Available online at www.sciencedirect.com

ScienceDirect

Procedia CIRP 135 (2025) 924–929

www.elsevier.com/locate/procedia

32nd CIRP Conference on Life Cycle Engineering (LCE 2025)

Life Cycle Assessment of Artificial Intelligence Applications: Research Gaps and Opportunities

Christiane Plociennik[a], Ponnapat Watjanatepin[b], Karel Van Acker[b,c], Martin Ruskowski[a]

[a] German Research Center for Artificial Intelligence (DFKI), Trippstadter Str. 122, 67663 Kaiserslautern, Germany
[b] Sustainable Materials Processing and Recycling (SeMPeR), Department of Materials Engineering, KU Leuven, Kasteelpark Arenberg 44, 3001 Leuven, Belgium
[c] Center for Economics and Corporate Sustainability (CEDON), KU Leuven, Warmoesberg 26, BE-1000 Brussels, Belgium

* Corresponding author. Tel.: 49 631 20575 3417. *E-mail address:* christiane.plociennik@dfki.de

Abstract

Applications based on Artificial Intelligence (AI) are increasingly being adopted, especially since tools like ChatGPT or Dall-E have become available to the public and have gained popularity. While its benefits are manifold, AI also has a significant environmental impact: AI technology usually relies on hardware which consumes a significant amount of energy, such as GPU servers. Furthermore, users normally access AI-based services via their own end devices, which also consume energy. The energy side of training and using AI models can be roughly estimated today, for example using libraries like CodeCarbon. However, a realistic estimate of AI's environmental impact needs to include not only the energy consumption from the use of the AI, but also factors such as the (critical) raw materials needed for the construction of the hardware, and the consumption of materials and energy for the maintenance and use of the hardware throughout its life cycle. This is not feasible today. Systematic life cycle assessment of AI technology is still in its infancy. It should not be limited to calculating the CO2-equivalent, but take into account different relevant environmental impact indicators from resource use depletion to human toxicity to water consumption. This paper reviews the literature on this topic, identifies gaps and opportunities for research and argues for developing an LCA methodology for AI. It further presents an adaptation of an existing LCA methodology to AI as a first step in this direction.

Keywords: Life Cycle Assessment, Life Cycle Analysis, Environmental Impact, Sustainability, Artificial Intelligence, LCA, AI

1. Introduction

Due to the fact that Artificial Intelligence (AI) is gaining huge popularity and is becoming ubiquitous, its relationship to sustainability is more and more in the focus. In the last years, two perspectives with regard to sustainability and AI have emerged: Employing *AI for sustainability* and evaluating the *sustainability of AI* [1]. Many studies address *AI for sustainability*, showing how AI can benefit sustainability in different areas, e.g. the fight against climate change [2]. Regarding the *sustainability of AI*, more and more research is being published as well, often with a focus on the environmental cost of AI. Most of these studies, however, have limitations: Concerning the environmental dimension, most studies focus on the energy consumption of AI, or its carbon footprint, respectively [3]. This, however, is not comprehensive since other indicators such as abiotic depletion potential (ADP) contribute signifcantly to AI's environmental impact as well [4]. Regarding the AI life cycle – which consists of the phases depicted in Figure 1 – many studies limit themselves to the learning/training phase of AI [3]. This view is also incomplete since, depending on the indicator considered, inference can have an even bigger impact than training [4].

Life cycle analysis, also called life cycle assessment, is an established tool to calculate the environmental impact of applications [5]. However, it has rarely been applied to AI yet, and there is currently no LCA methodology tailored to AI [6]. This paper makes the following contributions: First, it identifies pioneering work in the field of LCA of AI. Second, it lists the specific characteristics of AI to motivate why a dedicated LCA methodology for AI would be beneficial. Third, it shows how an existing LCA methodology can be adapted for AI applications. Our work can serve as a discussion basis to develop an LCA methodology completely tailored to AI.

10.1016/j.procir.2025.01.079

This paper is structured as follows: Section 2 lists related work, Section 3 describes the methodology used. Section 4 argues for an LCA methodology for AI and lists key points that should be considered when developing such a methodology. Section 5 proposes an adaptation of an LCA methodology to AI. Section 6 concludes the paper.

2. Related Work

Dokic et al. performed a systematic literature review and found out that there is currently no LCA methodology specifically for AI [6]. They furthermore identified the main factors for the negative environmental impact of AI during its life cycle: energy mix, timing of model training, efficiency of algorithms, hardware settings, model accuracy, and data center-related impacts.

Bouza et al. [3] presented a review study of different existing carbon calculators used to estimate the energy usage of AI as well as the conversion between the consumed wattage to carbon emissions based on the carbon intensity. Although this approach can be useful for a rough estimation of the carbon emissions per unit to construct the life cycle inventory (LCI), it is not a full LCA since only the energy consumption is estimated and the tools are largely black boxes. It is not known what is included in the impacts generated nor which impact assessment methodology was used to calculate them. Hence, it is not fully aligned with the standardized LCA methodologies described in ISO 14040 [7] and ISO 14044 [8] nor the ITU-T standard [9].

Existing LCA methods can be adapted to AI: Ligozat et al. adapted the ITU-T and ETSI standards for LCA of information and communication technologies [9] to AI [10]: They proposed to map tasks in AI development such as learning and inference to physical equipment that is needed for these tasks. Equipment can, for instance, be sensors, computers, supercomputers, and mobiles. An LCA can then be performed by considering the life cycle stages of the equipment (raw material acquisition, production, use, and end of life) and calculating emissions such as pollution or abiotic resources depletion for each life cycle stage. For estimating second-order impacts, Ligozat et al. assume that each AI application can be related to a corresponding reference application without AI. However, since AI enables new applications and creates new demand (see Section 4), this may be difficult in some cases.

Berthelot et al. used ITU-T [9] and Bordage et al. [11] as a basis and performed an LCA of the generative AI service Stable Diffusion [4]. They did not limit themselves to energy consumption or carbon footprint as an indicator, but calculated three different impact categories that Bordage et al. identified as the main impact categories for digital services: abiotic depletion potential (ADP), global warming potential (GWP), and primary energy (PE). They considered two functional units: a single use of the service (FU1) and the cost of the service for one year (FU2). They furthermore looked at the entire AI process (not only training) and considered all devices required during this process, including end-user terminals. This study is by far the most detailed and advanced concerning LCA of AI,

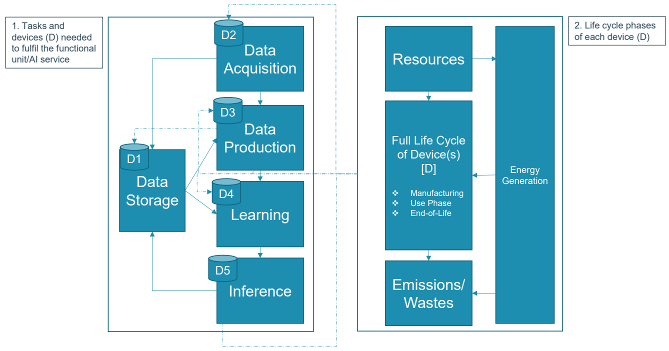

Fig. 1. AI life cycle (based on [10])

but it still has limitations: Several parameters used in the calculation are rough estimates. This includes impacts related to Graphical Processing Units (GPUs), which are estimated based on a methodology usually applied to Central Processing Units (CPUs) by the Boavizta working group [12]. Since GPU operation is usually more resource-intensive than CPU operation, the estimate is likely an underestimate. Average utilization rates (AURs) of equipment are rough estimations, too. Furthermore, as stated above, the main impact categories used here were originally devised for digital services. AI services might have additional or different impact categories not considered here. The study is based on the LCA principles and the equations presented can be used to construct estimates for the LCI. However, the impact assessment methodology used to calculate the impacts is not presented. This can hamper the reproducibility aspects mentioned in the ISO standards.

Both Ligozat et al. and Berthelot et al. state that they propose attributional LCA [13], but consequential would be preferable due to the complexities of the AI life cycle.

3. Methodology

The methodology for this paper employs a review of existing literature which is based on the systematic literature review (SLR) conducted by Dokic et al. [6]. This SLR lists the relevant literature on the topic of LCA and AI until 2023. We analyze all publications in Dokic et al.'s list. To update this with the newest relevant literature, we furthermore perform a forward-backward search [14], which yields the additional publications [4, 3, 15]. Overall, 51 publications related to LCA and AI are analyzed with regard to the research question: *"Which gaps exist in the current practice of environmental impact assessment of AI?"* The results are presented in the following section.

4. Towards a comprehensive environmental impact assessment of AI

This section presents the results of the literature review along with conclusions that follow from them in the form of theses.

A dedicated LCA methodology for AI is needed. AI has characteristics that differ from conventional information technology systems. AI has a specific process called the *AI life cycle*, not to be confused with the product life cycle (see Figure 1). It can involve a large number of different, heterogeneous devices [10]. Training and inferencing require not only CPU, but also GPU and in some cases TPU (Tensor Processing Unit) support. These devices have a different environmental impact compared to CPUs, the full extent of which can only be roughly estimated at the moment due to missing data [4]. Furthermore, this hardware is typically operational all the time, but it is not being used all the time. This is an important aspect since idle time can be an even bigger factor in terms of energy consumption than usage time [3]. End-user terminals must be considered as well. These may be battery-powered, which means they can have a different impact profile: The manufacturing phase typically dominates a battery-powered device's emissions, whereas always-connected devices produce more emissions via their energy consumption in the usage phase [16]. Moreover, AI needs a large amount of data, which must be acquired, transmitted, stored, processed and, in many cases, curated by a human. All of this requires resources, sometimes even additional devices [10].

Furthermore, AI has created new demand. For example, people write text using ChatGPT or generate images using Dall-E. This is a use of technology that is new – it simply was not possible before. The *way* people use this technology is also quite new: People often prompt AI systems multiple times because the first try is not good enough. This contributes to AI's environmental impact and should be taken into account in LCA. A similar phenomenon exists in the training phase of AI models: Here, AI practitioners often train a model several times with a slightly modified setup because the accuracy of the first run is not sufficient. Hence, emissions from each training run should be cumulated. Moreover, once a model has been released, it is often updated (cf. YOLOv1 – YOLOv7). The question is whether and to what extent the environmental impact of later versions should incorporate that of earlier versions.

Existing LCA methodologies for information technology [9, 11] do not take into account the specifics of AI detailed above. To account for these specifics, it would be beneficial to develop a dedicated LCA methodology for AI. In Section 5, we will show how an existing LCA methodology can be adapted for AI as a first proposal in this direction. However, the methodological side is but one part of the process to be addressed. This is a highly interdisciplinary field, and there are further open issues, e.g. concerning terminology and tooling. These are discussed in the following paragraphs.

AI and LCA experts should agree on a common terminology. In the future, AI practitioners will have to justify the resource usage of the models they develop. Hence, an LCA methodology for AI should not only be understandable for LCA experts, but must be easy to use for AI practitioners as well. This is not an easy task since it involves different fields that had little or no connection in the past. We propose to start with a common terminology. At the moment, this terminology is not standardized and can be confusing for either side. For example, existing studies use different terminology for the same or very similar concepts, which makes it hard to decide for non-experts whether these concepts are comparable – cf. the use of *average utilization rate* (AUR, [4]) vs. *average usage factor* [3]. Furthermore, AI and LCA in some cases use the same terms with a different meaning: the *use phase* in LCA denotes the life cycle phase of usage of the equipment. AI practitioners often call the use phase the *inferencing phase* of an AI model, which is called the *application phase* by Ligozat et al. [10]. Even the term *life cycle* itself is ambiguous: In LCA, it denotes all the phases of a product, in AI it denotes all the phases an AI model goes through (see Figure 1). These, in turn, are called *tasks* by Ligozat et al.

Tooling must be improved. As we have argued, measuring the energy consumption, or the carbon footprint, of AI is not enough. This is, however, the only indicator for which at least *some* tooling exists. There are tools like CodeCarbon [17], CarbonTracker [18], Green-Algorithms [19], eco2AI [20], Experiment-Impact-Tracker [21], MLCO2 [22] and Cumulator [23]. Most of them are cumbersome to use. They measure largely different things, they rely on different methodologies and assumptions, and their output is not comparable. In the process of measuring energy consumption, many parameters have to be roughly estimated because they cannot be measured directly. One example is the carbon intensity of the energy mix. The data sources for estimating this vary from tool to tool, leading to largely differing estimates. If measurements are based on wattmeters, usually the energy consumption of the whole computing node is measured, without taking into account how many and which processes are running on that node. This makes it difficult to allocate realistic portions of the energy consumption to a specific process. Another issue is the average usage factor. Some tools assume 100 % for this parameter if it is not available (Green-Algorithms, MLCO2). This leads to an overestimation of energy consumption. Consequently, it is almost impossible to judge which tools provide realistic output and which do not. Outputs can differ by as much as factor 4 between tools [3]. In addition, there are severe usability and accessibility issues. For access to energy consumption data of CPUs, several tools require read access to so-called RAPL files. This read access is restricted to root users. Users on supercomputers, however, usually do not have root permissions. Another issue is that RAPL files are only available under Intel CPUs and under Linux OS. When using virtual environments, it is often not even possible to find out the model of the CPU one is using. This, however, is needed by some tools. Concerning GPUs, some tools only work for Nvidia GPUs. Furthermore, tools are often buggy. One tool, Cumulator, does not even measure energy consumption – it is necessary to reverse-engineer this based on the carbon footprint. In summary, the tooling for measuring environmental impacts of AI should be improved in terms of accuracy, comparability, and usability, and it should accomodate indicators other than energy consumption or carbon footprint.

Categorization and standardization would streamline LCA of AI. Different studies employ different dimensions in the LCA of AI (see Figure 2). Ligozat et al. propose to carry out LCA along two dimensions: *life cycle phases* of the devices used for AI and *environmental indicators* [10]. Berthelot et al. use *equip-*

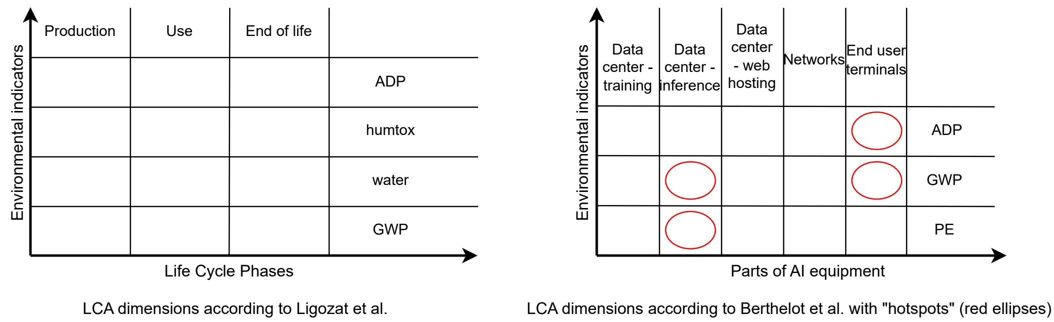

Fig. 2. LCA dimensions employed in different studies

ment and *environmental indicators* as the dimensions; life cycle phases are not explicitly included [4]. This makes different studies hard to compare. It would be desirable to standardize this as much as possible. Not all AI applications are the same, but it could be beneficial to categorize AI applications with respect to LCA. Then LCA methods could be developed for each category such that all applications in one category can be analyzed using the same LCA method. Developing these categories is a research effort of its own. One could start with proto-categories like *Large Language Models (LLMs)* or *Object Detectors*, or by grouping AI applications based on the set of devices they require. In the process, it might be possible to identify the most important elements (e.g. indicators, life cycle phases, equipment categories) for each AI application category with respect to environmental impact. Having such “hotspots” could mean that less influential indicators could be ignored to simplify the LCA process. Consider the right part of Figure 2: The red ellipses mark such hotspots identified by Berthelot et al [4]. Another potential hotspot: 75 % of emissions of end-user terminals occur during production [10]. Concerning system boundaries, Ligozat et al. consider data acquisition in their LCA process, while for Berthelot et al. it is outside the system boundaries. Hence, to improve comparability, it might be beneficial to standardize the system boundaries for each AI application category.

Different impacts should be considered. As mentioned, many studies limit themselves to looking at the energy consumption of training an AI model. To evaluate its full environmental impact, we must look at the entire AI process, including end user requests, inferencing, and possibly retraining, and we must take into account different environmental indicators, not just energy consumption or CO2 equivalent.

That said, it is not sufficient to look at the compute-related impacts of AI; the immediate application impacts and the system-level impacts [24] must be considered as well. This is due to the fact that AI nowadays permeates every walk of our lives. It has the potential to transform technology, human behavior, and society, and rebound effects become probable [10]. This can go along with fundamental shifts in material flows. For example, less profitable mines with a worse environmental impact might have to be opened to cater to the need for more lithium or cobalt. These are 3rd order effects that must be taken into account, which is beyond the scope of the studies already published on LCA of AI. Hence, a consequential LCA methodology [13] would be preferable over a purely attributional one [10]. However, at the moment, consequential LCA is not mature enough to be used in practice.

Consequently, since sustainability has 3 dimensions (ecologic, economic and social [25]), developing just an LCA methodology for AI is likely not comprehensive as it focuses on the ecologic dimension. For the economic dimension there is the concept of Life Cycle Costing Analysis (LCCA), and the social dimension is covered by Social LCA (S-LCA) [6]. For a complete sustainability evaluation of AI applications, it should be assessed whether the existing LCCA and S-LCA can be applied to AI applications or whether they need to be adapted.

Data quality must be improved. The existing data gaps make LCA of AI cumbersome and its results inaccurate. For example, reliable data concerning the GWP of GPUs and TPUs [10] is missing, and average utilization rates (AURs) must be roughly estimated since the big techs do not disclose them for the hardware they use. For models like GPT-4, information on the hardware that was used for training or training data is not available. To improve this situation, data gaps must be identified systematically. Then tech companies must increase their transparency concerning the data that is missing. If they refuse to do so, legislators should define which information needs to be disclosed.

5. Holistic Life Cycle Assessment of AI

As discussed in Section 2, there are no specific guidelines for LCA of AI. Therefore, this section shows how a comprehensive life cycle assessment of AI can be performed that is fully aligned with the ISO 14040 and 14044 standards and covers the ITU-T standard for LCA as well. Our proposal encompasses the four steps of LCA and covers what is missing in the discussed studies. The first step is illustrated in Figure 1 and the full process is depicted in Figure 3. As previously mentioned, the necessary foundational work to develop application-specific LCA methodologies for AI has not yet been completed. Hence, the approach proposed here is relatively generic and serves as an initial proposal. The four steps are as follows:

1. Goal and Scope Definition. In this step, the goal of the LCA as well as the scope of the study must be defined. The scope also sets the stage for the system boundary of the assessment – defining which processes are included in the assessment and

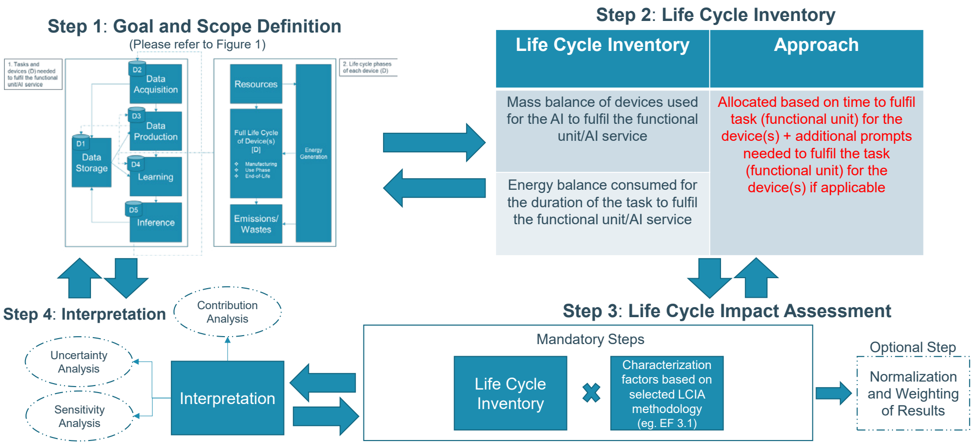

Fig. 3. Proposition for Life Cycle Assessment of AI

which are not. The functional unit (FU) of the assessment must be defined. The FU describes the function that the system must fulfil and provides the basis for the construction of the LCI. Many combinations of FU can be chosen according to the goal of the study, but the authors propose to include a functional unit based on the task to be performed by the AI. Examples are (depending on the type of AI application): detecting the objects in a given image or producing a text on a given subject.

2. Life Cycle Inventory (LCI). In this second step, the LCI must be constructed to match the FU defined in the first step. The LCI is constructed by collecting the mass and energy balance of all inputs and outputs that the system must perform to fulfil the chosen FU. Therefore, we propose to collect the inventory for both the input/output flows of the AI while it performs the task defined by the FU as well as to collect the data of the devices needed to perform each task in the life cycle of the AI. It is essential that the mass and energy balances obtained for the construction of the LCI should also include the full balances related to possible multiple prompts of the AI in order to fulfil its tasks/FU (see Figure 3). It is often logical to acquire the inventory based on the full life cycle of the device. However, to not overestimate the impacts, allocation of the impacts of the device should only be performed based on the time needed to perform the defined FU. While system expansion is preferred over allocation according to the ILCD handbook [26], the application of such an expansion is unlikely in the application field of AI.

3. Life Cycle Impact Assessment (LCIA). In this third step, the constructed LCI can be translated into environmental impacts by means of the characterization factors. Here, the choice of software or calculations as well as the life cycle impact assessment methodology depends on the LCA practitioner. However, the impact categories reported must match the goal and scope of the study defined in the first step. This way, the LCA model and the selected LCIA methodology can be used to transform the LCI into impacts. As an example, the use of the European Commission's Environmental Footprint 3.1 (EF 3.1) results in the calculation of 16 environmental impact categories – a wide range of indicators beyond carbon emissions/global warming potential [27]. The impact categories in EF 3.1 also comprise those recommended in the ITU-T standard. An optional normalization and weighting step is also included in Step 3, being fully in line with the process flow of an LCA of the ISO standards. This step can be introduced based on the questions the LCA practioners want to answer. Normalization uses factors to convert calculated impacts into relative values based on reference points, such as per capita, nation, area or baseline scenarios, aligning with the LCA goal and scope. For example, the EF 3.1 methodology provides normalization factors based on average yearly emissions per capita, which can help identify the most impactful categories for the system under study. Weighting is also available in case a specific need is to be addressed based on the quantified environmental impacts by giving preferences or additional weights to each impact category.

4. Interpretation. In this final step, the environmental impact categories can be further investigated using contribution analysis, uncertainty analysis, and sensitivity analysis. In combination with the optional normalization and/or weighting steps in Step 3, this process can therefore provide practical usage of the calculated LCA results in order to understand the impacts associated with the studied system as well as possibilities to improve its environmental sustainability. To fully align with the ISO 14040 standard series, it is important to remember that LCA is an iterative process, which can be repeated with better data or if the goal/scope of the assessment needs refinement. Additionally, iterations are possible between the different steps.

6. Conclusion and Future Work

In this paper, we have argued for a more comprehensive environmental impact assessment of AI: Due to the specific characteristics of AI, current LCA methodologies are not entirely suited for AI. In addition, issues like terminology and tooling must be addressed as well. Hence, we propose to develop a dedicated LCA methodology for AI. We have formulated a list of theses based on a literature review that can guide the development of such a methodology. We have proposed how an existing LCA methodology can be adapted to AI as a first step in that direction. It is important to stress that developing a comprehensive LCA methodology for AI is a big community effort that requires contributions from LCA experts and AI practitioners. We would like to kickstart this cross-discipline collaboration by inviting fellow scholars to discuss our ideas and findings.

In the following, some possibilities for future research are presented. We have stated above that data gaps must be identified and, in the next step, filled. Especially hardware-related environmental impacts are in the focus here, along with the training conditions of large commercial AI models. The Digital Product Passport (DPP) is a concept that can help collect data along the devices' value chain and life cycle. Once this data is available in the DPP, it can be leveraged for LCA. First concepts of an LCA method based on the DPP exist [28]. Thinking one step further, the DPP could integrate LCA results, e.g. in a machine-readable format of Environmental Product Declarations (EPDs) [29], to improve environmental data consistency.

A comprehensive LCA methodology for AI will make the environmental impacts of AI more transparent. It can then be used to systematically analyze different setups for AI applications. For example, parameters like batch size have an influence on energy consumption [3]. Developers can influence the environmental footprint of AI models by choosing the right setup and parameters, or by employing techniques to reduce model size or training time. Work on this already exists [15], but more systematic studies are required. One goal could be to develop recommendations for AI practitioners on how to achieve the smallest possible environmental impacts for their models given a desired output in terms of, e.g., training time or accuracy.

Acknowledgements

This work is funded by the BMUV, project Green-AI Hub, grant number 57-E-2110.

References

[1] A. van Wynsberghe, Sustainable AI: AI for sustainability and the sustainability of AI, AI and Ethics 1 (02 2021).

[2] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown, A. Luccioni, T. Maharaj, E. D. Sherwin, S. K. Mukkavilli, K. P. Körding, C. P. Gomes, A. Y. Ng, D. Hassabis, J. C. Platt, F. Creutzig, J. T. Chayes, Y. Bengio, Tackling Climate Change with Machine Learning, CoRR abs/1906.05433 (2019).

[3] L. Bouza, A. Bugeau, L. Lannelongue, How to estimate carbon footprint when training deep learning models? A guide and review, Environmental Research Communications 5 (11 2023).

[4] A. Berthelot, E. Caron, M. Jay, L. Lefèvre, Estimating the environmental impact of Generative-AI services using an LCA-based methodology, Procedia CIRP 122 (2024) 707–712.

[5] M. Hauschild, R. Rosenbaum, S. Olsen (Eds.), Life Cycle Assessment - Theory and Practice, Springer, 2018.

[6] D. Dokic, F. G. I. Woud, W. Maass, Towards Sustainability of AI: A Systematic Review of Existing Life Cycle Assessment Approaches and Key Environmental Impact Parameters of Artificial Intelligence, in: Proceedings of HICSS, 2024.

[7] DIN EN ISO 14040:2006, Environmental management: Life cycle assessment - Principles and framework.

[8] DIN EN ISO 14044, Environmental management: Life cycle assessment –Life cycle assessment – Requirements and guidelines (ISO 14044:2006 + Amd 1:2017 + Amd 2:2020).

[9] ITU-T, L.1410 : Methodology for environmental life cycle assessments of information and communication technology goods, networks and services. URL https://www.itu.int/rec/T-REC-L.1410

[10] A.-L. Ligozat, J. Lefevre, A. Bugeau, J. Combaz, Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions, Sustainability 14 (9) (2022).

[11] F. Bordage, L. de Montenay, S. Benqassem, J. Delmas-Orgelet, F. Domon, D. Prunel, C. Vateau, E. L. Perasso, Digital technologies in Europe: an environmental life cycle approach (2021).

[12] Boavizta working group, Boavizta. URL https://boavizta.org/en

[13] T. Schaubroeck, S. Schaubroeck, R. Heijungs, A. Zamagni, M. Brandão, E. Benetto, Attributional & Consequential Life Cycle Assessment: Definitions, Conceptual Characteristics and Modelling Restrictions, Sustainability 13 (13) (2021).

[14] FAU, Guide to Science Information Resources. URL https://libguides.fau.edu/science-resources/reference-searching

[15] K. Kirkpatrick, The Carbon Footprint of Artificial Intelligence, Commun. ACM 66 (8) (2023) 17–19.

[16] U. Gupta, Y. Kim, S. Lee, J. Tse, H. Lee, G. Wei, D. Brooks, C. Wu, Chasing Carbon: The Elusive Environmental Footprint of Computing, IEEE Micro 42 (4) (2022) 37–47.

[17] K. Lottick, S. Susai, S. A. Friedler, J. P. Wilson, Energy usage reports: Environmental awareness as part of algorithmic accountability, CoRR abs/1911.08354 (2019).

[18] L. F. W. Anthony, B. Kanding, R. Selvan, Carbontracker: Tracking and predicting the carbon footprint of training deep learning models, CoRR abs/2007.03051 (2020).

[19] L. Lannelongue, J. Grealey, M. Inouye, Green Algorithms: Quantifying the Carbon Footprint of Computation, Advanced Science 8 (12) (2021).

[20] S. Budennyy, V. Lazarev, N. Zakharenko, A. Korovin, O. Plosskaya, D. Dimitrov, V. Akhripkin, I. Pavlov, I. Oseledets, I. Barsola, I. Egorov, A. Kosterina, L. Zhukov, eco2AI: Carbon Emissions Tracking of Machine Learning Models as the First Step Towards Sustainable AI, Doklady Mathematics 106 (2023).

[21] 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, J. Mach. Learn. Res. 21 (1) (2020).

[22] A. Lacoste, A. Luccioni, V. Schmidt, T. Dandres, Quantifying the carbon emissions of machine learning, CoRR abs/1910.09700 (2019).

[23] T. Trébaol, Cumulator (2020). URL https://github.com/EPFLiGHT/cumulator

[24] L. H. Kaack, P. L. Donti, E. Strubell, G. Kamiya, F. Creutzig, D. Rolnick, Aligning artificial intelligence with climate change mitigation, Nature Climate Change 12 (6) (2022) 518–527.

[25] B. Purvis, Y. Mao, D. Robinson, D. Robinson, Three pillars of sustainability: in search of conceptual origins , Sustainability Science 14 (3) (2019) 681–695.

[26] European Commission - Joint Research Centre - Institute for Environment and Sustainability, International Reference Life Cycle Data System (ILCD) Handbook b - Framework and Requirements for Life Cycle Impact Assessment Models and Indicators.

[27] European Commission, European Platform on LCA — EPLCA. URL https://eplca.jrc.ec.europa.eu/LCDN

[28] C. Plociennik, A. Nazeri, M. H. Rimaz, S. Knetsch, A. do Carmo Precci Lopes, T. Hagedorn, J. Baehr, M. Vogelgesang, C. Li, W. Benner, B. Kellerer, E. Ionescu, M. Ruskowski, A. Weidenkaff, Connecting Producers and Recyclers: A Digital Product Passport Concept and Implementation Suitable for End-of-Life Management, Procedia CIRP 122 (2024) 831–836.

[29] A. Del Borghi, LCA and communication: Environmental Product Declaration, The International Journal of Life Cycle Assessment 18 (02 2013).