability index outlined here. Each indicator (represented by an axis) is assigned a score from 0 to 5 and aggregated into a radar chart. Through this chart, we can quickly grasp the model's performance on each sustainability-related aspect. On the design of the index, we adopt a graphic format to visualize the score of each indicator. It can be used as a badge to recognize outstanding conference papers, similar to the reproducibility badge

³We focus on carbon as it accounts for the vast majority of GHG emissions.

⁴<https://reproducible-agile.github.io/initiative/>

proposed by AGILE and ACM⁵. In the following sections, the three main components are substantiated.




Figure 2. A hypothetical example illustrating how the proposed indicators might be visualized in an index.

3.1 Energy

To shrink the carbon footprint of computing, most proposed solutions focus on the energy needed to operate devices. Their recommendations broadly follow two strategies: (1) increase energy efficiency and (2) minimize the carbon intensity of the energy source (Loft, 2020). Regarding energy efficiency, modelers have been encouraged to estimate and report the amount of energy consumed to develop, train, tune, infer with, or otherwise run their models and consider the energy costs of storing and transmitting large amounts of data. Practices have been advocated at the algorithm, hardware, and operating facility level to improve efficiency, as summarized in Fig. 3. Proposed measures include selecting efficient machine learning architectures (Gupta et al., 2022; Patterson et al., 2021; Lottick et al., 2019); improving hyperparameter tuning and inference efficiency (Anthony et al., 2020); using pre-trained models (Yosinski et al., 2014); reducing floating point-precision computation (Loft, 2020); developing and

⁵<https://www.acm.org/publications/policies/artifact-review-and-badging-current>

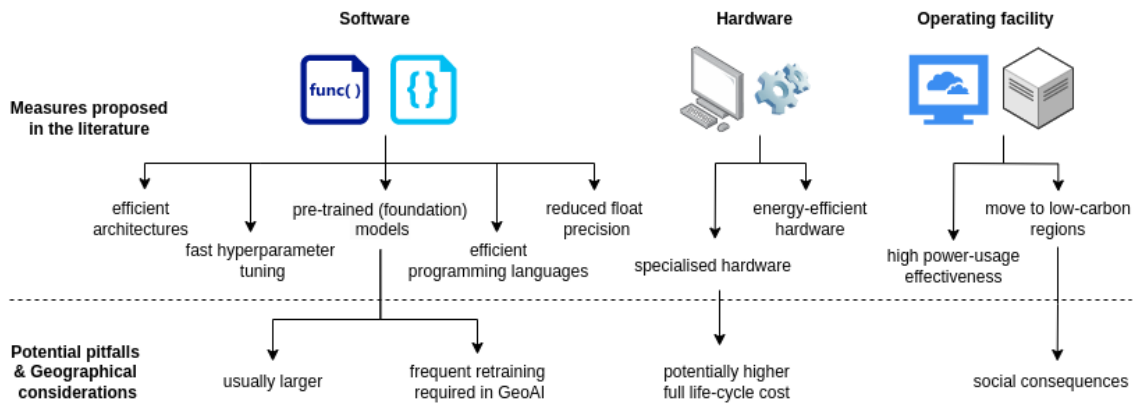


Figure 3. Proposed measures to improve the sustainability of AI models, and their potential general and GeoAI-specific pitfalls.

using more efficient programming languages, compilers, and code libraries (Gupta et al., 2022; Lanelongue et al., 2021); other strategies to modify algorithms in ways that increase their efficiency (Lacoste et al., 2019; Strubell et al., 2019; Wu et al., 2022); using hardware and settings that are energy efficient (Anthony et al., 2020; Patterson et al., 2021; Gupta et al., 2022; Henderson et al., 2020; Lacoste et al., 2019; Strubell et al., 2019; Shaikh et al., 2021; Lanelongue et al., 2021); and selecting cloud providers that operate facilities with an optimal power usage effectiveness (PUE) (Lacoste et al., 2019).

To minimize the carbon intensity of energy sources, recommendations include training and running models in geographic regions and at times of day that result in lower carbon emissions (Anthony et al., 2020; Patterson et al., 2021; Dodge et al., 2022; Lottick et al., 2019; Henderson et al., 2020; Lacoste et al., 2019; Shaikh et al., 2021; Lanelongue et al., 2021; Gupta et al., 2022; Wu et al., 2022). We argue that these measures are transferable to GeoAI to some extent but also come along with certain pitfalls that are shown in Fig. 3 and discussed below.

3.2 Social Consequences

Recognizing that energy production varies from region to region in terms of its carbon intensity, many have advocated shifting large computing jobs to carbon-friendly regions as a solution to reducing carbon emissions. Bender et al. (2021) note the geographic disparity between the people who typically benefit from large language models and those who are most vulnerable to the consequences of emitting carbon, i.e., climate change. Reducing carbon emissions slows the progression of climate change and, consequently, the risk exposure of communities most vulnerable to it. This argument recognizes the potential harm enacted through climate change to communities distant from those who benefit from the model. However, a second—largely ignored—mechanism by which harm can be done is through power generation itself. While outsourcing energy-intense computing to more carbon-friendly regions helps reduce carbon emissions globally,

it places a social and environmental burden on the place where the power is generated. Underlying social and environmental costs such as noise, biodiversity loss, appropriation of land and water resources that could be devoted to other uses, etc., are associated with any kind of power generation. For example, carbon-friendly energy sources like wind turbines can increase bird mortality and cause habitat loss (Marques et al., 2020; Smallwood, 2007). Whether through climate change or power generation, the consequences of computing—and their potential for geographic displacement—should be acknowledged.

To demonstrate current disparities, we compare the carbon intensity of production vs. consumption of selected countries in Table 1. Switzerland, for example, consumes more non-renewable energy than it generates, while Estonia is in the opposite situation, consuming less carbon-intensive energy than it produces. Such imbalances in energy imports and exports could increase with the growing importance of energy-intensive computing centers. Nowadays, the major cloud computing providers are taking action to reduce their carbon footprint by offsetting carbon emissions or investing in renewable energy. Even so, if the environmental burden outweighs the benefits of the computational task to the community where the power is generated, outsourcing power generation in this way could be considered an environmental injustice.

Table 1. An overview of carbon intensity in energy production and consumption in selected regions. Data are pulled from <https://app.electricitymaps.com/map> as yearly averages for 2022, in units of gCO_2eq/kWh .

Power Grid	Carbon Intensity (Production)	Carbon Intensity (Consumption)
North Sweden	13	16
Iceland	27	27
Switzerland	60	153
Estonia	607	390
Western India	622	600
Cyprus	949	949

Relocating computing centers to low-carbon regions will put local residents disproportionately at risk and is not the ultimate solution to AI sustainability. We should take into account whether the people there would benefit from the to-be-trained model and how much overlap there is between the population that benefits and the affected population. In the case of a global model, e.g., to predict the worldwide spread of COVID-19, shifting the model's training and deployment regions would not present the ethical concern we raise above, as the population affected by power generation would stand to benefit from the model, as well, if such benefits are distributed equally. However, a model that classifies road conditions in Boston should not necessarily be trained in Asia.

3.3 Transparency

While transparency and sustainability are often treated as separate values and principles (Jobin et al., 2019), we consider them to be intertwined if looking at them from a geographic perspective. Transparency, together with efficiency and accuracy, should be promoted as a crucial component of evaluation metrics. Information on model training time, sensitivity to hyperparameters, future fine-tuning needs, hardware requirements, spatio-temporal resolution and scope, where and when the energy powering computation was generated and consumed, etc., should be reported systematically, if possible. We acknowledge, however, that disaggregating and defining CO_2 emissions throughout a model's full life cycle remains a challenge (Luccioni et al., 2022). With such information, follow-up research and future adopters of the model will have a clear baseline from which to improve. To date, this discussion has not reached the GeoAI community at large. Therefore, we argue that available software for quantifying the energy and carbon footprint of machine learning models (Lottick et al., 2019; Anthony et al., 2020; Lacoste et al., 2019; Shaikh et al., 2021; Henderson et al., 2020) contributes an important step towards more transparency but should be extended to include social and geographical aspects.

4 Discussion

4.1 Performance vs. Resources: Trade-offs?

The trade-off between AI model performance and resources—e.g., hardware, energy—required to achieve that performance deserves scrutiny. For example, Geiping and Goldstein (2022) compared the BERT model's downstream performance to what they were able to achieve with only a single consumer-grade GPU and one day of training from scratch. This triggers more questions: how sensitive is a model to hyperparameter tuning for downstream tasks? How often does a model need to be retrained? Does a model require a long training time but no future fine-tuning or a relatively short training time with a constant

need for retraining? With the development of foundation models, these questions become even more critical as one advantage of using them is to answer questions relying on what is learned from the intensive pre-training process so that no more training is needed for new tasks. Therefore, pre-trained models may contribute to AI sustainability by potentially reducing GHG emissions.

4.2 The Spatio-temporal Resolution and Scope of Geographic Information

One distinguishing characteristic of geographic information is its resolution—or granularity, speaking more broadly—as one of the core concepts of spatial information proposed by Kuhn (2012). In contrast to language data, which can be reduced to a single character as its finest resolution, so to speak, geographic data can be collected and represented at infinitely fine spatial or temporal resolutions, at least theoretically. The resolution of data used in a model affects the computational resources needed. Beyond spatio-temporal resolution, we should also encourage authors of GeoAI work to improve model efficiency by critically selecting an appropriate spatio-temporal scope without compromising the results or introducing additional bias, e.g., representation bias (Liu et al., 2022). While a language model may be useful for years to come, a land-use model may need more frequent updating: language changes slowly (considering the large corpora used during training), while drivers behind the location of agricultural expansion, for example, change much faster (Versteegen et al., 2016). Our models in GeoAI may require an entirely new cycle of training, tuning, and deployment on a more frequent basis to remain relevant.

5 Conclusions

As machine learning models get increasingly powerful, we see tools like ChatGPT receiving unprecedented attention. However, for these large-scale, computationally intensive models, their energy consumption and associated carbon emissions are not as widely discussed.

In this paper, we echo the call for Green AI in the machine learning community and bring this issue to the attention of GeoAI research. In addition to transferring commonly used sustainability-evaluation metrics, such as efficiency and transparency, to the GeoAI field, we propose to further consider ethical factors related to the location of energy-consuming data centers. We should not simply suggest outsourcing computationally intensive training jobs to carbon-friendly regions in order to reduce the overall carbon footprint; the local residents in these regions might be negatively impacted by the underlying environmental and social costs of providing the power, even if it is carbon friendly. We include another spatio-temporal perspective and suggest that GeoAI models focus on improving efficiency by choosing their spatio-temporal reso-

lution and scope in relation to the affected vs. benefiting regions. Finally, we propose a sustainability framework for model evaluation that incorporates the aforementioned factors, hoping to foster GeoAI research that is more carbon conscious in the future.

We encourage researchers to systematically report on carbon emissions and information regarding the model setup, such as training time, sensitivity to hyperparameters, future fine-tuning needs, hardware and data center used, etc. More transparency can help follow-up research improve a model's efficiency, mitigate long-term carbon emissions, and promote social responsibility.

Finally, our work here is intended as a starting point for discussions, workshops, and community engagement in style similar to recent approaches that have brought multi-faceted ethical issues to our joint attention (Goodchild et al., 2022) before developing large-scale, resource-intensive GeoAI foundation models (Mai et al., 2022).

Acknowledgements

This work was partially funded via the KnowWhereGraph project (NSF #2033521).

References

- Anthony, L. F. W., Kanding, B., and Selvan, R.: Carbontracker: Tracking and predicting the carbon footprint of training deep learning models, arXiv preprint arXiv:2007.03051, 2020.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S.: On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?, in: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, pp. 610–623, 2021.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners, *Advances in neural information processing systems*, 33, 1877–1901, 2020.
- Cai, L., Janowicz, K., Mai, G., Yan, B., and Zhu, R.: Traffic transformer: Capturing the continuity and periodicity of time series for traffic forecasting, *Transactions in GIS*, 24, 736–755, 2020.
- Cong, Y., Khanna, S., Meng, C., Liu, P., Rozi, E., He, Y., Burke, M., Lobell, D. B., and Ermon, S.: SatMAE: Pre-training Transformers for Temporal and Multi-Spectral Satellite Imagery, in: *Advances in Neural Information Processing Systems*, 2022.
- Dodge, J., Prewitt, T., Tachet des Combes, R., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., and Buchanan, W.: Measuring the carbon intensity of ai in cloud instances, in: *2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1877–1894, 2022.
- Fuhrer, O., Chadha, T., Hoefler, T., Kwasniewski, G., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, C., Schulthess, T. C., et al.: Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0, *Geoscientific Model Development*, 11, 1665–1681, 2018.
- Geiping, J. and Goldstein, T.: Cramming: Training a Language Model on a Single GPU in One Day, arXiv preprint arXiv:2212.14034, 2022.
- Gelenbe, E. and Caseau, Y.: The impact of information technology on energy consumption and carbon emissions, *ubiquity*, 2015, 1–15, 2015.
- Gonzalez, O.: Bitcoin Mining: How Much Electricity It Takes and Why People Are Worried, <https://www.cnet.com/personal-finance/crypto/bitcoin-mining-how-much-electricity-it-takes-and-why-people-are-worried/>, 2022.
- Goodchild, M., Appelbaum, R., Crampton, J., Herbert, W., Janowicz, K., Kwan, M.-P., Michael, K., Alvarez León, L., Bennett, M., Cole, D. G., et al.: A white paper on locational information and the public interest, 2022.
- Gupta, U., Kim, Y. G., Lee, S., Tse, J., Lee, H.-H. S., Wei, G.-Y., Brooks, D., and Wu, C.-J.: Chasing carbon: The elusive environmental footprint of computing, *IEEE Micro*, 42, 37–47, 2022.
- Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., and Pineau, J.: Towards the systematic reporting of the energy and carbon footprints of machine learning, *The Journal of Machine Learning Research*, 21, 10 039–10 081, 2020.
- Heuvelink, G. B., Angelini, M. E., Poggio, L., Bai, Z., Batjes, N. H., van den Bosch, R., Bossio, D., Estella, S., Lehmann, J., Olmedo, G. F., et al.: Machine learning in space and time for modelling soil organic carbon change, *European Journal of Soil Science*, 72, 1607–1623, 2021.
- Hilty, L. M., Coroama, V., De Eicker, M. O., Ruddy, T., and Müller, E.: The role of ICT in energy consumption and energy efficiency, Report to the European Commission, DG INFSO, Project ICT ENSURE: European ICT Sustainability Research, Graz University, 1, 1–60, 2009.
- Hinsdale, J.: Cryptocurrency's Dirty Secret: Energy Consumption, <https://news.climate.columbia.edu/2022/05/04/cryptocurrency-energy/>, 2022.
- Jiao, J., Zhang, Y., Sun, H., Yang, X., Gao, X., Hong, W., Fu, K., and Sun, X.: A densely connected end-to-end neural network for multiscale and multiscene SAR ship detection, *IEEE Access*, 6, 20 881–20 892, 2018.
- Jobin, A., Ienca, M., and Vayena, E.: The global landscape of AI ethics guidelines, *Nature Machine Intelligence*, 1, 389–399, 2019.
- Kuhn, W.: Core concepts of spatial information for transdisciplinary research, *International Journal of Geographical Information Science*, 26, 2267–2276, 2012.
- Lacoste, A., Luccioni, A., Schmidt, V., and Dandres, T.: Quantifying the carbon emissions of machine learning, arXiv preprint arXiv:1910.09700, 2019.
- Lannelongue, L., Grealey, J., Bateman, A., and Inouye, M.: Ten simple rules to make your computing more environmentally sustainable, 2021.
- Liu, Z., Janowicz, K., Cai, L., Zhu, R., Mai, G., and Shi, M.: Geoparsing: Solved or Biased? An Evaluation of Geographic Biases in Geoparsing, *AGILE: GIScience Series*, 3, 9, 2022.

- Loft, R.: Earth system modeling must become more energy efficient, *Eos*, <https://doi.org/10.1029/2020EO147051>, 2020.
- Lottick, K., Susai, S., Friedler, S. A., and Wilson, J. P.: Energy Usage Reports: Environmental awareness as part of algorithmic accountability, arXiv preprint arXiv:1911.08354, 2019.
- Luccioni, A. S., Viguier, S., and Ligozat, A.-L.: Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model, arXiv preprint arXiv:2211.02001, 2022.
- Mai, G., Cundy, C., Choi, K., Hu, Y., Lao, N., and Ermon, S.: Towards a foundation model for geospatial artificial intelligence (vision paper), in: *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pp. 1–4, 2022.
- Malmodin, J., Bergmark, P., and Lundén, D.: The future carbon footprint of the ICT and E&M sectors, *On Information and Communication Technologies*, 12, 2013.
- Malof, J. M., Collins, L. M., and Bradbury, K.: A deep convolutional neural network, with pre-training, for solar photovoltaic array detection in aerial imagery, in: *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 874–877, IEEE, 2017.
- Marques, A. T., Santos, C. D., Hanssen, F., Muñoz, A.-R., Onrubia, A., Wikelski, M., Moreira, F., Palmeirim, J. M., and Silva, J. P.: Wind turbines cause functional habitat loss for migratory soaring birds, *Journal of Animal Ecology*, 89, 93–103, 2020.
- Mayfield, H., Smith, C., Gallagher, M., and Hockings, M.: Use of freely available datasets and machine learning methods in predicting deforestation, *Environmental modelling & software*, 87, 17–28, 2017.
- Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., et al.: Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators, arXiv preprint arXiv:2202.11214, 2022.
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., and Dean, J.: Carbon emissions and large neural network training, arXiv preprint arXiv:2104.10350, 2021.
- Samuel, G. and Lucassen, A. M.: The environmental impact of data-driven precision medicine initiatives, *Cambridge Prisms: Precision Medicine*, 1, e1, 2023.
- Schmidt, J. and Powell, F.: Why Does Bitcoin Use So Much Energy?, <https://www.forbes.com/advisor/investing/cryptocurrency/bitcoins-energy-usage-explained/>, 2022.
- Schwartz, R., Dodge, J., Smith, N. A., and Etzioni, O.: Green ai, *Communications of the ACM*, 63, 54–63, 2020.
- Shaikh, O., Saad-Falcon, J., Wright, A. P., Das, N., Freitas, S., Asensio, O., and Chau, D. H.: EnergyVis: interactively tracking and exploring energy consumption for ML models, in: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–7, 2021.
- Smallwood, K. S.: Estimating wind turbine-caused bird mortality, *The Journal of Wildlife Management*, 71, 2781–2791, 2007.
- Sønderby, C. K., Espeholt, L., Heek, J., Dehghani, M., Oliver, A., Salimans, T., Agrawal, S., Hickey, J., and Kalchbrenner, N.: Metnet: A neural weather model for precipitation forecasting, arXiv preprint arXiv:2003.12140, 2020.
- Strubell, E., Ganesh, A., and McCallum, A.: Energy and policy considerations for deep learning in NLP, arXiv preprint arXiv:1906.02243, 2019.
- Verstegen, J. A., Karssenbergh, D., van der Hilst, F., and Faaij, A. P.: Detecting systemic change in a land use system by Bayesian data assimilation, *Environmental modelling & software*, 75, 424–438, 2016.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., and Zhou, D.: Chain of thought prompting elicits reasoning in large language models, arXiv preprint arXiv:2201.11903, 2022.
- Wu, C.-J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Aga, F., Huang, J., Bai, C., et al.: Sustainable ai: Environmental implications, challenges and opportunities, *Proceedings of Machine Learning and Systems*, 4, 795–813, 2022.
- Yin, X., Wu, G., Wei, J., Shen, Y., Qi, H., and Yin, B.: Deep learning on traffic prediction: Methods, analysis, and future directions, *IEEE Transactions on Intelligent Transportation Systems*, 23, 4927–4943, 2021.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H.: How transferable are features in deep neural networks?, *Advances in neural information processing systems*, 27, 2014.