EMERALDS
Extreme-scale Urban Mobility Data Analytics as a Service

Ethics of Mobility AI Development

Anita GRASER
AIT Austrian Institute of Technology

EMERALDS Webinar

A Sustainability Framework for Model Evaluation

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

AI Carbon Footprint

Year	Model	Number of Parameters	CO_2 Emissions
2017	$\text{Transformer}_{base}$ [1]	65M	12 kg [2]
2019	BERT_{base} [3]	110M	652 kg [2]
2020	GPT-3 [4]	175B	552 mt [5]
2023	GPT-4 [6]	1T? 100T?	?

~ 50x more

~ 1000x more

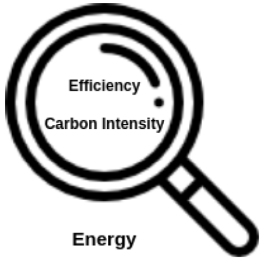

Table 1: Overview of recent large language models and their estimated carbon costs of training

Figure 1: The carbon cost of training a BERT language model is comparable to a round trip trans-American jet travel.

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

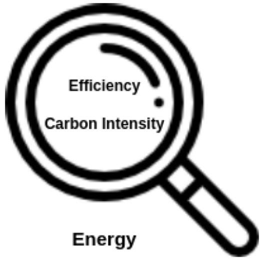

Red AI refers to AI research that improves accuracy through massive computational power while disregarding the costs.

Green AI refers to AI research that yields novel results while taking computational costs into account and encouraging reduced resource consumption.

- Green Algorithms [1]
- Green Software Foundation [2]
- Carbon Tracker [8]
- Machine Learning Emissions Calculator Tool [9]
- ...

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

Proposed Measures

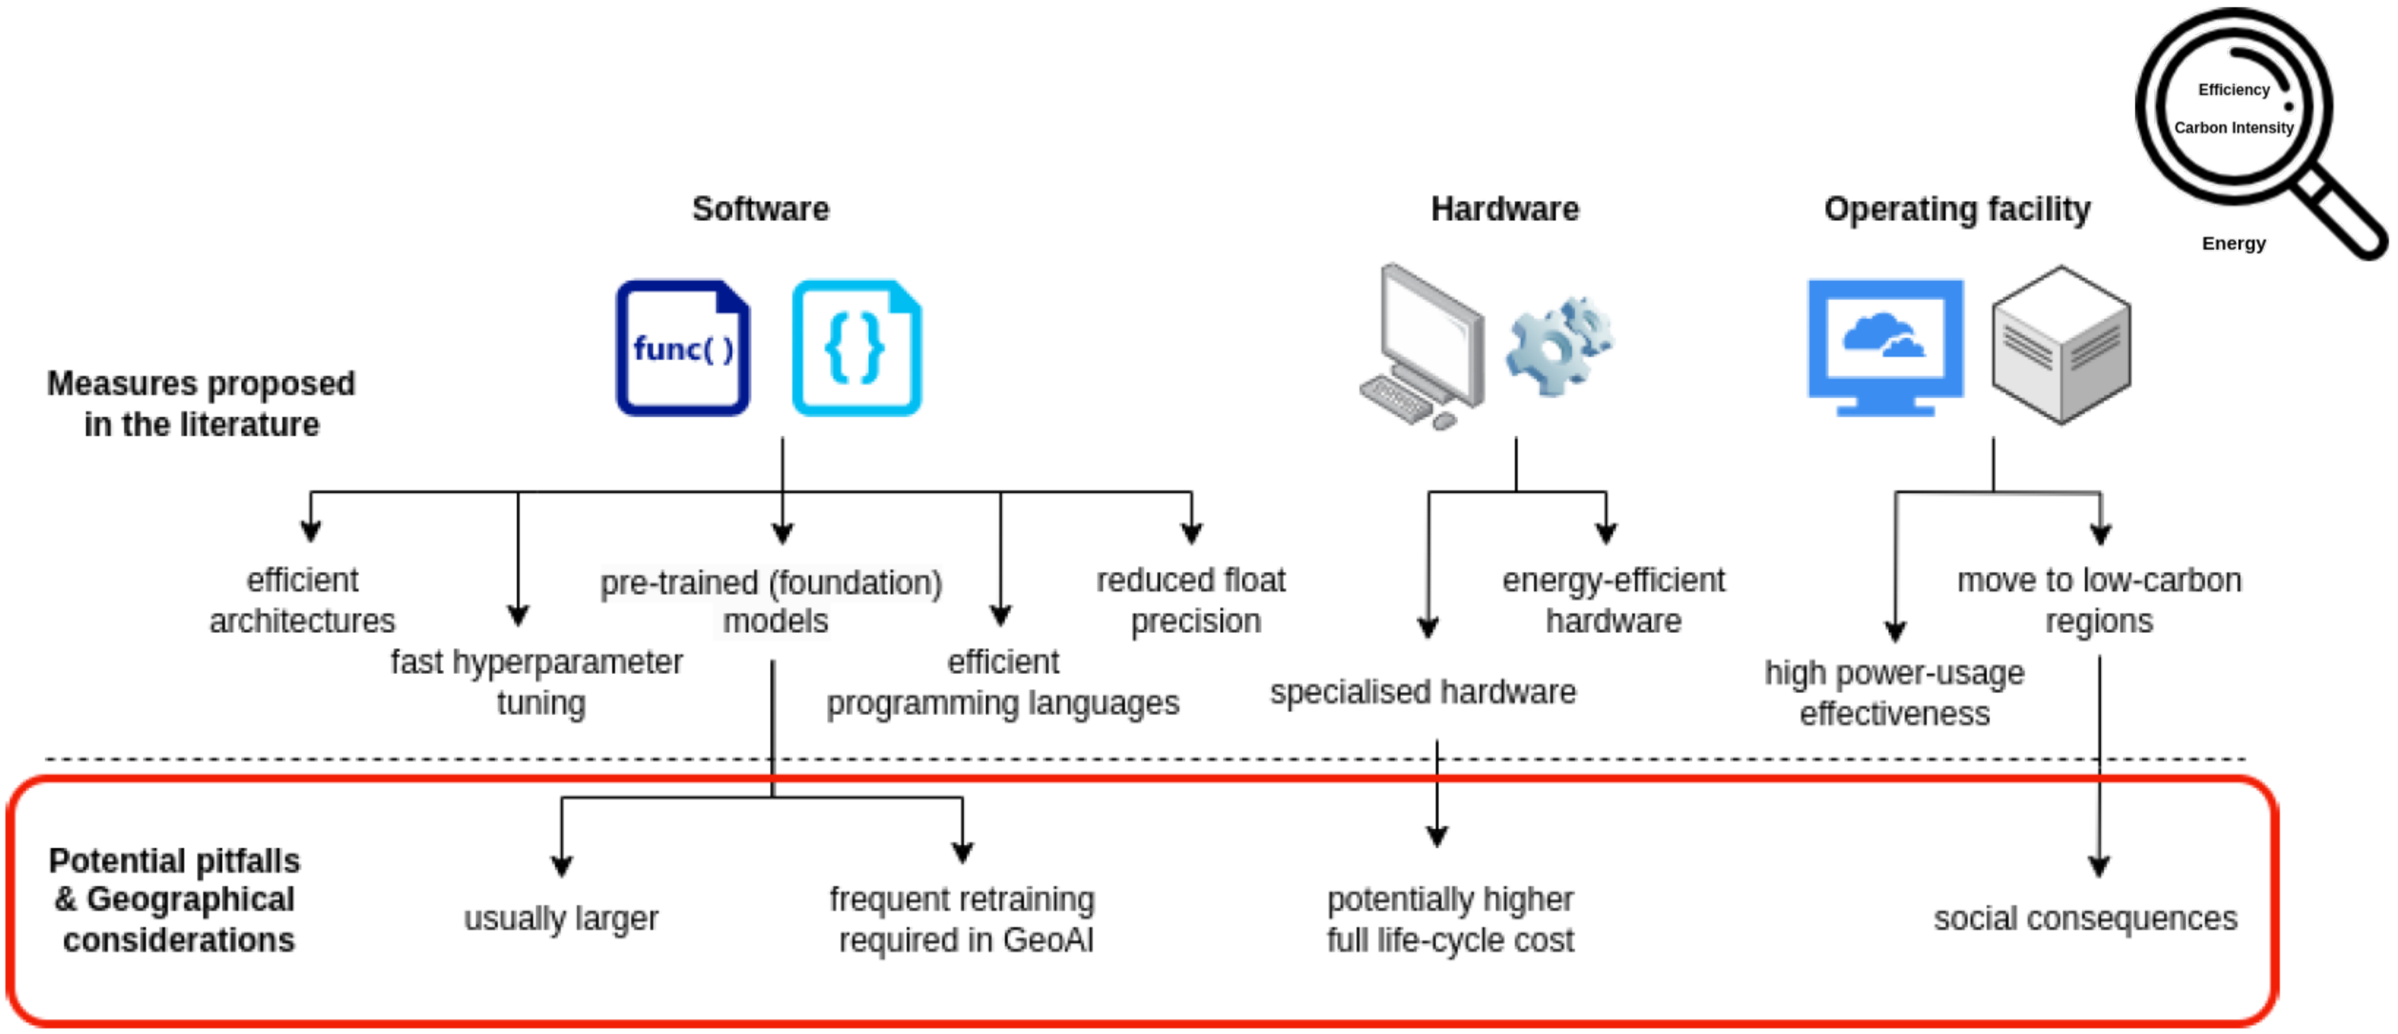

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

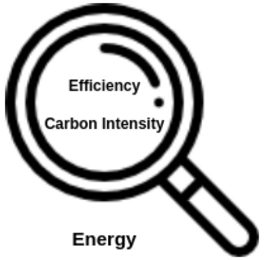

- Efficiency
  - Amount of energy consumed to train and re-train a model
  - Indicators of energy consumption, e.g., runtime, hardware used
- Carbon Intensity
  - Power-generation source, e.g., solar, hydro, wind, fossil fuels
  - An estimate of the carbon emissions generated
  - Carbon intensity of the data centers used

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

- Relocating computing centers to low-carbon regions may put local residents disproportionately at risk and is not the ultimate solution to AI sustainability.

- Take into account whether the people there would benefit from the model to be trained and how much spatial overlap there is between the population that benefits and the population that is affected.

→ evaluate potential risks and benefits to populations that may be affected

- Protecting privacy in mobility data is crucial as location information can reveal sensitive details like
  - home,
  - workplace, and
  - personal habits

- Challenges in keeping mobility data private include dealing with
  - patterns in location information,
  - understanding complex location meanings, and
  - addressing various uses like contact tracing

Mokbel, M., Sakr, M., Xiong, L., Züfle, A., Almeida, J., Aref, W., …, Graser, A., … & Zimányi, E. (2023). Towards Mobility Data Science (Vision Paper). arXiv preprint arXiv:2307.05717.

Social Consequences – Mobility Data Privacy

1. **Local**: Local differential privacy (LDP) allows individuals to protect their data before sharing it with services
   → Privacy at data collection stage

2. **Global**: Differential privacy (DP) is used to anonymize and share aggregate mobility data, ensuring your privacy even when data is collected on a larger scale
   → Commonly used to train machine learning models and create synthetic data, ensuring privacy while still allowing for useful analysis

Strong privacy protections are needed even for seemingly anonymous data, to prevent privacy breaches and misuse

Mokbel, M., Sakr, M., Xiong, L., Züfle, A., Almeida, J., Aref, W., …, Graser, A., … & Zimányi, E. (2023). Towards Mobility Data Science (Vision Paper). arXiv preprint arXiv:2307.05717.

- Mobility AI decisions should be understandable and NOT black boxes

e.g. LIME

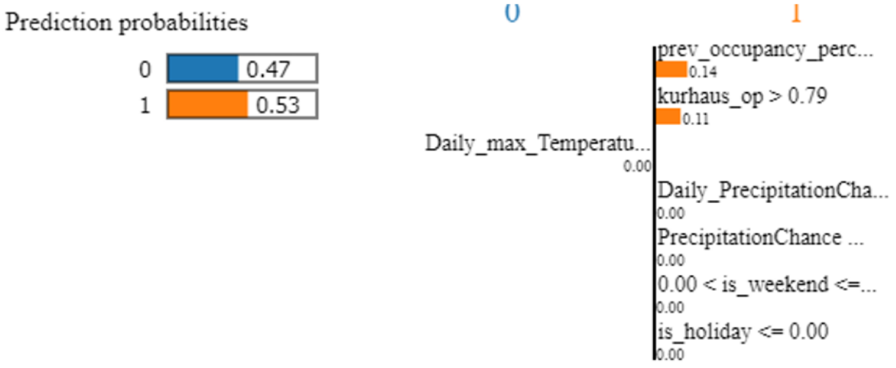

Feature	Value
prev_occupancy_percentage	1.00
kurhaus_op	1.00
Daily_max_Temperature	0.84
Daily_PrecipitationChance	0.00
PrecipitationChance	0.00
is_weekend	1.00
is_holiday	0.00

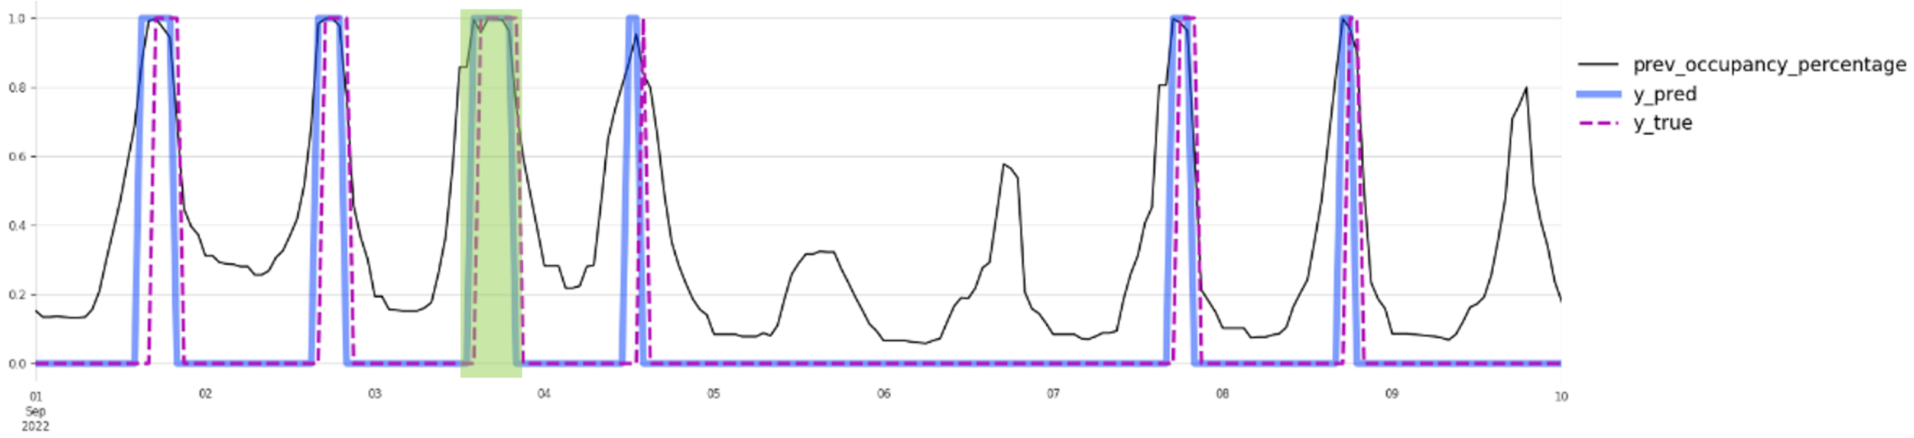

Jalali, A., Graser, A., & Heistracher, C. (2023). Towards eXplainable AI for Mobility Data Science. arXiv preprint arXiv:2307.08461.

In addition to data and software availability, also include:

- Training time
- Sensitivity to hyperparameters
- Hardware used
- Spatio-temporal resolution and scope (if applicable)
- Potential usage of pre-trained models
- ...

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

Performance vs. Resources: Trade-offs?

- How reusable is the model?
- How sensitive is it to hyperparameter tuning for downstream tasks?
- How often does it need to be retrained?
- Does it require a long training time but no future fine-tuning or a relatively short training time with a constant need for retraining?
- Is the contribution worth the resources consumed?

- Explainability approaches are resource-intensive too

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

- Mobility data needs strong privacy protections to prevent privacy breaches and misuse
- Mobility AI may require more frequent training, tuning, and deployment to remain relevant
- Ethical AI development needs to consider who profits and who has to carry the burden
- Mobility AI should be trustworthy and explainable

Shi, M., Currier, K., Liu, Z., Janowicz, K., Wiedemann, N., Verstegen, J., McKenzie, G., Graser, A., Zhu, R., and Mai, G. (2023) Thinking Geographically about AI Sustainability, AGILE GIScience Ser., 4, 42, https://doi.org/10.5194/agile-giss-4-42-2023.

EMERALDS

Extreme-scale Urban Mobility Data Analytics as a Service

Follow us

[1] A. Vaswani, N. Shazeer, N. Parmar, *et al.*, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[2] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in nlp," *arXiv preprint arXiv:1906.02243*, 2019.

[3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[4] T. Brown, B. Mann, N. Ryder, *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.

[5] D. Patterson, J. Gonzalez, Q. Le, *et al.*, "Carbon emissions and large neural network training," *arXiv preprint arXiv:2104.10350*, 2021.

[6] OpenAI, *Gpt-4 technical report*, 2023. arXiv: `2303.08774 [cs.CL]`.

[7] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green ai," *Communications of the ACM*, vol. 63, no. 12, pp. 54–63, 2020.

[8] L. F. W. Anthony, B. Kanding, and R. Selvan, "Carbontracker: Tracking and predicting the carbon footprint of training deep learning models," *arXiv preprint arXiv:2007.03051*, 2020.

[9] A. Lacoste, A. Luccioni, V. Schmidt, and T. Dandres, "Quantifying the carbon emissions of machine learning," *arXiv preprint arXiv:1910.09700*, 2019.