

Towards Sustainability Model Cards

Gwendal Jouneaux¹, Jordi Cabot^{1,2}

¹*Luxembourg Institute of Science and Technology, Esch-sur-Alzette, Luxembourg*

²*University of Luxembourg, Esch-sur-Alzette, Luxembourg*

Abstract

The growth of machine learning (ML) models and associated datasets triggers a consequent dramatic increase in energy costs for the use and training of these models. In the current context of environmental awareness and global sustainability concerns involving ICT, Green AI is becoming an important research topic. Initiatives like the AI Energy Score Ratings are a good example. Nevertheless, these benchmarking attempts are still to be integrated with existing work on Quality Models and Service-Level Agreements common in other, more mature, ICT subfields. This limits the (automatic) analysis of this model energy descriptions and their use in (semi)automatic model comparison, selection, and certification processes.

We aim to leverage the concept of quality models and merge it with existing ML model reporting initiatives and Green/Frugal AI proposals to formalize a Sustainable Quality Model for AI/ML models. As a first step, we propose a new Domain-Specific Language to precisely define the sustainability aspects of an ML model (including the energy costs for its different tasks). This information can then be exported as an extended version of the well-known Model Cards initiative while, at the same time, being formal enough to be input of any other model description automatic process.

Keywords

AI Models, Model Cards, Sustainability, Energy, Quality model, Domain-Specific Language

1. Introduction

The large adoption of AI technologies, increase in model complexity, and dataset size has led to a dramatic rise in the computational power required to run and train the models and their overall energy cost and sustainability impact. In the last few years, the question of the carbon footprint of AI models became a priority inside the research community. The seminal paper from Strubell *et al.* [1] analyzed the carbon impact for training four state-of-the-art NLP models. The resulting conclusion is that the carbon footprint for training and using AI models should be reduced.

While efforts to benchmark, monitor and fine-tune AI models in the context of Green AI has been made [2], this information is usually not readily available for model users. Approaches such as Model Cards [3] provide detailed information to model users, such as performance metrics or ethical concerns. Yet, few approaches bridge the gap and present sustainability information. The most notable is the recent AI Energy Score [4] proposed by HuggingFace. However, this energy score only provides information on inference energy consumption, disregarding carbon emission, water usage or training impact.

To address the aforementioned issue, this paper proposes the definition of Sustainability Model Cards. Sustainability Model Cards complement the already existing Model Cards [3] with the additional concern of sustainability of AI models. We propose a domain-specific language (DSL) to precisely define these cards and pave the way for future automated use of this information.

The rest of the paper is organized as follows. Section 2 reviews the state of the art on the modeling of quality models for AI, formalism to describe AI models and related artifacts, and existing approaches to make model users aware of sustainability aspects of the models. Sections 3 and Section 4 detail the Sustainability Model Cards and associated DSL, respectively. Finally, Section 5 discusses future work for Sustainability Model Cards in the form of a research roadmap, while Section 6 concludes the paper.

2nd Workshop on Green-Aware Artificial Intelligence, 28th European Conference on Artificial Intelligence (ECAI 2025), October 25–30, 2025, Bologna, Italy

✉ gwendal.jouneaux@list.lu (G. Jouneaux); jordi.cabot@list.lu (J. Cabot)

🌐 <https://www.gwendal-jouneaux.fr/> (G. Jouneaux)

>ID 0000-0003-1158-9335 (G. Jouneaux); 0000-0003-2418-2489 (J. Cabot)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. State of the art

Quality Assurances (QA) and related quality models have been identified as a challenge for current AI software research [5]. In the past years, the research community proposed multiple quality models [6, 7]. However, most of those models do not present a full picture of AI software quality. Among the twenty-nine papers studied by Gezici *et al.* [6], only three discuss sustainability as a relevant quality aspect [8, 9, 10].

In addition to quality models, other formalisms to describe AI models or related artifacts have been developed. For datasets, Dataset Cards [11] allows defining standard information such as provenance, authorship, license, and tasks for which the dataset is suitable. DescribeML [12] additionally describes social concerns potentially leading to bias and provides a dedicated language and tool support for its specification. Finally, Croissant [13] uses a JSON notation that is both human-readable and compatible with existing tools and frameworks, providing additional interoperability, portability, and discoverability to datasets. For AI models, Mitchell *et al.* proposed Model Cards [3] used for model reporting. Model Cards describe the AI model in terms of intended use, data, performances and ethical considerations, among other things. However, none of those approaches allows the description of the sustainability concern.

On the other hand, approaches such as the ones from Hugging Face try to specifically assess and describe sustainability aspects of machine learning models. They extended the Model Cards approach with sustainability data [14] such as cloud provider location, training time, hardware, and estimated carbon emissions. They also created the AI Energy Score Ratings [4] to evaluate the energy cost when using a model. This method compute the average inference cost (in Wh) over one thousand requests.

However, these approaches neither provide a formal description that could be used to check the syntactic and semantic correctness of the card information nor an easy way to automatically process such information as part of a MLOps pipeline that should take into account energy concerns [15]. Furthermore, while carbon emissions of the training phase and energy consumption of the inference task are important metrics for model selection, our DSL covers a larger variety of energy-related information, such as the water consumption or counter measures taken by the platforms (and platform providers) to counter such energy impact. Finally, the use of DSL technology facilitates reusing many other existing tools in the DSL realm to simplify the creation of new functionality (e.g. verbalization of the card information) around this new Sustainability Model Card ecosystem.

3. Sustainability Model Cards

To address the aforementioned problem, we propose Sustainability Model Cards, inspired by the well-known Model Cards [3] concept. In what follows, we describe the dimensions that are part of our Sustainability definition based on an analysis of the relevant literature in this field. More specifically, we group Sustainability information of a ML model into four main sections: Metadata, Training, Inference, and Platform.

Sustainability Model Card			
Metadata	Training	Inference (for each task)	Platform
– Name	– Training Duration	– Inference Type	– Hardware Details
– Version	– Energy Consumption	– Energy Consumption	– Platform Provider
– Type	– Carbon Emissions	– Carbon Emissions	– Platform Region
– Provider	– Water Consumption	– Water Consumption	– Carbon Offset Credit
– License	– Platform	– Platform	– Energy Sources

The **Metadata** section contains information about the model itself. This includes the identification of the model through a name and version, the type of the model (e.g., decision tree, CNN, regression), identification of the provider and license of the model. This information allows establishing the link between the Sustainability Model Cards and already existing Model Cards for the same model to later do some combined analysis. The model type is particularly important, as studies such as Yu *et al.* shows potential correlation between model type and inference energy consumption [16].

The **Training** section regroup information concerning the environmental impact of the training phase of the model. More important aspects are the energy consumption, carbon emission and water consumption resulting from this training phase. While carbon emissions and energy consumption are already studied in the context of Green AI [1], water consumption is often overlooked, yet it is still regarded as an important sustainability aspect for datacenters [17]. To complement this information, this section also includes a reference (represented italicized) to the definition of the platform used, and the time spent to train the model (Which could be used to infer some energy consumption a posteriori with a certain degree of approximation).

The **Inference** section allows defining the impact of the different inference tasks a model can be used for. For each task supported by the model (e.g., text generation, text summarization), this section reports the inference task type, the average energy consumption, the average carbon emission, the average water consumption, and a reference (represented italicized) to the platform used to compute these estimations.

Finally, the **Platform** section describes the platforms used to train or execute the model. This description includes details about the hardware used, the region (e.g., Azure and EU-west), and the energy mix used for the training (the ratio of renewable and fossil energy used), as these aspects are useful when choosing a deployed model infrastructure. In addition, carbon offsetting has become a common practice for platform provider to reduce their environmental impact. Carbon offset credits that represent the quantity of greenhouse gas emission avoided or countered by other means (e.g., financing renewable energy projects). In the Sustainability Model Cards, carbon offset credits can be described either as a quantity (in $kgCO_2eq$) offsetting the platform emission, or as a percentage of the emissions offset afterward, as typically done by cloud platform providers such as Amazon, Google or Microsoft.

4. Description of a DSL for ML Sustainability reporting

This section describes the proposed domain-specific language (DSL) to support the definition of Sustainability Model Cards, including all dimensions mentioned in the previous section. A DSL is a language specially designed to model apps for a certain domain (health, finance, or, such as in this case, AI).

Once described with our DSL, the Sustainability information can be automatically processed as part of a ML workflow or MLOps approach (e.g., to search for sustainable alternatives, compare the sustainability of equivalent models, analysis of a set of models with a energy perspective, ...).

Each DSL has two main components:

- Abstract syntax: Describing the structure of the language and the way different language primitives can be combined, independently of any particular representation.
- Concrete syntax. Describes one or more specific notations for the language, covering the encoding and/or the visual/textual appearance of the elements in the abstract syntax

As we will see, our DSL uses a YAML based concrete syntax to facilitate the adoption of the language and its integration with the existing Model Cards of the Hugging Face repository.

4.1. Abstract Syntax of the Language

Figure 1 presents the abstract syntax of the proposed DSL in the form of a metamodel (i.e. the schema or grammar of the language expressed using an object-oriented perspective) structuring the language elements, the properties of every element and the possible relationships among them.

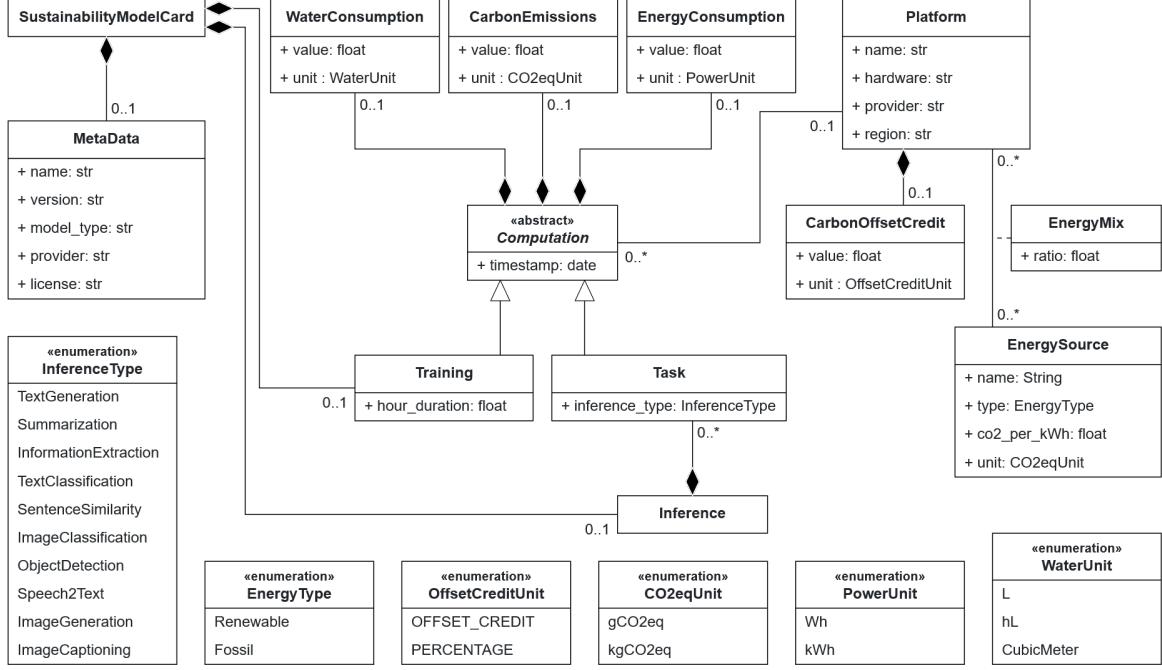


Figure 1: Metamodel of the Sustainability Model Cards DSL

Let's describe in more details the different elements of the language. The `SustainabilityModelCard` class is the root concept representing the whole Sustainability Model Card. This card is composed of three subcomponents: `MetaData`, `Training` and `Inference`.

The `MetaData` class represents the metadata section of the card defining the name, version, model type, provider, and license as strings.

The `Training` class models the training section of the card and defines the duration of the training. The `Inference` class encompasses all the inference tasks (represented by the `Task` class) addressed by the model. Each `Task` defines the inference type, based on the list used in the AI Energy Score [4]. Both the `Training` and `Task` classes represent computations that have an environmental impact. This is materialized through their inheritance from the `Computation` abstract class. This class models the impact of the computation when executed on a given `Platform`. The impact is materialized through associations to water consumption, carbon emissions and energy consumption metrics. These three metrics are reified in their own class and are represented through a value and its associated unit. Finally, the `Computation` a timestamp denoting the moment of the measurements.

In addition, a `Computation` is associated to a `Platform`. When reporting `Training` impact, the `Platform` represents the infrastructure used to train the model. While for `Task` impact, the `Platform` represents the infrastructure used when computing the metrics. The `Platform` class is characterized by a name, hardware description, provider and compute region. The name is used by computations to reference their platform in the textual notation. The hardware description, provider and compute region are used to describe the hardware used and its provider (*i.e.*, local or cloud provider). For a more fine-grained representation of the platform carbon impact, the `Platform` is associated to a `CarbonOffsetCredit` and a set of `EnergySource`. The `CarbonOffsetCredit` represent either the quantity of greenhouse gas emission mitigated in number of credits (one credit is 1000kgCO₂eq), or the percentage of the emissions mitigated afterward. The energy sources used by the platform are represented in the form of an energy mix, denoting the ratio of energy provided by a source. The `EnergySource` is characterized by the type of energy and carbon efficiency.

While `Training` and `Task` classes are mostly identical, we explicitly make a distinction based on their semantic difference and intent to have a more fine-grained description of the training phase, as presented in Section 5.

4.2. Concrete Syntax of the Language

The concrete syntax used to specify the Sustainability Model Card is based on YAML. A YAML structure is represented as a set of key-value pairs, where the value can be of three types: Scalar, Sequence, and Mapping. Scalar values are represented as a series of zero or more characters. Sequence values are represented as a list of values. Finally, mapping values are represented as a set of key-value pairs.

To encode Sustainability Model Cards in this format, we defined a set of rules.

1. Class instances are represented using the class name in snake case as key and a mapping as value
2. Attributes are part of instances mapping and represented using the attribute name in snake case as key and a scalar as value
3. Compositions are defined by having the composed class instance nested in the containing class mapping
4. Multiplicity higher than one is managed using YAML sequence
5. Simple associations are defined using the associated class name in snake case as key and the object name attribute as scalar value

In addition, there are three special cases to these rules: Platform and EnergySource, Inference, and EnergyMix. Even if platforms and energy sources are not direct components of the Sustainability Model Card, they are defined as a list contained in the card (See Listing ?? *platforms* and *energy_sources* lists). As Inference only contains the list of tasks, the intermediate “*task*” attribute containing the list is bypassed, making the YAML *inference* section a sequence. Finally, the EnergyMix association class is represented as a sequence of EnergyMix class instances containing the ratio as attribute and the EnergySource class instance.

```
1 sustainability_model_card:
2   meta_data:
3     name: GPT-3 175B
4     model_type: LLM
5     provider: OpenAI
6     platforms:
7       - platform:
8         name: Infrastructure
9         hardware: Multiple V100
10        provider: Microsoft Azure
11        region: US
12        carbon_offset_credit:
13          value: 100.0
14          unit: PERCENTAGE
15        energy_mix:
16          - energy_mix:
17            ratio: 100.0
18            energy_source: Azure US
19        energy_sources:
20          - energy_source:
21            name: Azure US
22            type: Fossil
23            co2_per_kWh: 0.3496 0.4759
24            unit: kgCO2eq
25
26   training:
27     platform: Infrastructure
28     carbon_emissions:
29       value: 449935.2
30       unit: kgCO2eq
31     energy_consumption:
32       value: 1287000
33       unit: kWh
34     water_consumption:
35       value: 5439
36       unit: CubicMeter
37     timestamp: 2025-01-02T09:00:00
38
39   inference:
40     - task:
41       inference_type: TextGeneration
42       platform: Infrastructure
43       carbon_emissions:
44         value: 1.398
45         unit: gCO2eq
46       energy_consumption:
47         value: 4
48         unit: Wh
49       water_consumption:
50         value: 0.017
51         unit: L
52     timestamp: 2025-01-21T09:00:00
```

Listing 1: Sustainability Model Card of a GPT-3 175B model trained and used on Microsoft Azure infrastructure based in the United States

To demonstrate Sustainability Model Cards, we provide Listing 1 as a concrete example using our DSL syntax to report the sustainability aspects of the GPT-3 model with 175 billion parameters.

The infrastructure described is using Microsoft Azure and multiple V100 GPUs as described in the related paper [18]. Data on water and energy consumption of the model comes from the study conducted by Li *et al.* [19] on GPT-3 training water consumption. These metrics were based on the average consumption across Azure datacenters in the US. Consequently, we used the average carbon efficiency in the US reported in eGRID 2023 Data provided by the US Environmental Protection Agency. Based on the consumption and carbon efficiency, we were able to compute the carbon emission for both the training phase and the text generation inference.

By processing this card, we observe that the inference carbon emissions exceed the training emissions after roughly 325 million inferences. In a context where the number of expected inferences across the model lifecycle is vastly superior to this number, this model could be compared to other based on the inference emissions, as its impact on the overall emission is greater.

4.3. Tool support

To support the formal definition of Sustainability Model Cards, we provide a Python implementation of our DSL. This implementation is composed of a validating parser and a set of classes implementing the metamodel, and is available in open-source on GitHub¹.

To implement the metamodel, we relied on the BESSER Low-Code platform [20] to create the DSL metamodel and generate a Python implementation of it. The Python classes generated can then be used to instantiate any model conforming to the specified metamodel, allowing its manipulation or creation by other tools. For instance, BESSER offers both a language to specify and generate implementation of neural networks [21], and a deployment language [22]. In the future, BESSER could use this infrastructure to automatically benchmark specified neural networks and generate their Sustainability Model Card as part of the design and execution process of the network.

In addition, we have implemented a parser validating and transforming the YAML description to model instances. First, we use an existing YAML parser implementation, transforming the textual description into a manipulable Python object. Then, we traverse the object structure to assess the conformance of the structure to the metamodel. These validation checks ensure: (1) the presence of units when required, (2) the correspondence of these units to the ones defined in the metamodel, (3) the correspondence of the inference and energy types to the ones defined in the metamodel, and (4) that the values representing percentage are bound to the [0,1] interval. Finally, the validated structure is transformed to a corresponding model instance using the metamodel classes.

5. Research Roadmap

Our DSL is a first step towards a more ambitious goal towards the adoption, formalization, analysis and improvement of AI sustainability. In this section, we discuss a research roadmap and possible next steps for the evolution of the Sustainability Model Cards initiative, including a number of application scenarios. We hope to extend and prioritize this list based on discussions with the community.

Extending the coverage and granularity of the Sustainability DSL. When creating the Sustainability Model Cards and associated DSL, we aimed to be as complete as possible by making the union of sustainability concepts mentioned in the surveyed papers [1, 23, 14, 4, 24, 16, 25, 26, 17]. Yet, more details could be added such as a more granular description of the training phase diving in the pre-training and fine-tuning part of the training, the hyperparameters values, or the dataset used. Furthermore, the Sustainable AI research field is highly active and propose new frameworks and new metrics that should be included in Sustainability Model Cards along the way. As expressed in the recent article of Cruz *et al.* [15], there is a need for standardized metrics in the evaluation of AI sustainability, and our framework should evolve to match (and possibly, influence) these upcoming metrics.

¹Implementation: <https://www.gwendal-jouneaux.fr/SustainabilityModelCards-Parser>

Graphical notation. When using a language, different users prefers different syntaxes, also depending on their technical profile. For instance, less technical users tend to prefer more graphical notations. For this reason, we plan to extend the language with other concrete syntaxes, such as a graphical notation, and even a conversational-based one, for input and a verbalization in Controlled English as output to help all types of users to write and read sustainability cards.

Tighter integration with Model Cards. Sustainability Model Cards focus only on the sustainability aspect of AI models. One of the ways to have a complete view of the model data would be through the integration of this card with other existing cards. A first step in this direction has already been made, as our DSL concrete syntax has been built on YAML to integrate seamlessly with the existing Hugging Face model cards. The next step in this direction would be to propose a formal description language integrating both cards, allowing automatic processing using all the available information. Another direction would be to extend the sustainability concern with ethical concern to allow more transparency in the ethical and environmental sustainability impact of AI models [27].

Analyzing impact on model users. An additional aspect of these cards is the social aspect. The final choice of using the model with the most accuracy, the least carbon emissions, or something in the middle belongs to the model user. Conducting a user study on how they decide on a model, especially among models that offer similar features and performance, would allow assessing the impact of providing more sustainability data in the choice of model users.

Application on different scenarios. The precise models of the Sustainability Model Cards resulting from the use of our DSL allow for the automatic processing of the information reported in the cards. This can be useful for a range of different scenarios, including many MLOps ones. For instance, a first scenario envisioned is to perform automatic model selection based on the model environmental impact. Another scenario would be to optimize model deployment based on an impact analysis using information on locations' energy providers and carbon efficiency, energy consumption of inference, and hardware. Finally, the provided information could be used and monitored at runtime to enforce sustainability aware Service Level Agreements (SLA) that could be established between the users and the model providers as typically done for other types of IT services following existing quality models.

6. Conclusion

In this paper, we have presented the Sustainability Model Cards, a new DSL allowing the description of AI models sustainability aspects including energy consumption, carbon emissions and water consumption. This DSL enables the formal definition of this information and facilitates the automatic processing of sustainability information as part of a MLOps pipeline while still exporting it as an extension of the Models Card formalism for better readability and integration with similar initiatives.

As further work, we plan to work on the aspects discussed as potential future roadmap for this initiative.

Acknowledgments

This work has been funded by the European Union under the Grant Agreement No 101189664 (MOSAICO project). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them.

Jordi Cabot is supported by the Luxembourg National Research Fund (FNR) PEARL program, grant agreement 16544475

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

- [1] E. Strubell, A. Ganesh, A. McCallum, Energy and policy considerations for deep learning in NLP, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, 2019, pp. 3645–3650. doi:10.18653/v1/P19-1355.
- [2] R. Verdecchia, J. Sallou, L. Cruz, A systematic review of green ai, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 13 (2023) e1507.
- [3] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, T. Gebru, Model cards for model reporting, in: *Proceedings of the conference on fairness, accountability, and transparency*, 2019, pp. 220–229.
- [4] HuggingFace, AI Energy Score, 2025. URL: <https://huggingface.github.io/AIEnergyScore>, [Online; accessed 26. May 2025].
- [5] M. Felderer, R. Ramler, Quality assurance for ai-based systems: Overview and challenges, in: *Software Quality: Future Perspectives on Software Engineering Quality: 13th International Conference, SWQD 2021, Vienna, Austria, January 19–21, 2021, Proceedings* 13, Springer, 2021, pp. 33–42.
- [6] B. Gezici, A. K. Tarhan, Systematic literature review on software quality for ai-based software, *Empirical Software Engineering* 27 (2022) 66.
- [7] M. A. Ali, N. K. Yap, A. A. A. Ghani, H. Zulzalil, N. I. Admodisastro, A. A. Najafabadi, A systematic mapping of quality models for ai systems, software and components, *Applied Sciences* 12 (2022) 8700.
- [8] L. Pons, I. Ozkaya, Priority quality attributes for engineering ai-enabled systems, *arXiv preprint arXiv:1911.02912* (2019).
- [9] J. Siebert, L. Joeckel, J. Heidrich, K. Nakamichi, K. Ohashi, I. Namba, R. Yamamoto, M. Aoyama, Towards guidelines for assessing qualities of machine learning systems, in: *Quality of Information and Communications Technology: 13th International Conference, QUATIC 2020, Faro, Portugal, September 9–11, 2020, Proceedings* 13, Springer, 2020, pp. 17–31.
- [10] J. Horkoff, Non-functional requirements for machine learning: Challenges and new directions, in: *2019 IEEE 27th international requirements engineering conference (RE)*, IEEE, 2019, pp. 386–391.
- [11] HuggingFace, Dataset Cards, 2025. URL: <https://huggingface.co/docs/hub/en/datasets-cards>, [Online; accessed 26. May 2025].
- [12] J. Giner-Miguel, A. Gómez, J. Cabot, A domain-specific language for describing machine learning datasets, *Journal of Computer Languages* 76 (2023) 101209.
- [13] M. Akhtar, O. Benjelloun, C. Conforti, L. Foschini, J. Giner-Miguel, P. Gijsbers, S. Goswami, N. Jain, M. Karamousadakis, M. Kuchnik, et al., Croissant: A metadata format for ml-ready datasets, *Advances in Neural Information Processing Systems* 37 (2024) 82133–82148.
- [14] HuggingFace, CO2 Emissions and the Hugging Face Hub: Leading the Charge, 2022. URL: <https://huggingface.co/blog/carbon-emissions-on-the-hub>, [Online; accessed 26. May 2025].
- [15] L. Cruz, J. P. Fernandes, M. H. Kirkeby, S. Martínez-Fernández, J. Sallou, H. Anwar, E. B. Roque, J. Bogner, J. Castaño, F. Castor, A. Chasmawala, S. Cunha, D. Feitosa, A. González, A. Jedlitschka, P. Lago, H. Muccini, A. Oprescu, P. Rani, J. Saraiva, F. Sarro, R. Selvan, K. Vaidhyanathan, R. Verdecchia, I. P. Yamshchikov, Greening ai-enabled systems with software engineering: A research agenda for environmentally sustainable ai practices, 2025. URL: <https://arxiv.org/abs/2506.01774>. arXiv:2506.01774.
- [16] J.-R. Yu, C.-H. Chen, T.-W. Huang, J.-J. Lu, C.-R. Chung, T.-W. Lin, M.-H. Wu, Y.-J. Tseng, H.-Y. Wang, Energy efficiency of inference algorithms for clinical laboratory data sets: green artificial intelligence study, *Journal of Medical Internet Research* 24 (2022) e28036.
- [17] B. Ristic, K. Madani, Z. Makuch, The water footprint of data centers, *Sustainability* 7 (2015) 11260–11284.
- [18] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh,

D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, 2020. URL: <https://arxiv.org/abs/2005.14165>. arXiv:2005.14165.

- [19] P. Li, J. Yang, M. A. Islam, S. Ren, Making ai less "thirsty": Uncovering and addressing the secret water footprint of ai models, 2025. URL: <https://arxiv.org/abs/2304.03271>. arXiv:2304.03271.
- [20] I. Alfonso, A. Conrady, A. Sulejmani, A. Nirumand, F. Ul Haq, M. Gomez-Vazquez, J.-S. Sottet, J. Cabot, Building BESSER: an open-source low-code platform, in: International Conference on Business Process Modeling, Development and Support, Springer, 2024, pp. 203–212.
- [21] N. Daoudi, I. Alfonso, J. Cabot, Modelling neural network models, in: International Conference on Research Challenges in Information Science, Springer, 2025, pp. 130–139.
- [22] F. Ul Haq, I. Alfonso, A. Sulejmani, J. Cabot, Extending a low-code tool with multi-cloud deployment capabilities, in: European Conference on Software Architecture, Springer, 2024, pp. 39–46.
- [23] E. Strubell, A. Ganesh, A. McCallum, Energy and policy considerations for modern deep learning research, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020, pp. 13693–13696.
- [24] A. Guldner, S. Kreten, S. Naumann, Exploration and systematic assessment of the resource efficiency of machine learning, in: INFORMATIK 2021, Gesellschaft für Informatik, Bonn, 2021, pp. 287–299.
- [25] L. F. W. Anthony, B. Kanding, R. Selvan, Carbontracker: Tracking and predicting the carbon footprint of training deep learning models, arXiv preprint arXiv:2007.03051 (2020).
- [26] E. García-Martín, C. F. Rodrigues, G. Riley, H. Grahn, Estimation of energy consumption in machine learning, Journal of Parallel and Distributed Computing 134 (2019) 75–88.
- [27] A. S. Luccioni, G. Pistilli, R. Sefala, N. Moorosi, Bridging the gap: Integrating ethics and environmental sustainability in ai research and practice, 2025. URL: <https://arxiv.org/abs/2504.00797>. arXiv:2504.00797.