

The Green AI Ontology: An Ontology for Modeling the Energy Consumption of AI Models

Michael Färber*, David Lamprecht

Karlsruhe Institute of Technology (KIT), Institute AIFB, Germany

Abstract

Modeling AI systems' characteristics of energy consumption and their sustainability level as an extension of the FAIR data principles has been considered only rudimentarily. In this paper, we propose the *Green AI Ontology* for modeling the energy consumption and other environmental aspects of AI models. We evaluate our ontology based on competency questions. Our ontology is available at <https://w3id.org/Green-AI-Ontology> and can be used in a variety of scenarios, ranging from comprehensive research data management to strategic controlling of institutions and environmental efforts in politics.

Keywords

Machine Learning, Green AI, Energy Consumption, Ontology Engineering

1. Introduction

Pre-trained language models such as GPT have been commended for their artificial general intelligence capabilities and are nowadays widely used for tasks such as question answering, information extraction, and text summarization. However, in the case of GPT-3 with its 175 billion parameters, the training required 10,000 GPUs and cost 552 metric tons of carbon dioxide.¹ Thus, the question arises of how “green” AI models are. Regardless of an ethical assessment, we argue that it is useful to model AI systems' characteristics of energy consumption and sustainability (e.g., operating costs), extending the FAIR data principles [1], which focus on the availability and reuse of research data and other artifacts. Existing ontologies and knowledge graphs focus on the modeling of the research landscape, modeling publications, authors, and venues (e.g., FaBiO, ORKG, MAKG) [2]. Furthermore, ontologies for modeling software and neural networks have been proposed. For instance, the *Ontology for Informatics Research Artifacts* (OIRA) [3] provides a way to model software and datasets. In FAIRnets [4], the authors propose a schema for modeling neural networks. However, surprisingly, none of these ontologies allow the modeling of the energy consumption of AI models (e.g., runtime or CO2 footprint of pretrained language models, which can be measured via tools [5]).

21st International Semantic Web Conference, ISWC 2022, Virtual Event, October 23–27, 2022

*Corresponding author.

✉ michael.farber@kit.edu (M. Färber); david.lamprecht@student.kit.edu (D. Lamprecht)

ORCID [0000-0001-5458-8645](https://orcid.org/0000-0001-5458-8645) (M. Färber); [0000-0002-9098-5389](https://orcid.org/0000-0002-9098-5389) (D. Lamprecht)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://fortune.com/2021/04/21/ai-carbon-footprint-reduce-environmental-impact-of-tech-google-research-study/>

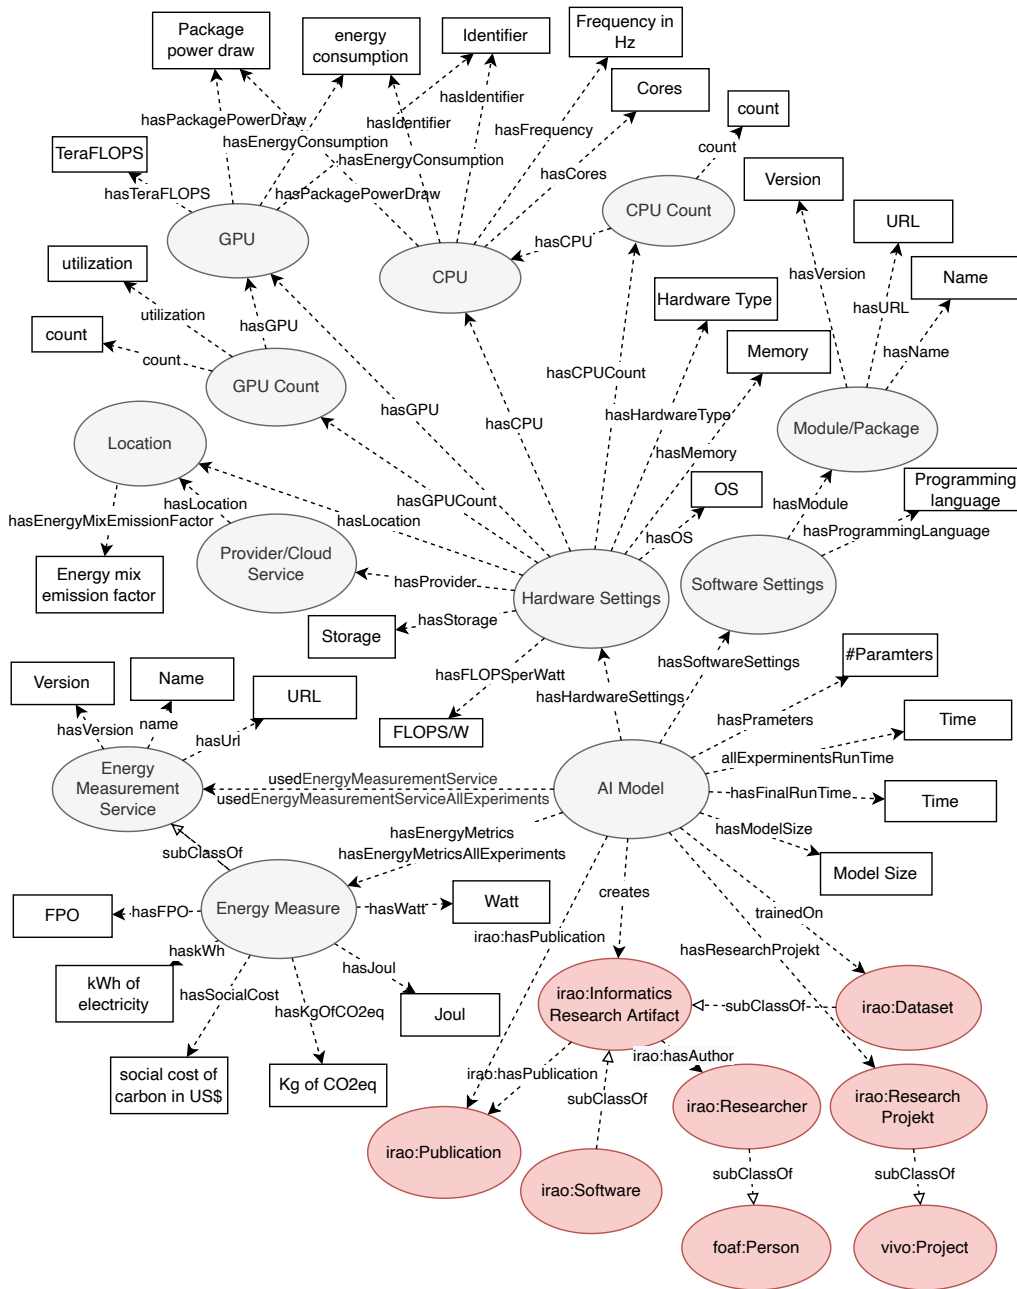


Figure 1: Main classes and properties of the *Green AI Ontology*.

In this paper, we propose – to our knowledge – the first ontology for modeling the energy consumption of AI models. It is available at <https://w3id.org/Green-AI-Ontology> (OWL file at <https://w3id.org/Green-AI-Ontology/ontology>). We create a knowledge graph based on our ontology and evaluate our ontology based on competency questions. Our ontology can be used in various scenarios, ranging from improved research data management to strategic controlling of institutions and implementation of standards.

2. The Green AI Ontology

Ontology Design. Figure 1 shows the main classes and properties of the ontology. The corresponding documentation is linked in our repository. Overall, our ontology is designed to model the following aspects:

1. *Metrics and tools:* This part addresses the metrics that are used to measure the energy consumption of AI models. Apart from the pure values (*Energy Measure*), we consider the online services (*Energy Measurement Service*) with which the values can be determined. *Energy Measurement Service* is defined as a subclass of *Energy Measurement* so that all relevant key figures are modeled in addition to information about the service.
2. *Hardware settings, cloud service, and location:* The property *hasHardwareSettings* shows information about the hardware used. In addition to the modeling of private infrastructures, services/cloud providers are represented here. The location of the hardware (e.g., city, country), which may have an impact on the environmental balance, can also be taken into account.
3. *Software settings:* The property *hasSoftwareSettings* shows information about the software used, including software packages and modules.
4. *Linking to scholarly linked data:* This part of the ontology is designed to integrate the modeled energy consumption information into the modeling of the scientific landscape. Since the AI models are closely linked to further computer science artifacts (e.g., datasets, software), we reuse the *Ontology for Informatics Research Artefacts* [3] considering best practice for reusing existing ontologies. As a result, the modeled information is not a silo but is closely linked to papers, data sets, and researchers. In this way, novel queries and strategic controlling are possible (e.g., answering: *What is the average energy consumption of AI models developed and trained at my institution over the last five years?*).

Knowledge Graph Construction. To create a knowledge graph based on our ontology, we first applied 10 regex patterns (an extension of [6]) on all 217,000 arXive computer science papers as of July 31, 2020 (from <http://unaxr.org>). In this way, we obtained 3,016 energy information units. However, we noticed that the accuracy of the matched patterns is insufficient due to the low precision of the information extraction approach. For instance, a large portion of the extracted energy information refers to non-AI models, such as mobile phones and e-cars.

Thus, we refrained from this approach and instead asked AI researchers via a questionnaire to report the energy consumption of AI models published in papers. In this way, we obtained a proof-of-concept knowledge graph, modeling 40 AI models and 1,975 statements.

Ontology Evaluation. Following the best practices of ontology engineering (e.g., the NeOn ontology engineering methodology), we identified 15 competency questions (see our repository; based on 79 Green AI-related papers that are listed in our repository) that our ontology should be able to answer. Based on our created knowledge graph and created SPARQL queries (see our repository and Listing 1), we were able to answer all competence questions.

Use Cases. In the following, we outline several potential use cases of our ontology.

Research Data Management. The FAIR principles [1] have been proposed to ensure that resources are findable, accessible, interoperable, and reusable. Our ontology can be considered

```
SELECT * WHERE { ?AIModel a gai:AIModel .  
                 ?AIModel gai:hasEnergyMetrics ?EnergyMetrics .  
                 ?EnergyMetrics gai:hasFPO ?FloatingPointOperations . }
```

Listing 1: SPARQL query answering “How many floating point operations (FPO) do the AI models have?”

an extension of these principles, allowing the modeling of usage information next to existing ontologies and knowledge graphs.

AI Systems. Engineers training and deploying AI models as well as end users may be increasingly interested in knowing the environmental background of given AI models [5] in order to assess them more thoroughly than merely for their effectiveness. Our modeling of the energy consumption of AI models is not restricted to one metric (e.g., CO₂, run time); instead, our ontology allows the modeling of several measurements for each AI model.

Society. From the perspective of popular sciences and politicians, our ontology complies with the rising public awareness of Green AI and environmental studies. The ontology enables energy consumption to be put into perspective (e.g., comparing energy consumption of language models and bitcoin mining).

3. Conclusion

In this paper, we proposed the *Green AI Ontology* for modeling the energy consumption of AI models. It can be used to extend academic knowledge graphs, to encourage researchers to provide information on the energy consumption of their AI models, and to ensure that the community appreciates this information.

References

- [1] M. D. Wilkinson, et al., The FAIR Guiding Principles for scientific data management and stewardship, *Scientific Data* (2016) 2052–4463.
- [2] V. B. Nguyen, V. Svátek, G. Rabby, Ó. Corcho, Ontologies Supporting Research-Related Information Foraging Using Knowledge Graphs: Literature Survey and Holistic Model Mapping, in: *Proc. of EKAW*, 2020, pp. 88–103.
- [3] V. B. Nguyen, V. Svátek, Ontology for informatics research artifacts, in: *Proceedings of the 18th Extended Semantic Web Conference, ESWC’21*, 2021, pp. 126–130.
- [4] A. Nguyen, T. Weller, M. Färber, Y. Sure-Vetter, Making neural networks FAIR, in: *Proceedings of the Second Iberoamerican Conference and First Indo-American Conference on Knowledge Graphs and Semantic Web, KGSWC’20*, 2020, pp. 29–44.
- [5] A. Lacoste, A. Luccioni, V. Schmidt, T. Dandres, Quantifying the Carbon Emissions of Machine Learning, *CoRR abs/1910.09700* (2019).
- [6] P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, J. Pineau, Towards the systematic reporting of the energy and carbon footprints of machine learning, *CoRR abs/2002.05651* (2020).