Artificial Intelligence in a degrowth context.

A conviviality perspective on machine learning

The degrowth movement lacks a concrete vision for technology, thereby disregarding a crucial aspect of the green growth narrative. This paper helps fill this gap by exploring the compatibility of Artificial Intelligence with a degrowth-related concept: convivial tools – tools that promote autonomy, creativity, and relationships among humans and with nature.

Marion Meyers

Artificial Intelligence in a degrowth context.
A conviviality perspective on machine learning
GAIA 33/1 (2024): 186 – 192

Abstract
Degrowth has emerged as a strong voice against the green growth narrative. However, it has so far left largely unshaped its vision for technology, thereby overlooking a pivotal element of the green growth narrative. This article contributes to filling this gap by analyzing the appropriateness of a digital technology, Artificial Intelligence, to a degrowth context. It does so through the angle of conviviality, a concept introduced by Ivan Illich and frequently used by degrowth scholars, which states that convivial tools should foster autonomy, creativity, and relationships among humans and with nature. This paper specifically applies Vetter's Matrix of Convivial Technology to an application of machine learning with potential environmental benefits: predictive maintenance – a proactive maintenance technique based on real-time sensor monitoring. Three key limitations to its conviviality are identified: 1. the high complexity of machine learning, 2. its environmental impacts, and 3. the size of the infrastructure it relies on. These limitations prompt critical reflections on the appropriateness of machine learning (as a part of Artificial Intelligence) to degrowth but also act as inspirations for reshaping the technology towards more conviviality.

Keywords
Artificial Intelligence, conviviality, degrowth, machine learning, technology

Marion Meyers, MSc | ETH Zurich | Zurich | CH | marion.meyers@hotmail.com

https://doi.org/10.14512/gaia.33.1.13
Received November 27, 2023; revised version accepted March 28, 2024 (editorial board peer review).

Degrowth has emerged in the last few years as an important discourse within academic and activist circles. Based on an increasing amount of evidence indicating that absolute decoupling between economic growth and natural resource usage is highly unlikely to happen at the needed scale to remain within the 1.5-degree scenario (Parrique et al. 2019), degrowth calls for the need to re-imagine a society beyond the growth paradigm. It stands in strong opposition to the green growth narrative which argues that technology will bring the necessary efficiency improvements to keep pursuing economic growth while facing the various environmental crises. In that regard, digital technologies are frequently described as promising (Pollex and Lenschow 2016). Interestingly, while degrowth addresses key aspects of a new social imaginary by promoting principles of environmental and social justice, it has so far left largely unshaped its vision of technology, thereby overlooking a pivotal element of the green growth narrative. This is the research gap that this paper aims to contribute to: shaping a degrowth perspective on technology. Specifically, this paper focuses on a digital technology that lies at the center of political support, promises of societal and environmental benefits, and extremely high investments and research interest: Artificial Intelligence (AI), and even more specifically machine learning (ML). It aims to tackle the following research question: Could AI be appropriate to a degrowth context? If so, how?

To tackle this question, I use a central concept within the existing literature on degrowth and technology: conviviality (Kerschner et al. 2018, Zoellick and Bisht 2018). Ivan Illich first introduced the concept of convivial tools, which he defined as "those which give each person who uses them the greatest opportunity to enrich the environment with the fruits of his or her vision" (Illich 1973, p. 21). Defined in opposition to industrial tools, the concept has been utilized as inspiration for what technologies aligned with degrowth could look like. Vetter (2018) for example combined Illich's concept with a focused ethnography work with degrowth-related groups to develop the Matrix of Convivial Technology (MCT), a tool that has already been applied in the degrowth literature (Bobulescu and Fritscheova 2021, Pansera and Fressoli 2021, Priavolou et al. 2022, Ralph 2021). The literature on degrowth and technology has however so far mostly focused

GAIA Masters Student Paper Award

Marion Meyers is the winner of the 2024 *GAIA Masters Student Paper Award (GMSPA)*[a]. Her paper *Conviviality perspective on machine learning: Discussing the appropriateness of Artificial Intelligence to a degrowth context* was selected by an international jury and is now published in *GAIA* after successful editorial board peer review.

a For more details see *www.oekom.de/publikationen/zeitschriften/gaia/c-131*.

on physical technologies. A few authors external to degrowth have used conviviality to critique the opacity of digital tools such as ML (Beinsteiner 2021), or social media platforms (Christiaens 2022). However, they do not focus on environmentally relevant applications of digital technologies and do not systematically assess the conviviality of a particular application.

As there is little work on the appropriateness of digital technologies to a degrowth context, I contribute to filling this gap by applying Vetter's MCT as an evaluation tool to assess ML, thereby proposing an adaptation of this framework to a digital technology. This – to the best of my knowledge – was not previously done. I particularly focus on an application of ML with claimed environmental benefits, namely predictive maintenance (PdM) (Rolnick et al. 2023).

Vetter's Matrix of Convivial Technology (MCT)

The MCT (Vetter 2018) lists five dimensions important to convivial technologies: relatedness, access, adaptability, bio-interaction, and appropriateness, which are analyzed on four levels corresponding to typical technologies' lifecycle levels: material, production, use, and infrastructure. Each square of the MCT is filled with antagonist terms that help identify important characteristics of convivial technologies (table 1, pp. 190 f.). The conviviality dimensions are described by Vetter (2018, pp. 1782 ff.) as follows:

- The ***relatedness dimension*** is defined by the question: "What does it [the technology] bring between people?". Convivial technologies are considered to be those that support and enhance human relationships.
- The ***access dimension*** is defined by the question: "Who can build or use it where and how?". It doesn't only involve access to the material and the means of production of the technology but also to the knowledge and skills needed to build or use it. Convivial technologies should strive to be accessible to anyone.
- The ***adaptability dimension*** asks the question: "How independent and linkable is it?". "Independent" involves "the autonomy to decide whether to use a technological device or not". "Linkable" means that people should be "able to decide whether one wants to be independent or linked".
- The ***bio-interaction dimension*** is defined by the question: "How does it interact with living organisms?". This aspect of conviviality goes further than just trying to reduce the environmental impacts of the technology but includes the fact that it should be useful to natural processes.
- The ***appropriateness dimension*** focuses on "the relation between input and output considering the context". Following Vetter, appropriateness means "to take the whole situation into account, consider the local availability of materials and skills, and then to decide where a technology makes sense and where not". I understand this dimension as a reminder that the analysis of conviviality (i. e., the four dimensions above) has to be tailored to the specific context in which a technology is situated. I will therefore not analyze ML on the appropriateness dimension but will instead make sure that the analysis of the other dimensions is adapted to a local context. I do so by selecting a case example of ML in the manufacturing sector: PdM.

Case example: Predictive maintenance (PdM) in manufacturing

PdM refers to the process of predicting when a machine will need to be repaired or replaced through the continuous monitoring of sensors integrated within the machines or externally installed such as power meters, cameras, humidity, or temperature sensors. The use of ML for PdM (ML-based PdM) has received growing attention in the last ten years, especially in the context of the Industry 4.0 narrative (Carvalho et al. 2019, Çinar et al. 2020). Rolnick et al. (2023) consider ML-based PdM a high-leverage application to tackle climate change because 1. it can reduce production waste by detecting early when a machine is malfunctioning and thereby prevent it from creating defective products and 2. it can prevent environmentally damaging leaks, for example by accurately predicting pipes' failure. Other authors argue that PdM can prolong machines' lifetime, such as Abidi et al. (2022) who have trained an ML model on lithium-ion battery data sets that can successfully predict the health conditions of various components, thereby allowing for improved maintenance planning. The three core characteristics of PdM as a case example in this article are 1. the data are generated by sensors, 2. they are collected through real-time monitoring of these sensors, and 3. the objective is to reduce the overall environmental impacts of an industrial context.

Adapting the levels of the Matrix of Convivial Technology to predictive maintenance

For the MCT to reflect the structure of ML, its levels have to be specified to the case example. Since the design of the MCT was based on an analysis of mainly physical products, its application to a digital product is not so straightforward. It involves for example defining what counts as "material" for an ML algorithm. For this analysis, I focus on two crucial sets of physical devices on which PdM systems rely: data-collecting sensors and one or multiple server(s), resulting in the following adaptation of the MCT levels described by Vetter (2018, p. 1780):

>

- The ***material level*** corresponds to the "harvesting, processing and disposal of raw matter" needed for the sensors and the server pieces.
- The ***production level*** corresponds to the "assembling of raw materials and pre-products" for the production of the sensors and the server pieces.
- The ***use level*** corresponds to the set of tasks needed for the development, deployment, and use of an ML algorithm for PdM, including data acquisition, data production, learning, data storage, and inference (running the trained algorithm to issue the maintenance alerts).
- The ***infrastructure level*** corresponds to the "needed environment for using" the ML system, which I understand as the type of machines in the industrial setting that can be subject to PdM and the infrastructure the algorithm relies on.

Key limiting factors to the conviviality of machine learning

At each of the adapted MCT levels, I have identified various components of ML-based PdM that contradict some dimensions of conviviality: 1. the high complexity of ML, 2. the environmental impacts of ML and 3. the (size of) the infrastructure needed. Table 1 (pp. 190 f.) presents an overview of these components as well as the antagonist terms from Vetter's MCT which contributed to identifying them.

1 The high complexity of machine learning

The first limiting factor to the conviviality of ML is its high complexity, englobing both the highly complex production of electronic devices (production level of the MCT), and the technical opacity of ML algorithms (use level of the MCT).

Highly complex devices production (production level)

Most electronic devices on which PdM applications rely require highly specific production processes and therefore result from a highly centralized production system. I argue that this constitutes a key limiting factor to conviviality according to the dimensions of relatedness, access, and adaptability.

Relatedness: It contradicts the relatedness dimension firstly because ordering devices from far-away centralized companies creates distance instead of a conjoint experience *[R.P1, R.P2]*[1]. Secondly, today's centralized electronics production relies on labor exploitation, for example in semiconductor factories in Malaysia where migrant workers have reported cases of abuse, illegal wage reductions, and forced labor (Mortensen 2019), emphasizing the contradiction with a dimension that aims at enhancing human relationships.

Access: A complex device production also challenges the access dimension. Taking the example of semiconductor chips, central to most electronic devices, the costs of investing in the production capabilities are very high *[A.P2]* and the machines used in the process are very complex to handle which results in only a handful of companies producing these chips *[A.P1, A.P3]* (Miller 2022).

Adaptability: The adaptability dimension is also challenged. As mentioned above, some components of electronic devices such as chips require highly special machines to be produced *[Ad.P2]* and are only economical at a large scale because the infrastructure needed is highly expensive [Ad.P1], thereby preventing the production from being possible everywhere *[Ad.P3]*.

Technical opacity of machine learning algorithms (use level)

ML algorithms can be very opaque due to their high complexity. Burrell (2016) highlights two aspects contributing to this opacity: technological illiteracy and cognitive mismatch. Technological illiteracy means that ML development pipelines are typically only understood by a minority. Cognitive mismatch means that many ML models present inherent black-box characteristics resulting both from characteristics of the models and from the scale at which they operate. Many applications of PdM use models with inherent opaque features such as neural networks, support vector machines, and random forests (Carvalho et al. 2019, Çinar et al. 2020). I argue that the technical opacity of ML-based PdM is a key limiting factor to conviviality according to the dimensions of relatedness, access, and adaptability.

Relatedness: Technical opacity challenges this dimension because a PdM system based on ML where the efforts of providing explanations of its functioning have not been made or are impossible due to cognitive mismatch can create feelings of alienation in the workplace *[R.U2]* (Vredenburgh 2022). It can also impair the self-determination of the workers because they lack the necessary power to control, adapt, and creatively use the technology *[R.U3, R.U1]*.

Access: When it comes to the access dimension, the issues are numerous: technical illiteracy implies that ML is a tool that is often restricted to an educated elite *[A.U1]*; cognitive mismatch implies that some ML models are abstract rather than comprehensible *[A.U3]*; the high complexity of ML often forces organizations to use the services of AI consultancy firms because the costs required to develop this expertise "in-house" are too high (Saha 2024), thereby making local organizations increasingly reliant on foreign experts *[A.U2]*.

Adaptability: Technical opacity limits the possibilities for the workers impacted by the technology to change it *[Ad.U1]* or to repair it *[Ad.U2]* which challenges the adaptability dimension.

1 The annotations are used to refer to the antagonist terms in table 1. Here for example *[R.P1]* refers to the term "organization centralized – organization distributed" at the intersection between the production level and the relatedness dimension.

2 The environmental impacts of machine learning

The second limiting factor is the environmental impact of ML. This includes the environmental impacts of ML's supply chain (material, production and use level of the MCT), and specifically its use of rare earth elements (material level). It also includes the fact that ML classifies as an eco-innovation (use level).

Environmental impacts of machine learning throughout its supply chain (material, production and use level)

ML has numerous environmental consequences challenging the bio-interaction dimension.

Bio-interaction: Firstly, the development of electronics comes with great environmental impacts. First-order environmental effects include the environmentally destructive extraction of rare earth elements *[B.M1]* but also the resource-intensive processes needed for the production of these devices *[B.P1]* (Bieser et al. 2023) and the immense stream of electronic waste generated (Abdelbasir et al. 2018). Secondly, ML systems typically use a lot of energy through their high server utilization *[B.U1]* (Strubell et al. 2019).

Eco-innovation rather than ecosystem beneficial (use level)

ML-based PdM perfectly fits the definition of eco-innovation (Pansera 2011): it is aimed at reducing the environmental impacts of a certain process but not at contributing something useful to the ecosystem. This contradicts the bio-interaction dimension.

Bio-interaction: Convivial technologies should allow co-productivity which is not the case for PdM *[B.U2]*.

The use of rare earth elements in electronics (material level)

Electronic devices are based on rare earth elements. The current recycling rate of electronics is low (Abdelbasir et al. 2018) which means that most electronic devices result from mining. I argue that this constitutes a key limiting factor to conviviality according to the dimensions of relatedness, access, adaptability, and bio-interaction.

Relatedness: This dimension is challenged firstly because mining companies, driven by profit incentives *[R.M1],* have a history of human rights abuses (McKie 2021, Raid 2021). Secondly, mining practices sustain colonialist relationships of domination *[R.M2]*. Indeed, the Global North benefits most from the extracted resources while most mines in which workers and nature are exploited are located in the Global South (Jerez et al. 2021). The above characteristics contradict the dimension of relatedness following which human relationships should be enhanced.

Access: The use of rare earth elements challenges the access dimension because their extraction happens today for the most part in a highly centralized market where a few countries dominate the resources or a few companies dominate the means of access to these resources *[A.M1]*. Additionally, mining is a highly energy-intensive process (Althaf and Babbitt 2021) implying high costs and a major barrier to access *[A.M2]*.

Adaptability: This dimension is also contradicted. Indeed, rare earth elements are regarded as special materials rather than standardized materials *[Ad.M3]* and accessing them requires special tools *[Ad.M1]*. The high costs and energy intensity of the extraction processes also imply that accessing these materials is a large-scale operation *[Ad.M2]*.

Bio-interaction: The contradiction between the extraction of rare earth elements and the bio-interaction was mentioned in the previous section.

3 The (size of the) infrastructure needed for machine learning

The third aspect of ML-based PdM systems that challenges numerous dimensions of conviviality is the size of the infrastructure they rely on, specifically their need for big data infrastructure (use and infrastructure levels of the MCT).

Real-time monitoring requires big data infrastructure (use and infrastructure level)

The type of data gathered through the real-time monitoring of sensors, such as in PdM, is often considered Big Data (Hashem et al. 2015). I argue that the need for big data infrastructure poses issues to the relatedness, access, adaptability, and bio-interaction dimensions.

Relatedness: It is complex and costly for local organizations to maintain the data storage capacities and computing power, as well as to develop the skills needed to deal with big data applications (Hashem et al. 2015). This leads many to resort to the use of cloud services, which I argue challenges the relatedness dimension. Indeed, it creates physical and mental distance between the users of the algorithm and the infrastructure on which it relies *[R.I1]*. I also argue that it prevents workers' self-determination because they lose control over their data *[R.U3]* (De Filippi and McCarthy 2012).

Access: The complexity and cost of big data infrastructures also restrict access to these applications to those with the resources to invest in such servers and expertise *[A.I1, A.I2]*. While the emergence of cloud service providers could be seen as a way to democratize this access, it also exacerbates the reliance of the technology users on foreign experts *[A.U2]* and makes the technology more abstract *[A.U3]*.

Adaptability: The need for big data infrastructure also challenges the adaptability dimension. Firstly, it prevents PdM from being operable without additional infrastructure *[Ad.U3]*. Secondly, the outsourcing of tasks to cloud services means that the infrastructure is not fully locally operable *[Ad.I2]* which hinders its self-repairability *[Ad.I1]*. Finally, the current tendency of cloud service providers to be dominated by a few multinational com- >

TABLE 1: Result table adapted from Vetter (2018, p. 1780) presenting the aspects of machine learning-based predictive maintenance (ML-based PdM) that contradict conviviality per level and dimension of the Matrix of Convivial Technology (MCT): the environmental impacts of ML, the high complexity of ML, and the (size of the) infrastructure needed ("><" represents a contradiction between the antagonist terms of the MCT in this square and the

LEVELS **DIMENSIONS**	**MATERIALS** *harvesting, processing and disposal of raw matter*	**PRODUCTION** *assembling raw materials and preproducts*
RELATEDNESS *What does it bring about between people?*	>< **environmental impacts of ML:** use of rare earth elements	>< **high complexity of ML:** highly complex devices production
	[R.M1] market-driven ----- need-driven *[R.M2]* alien implementation ----- respects local traditions	*[R.P1]* organization centralized ----- organization distributed *[R.P2]* distance-creating ----- conjoint experience
ACCESS *Who can produce/ use it where and how?*	>< **environmental impacts of ML:** use of rare earth elements	>< **high complexity of ML:** highly complex devices production
	[A.M1] elitist ----- open to anyone *[A.M2]* cost-intensive ----- low-cost	*[A.P1]* elitist ----- open to anyone *[A.P2]* cost intensive ----- low cost *[A.P3]* secret or patented ----- knowledge freely accessible
ADAPTABILITY *How independent and linkable is it?*	>< **environmental impacts of ML:** use of rare earth elements	>< **high complexity of ML:** highly complex devices production
	[Ad.M1] special machines ----- everyday tools *[Ad.M2]* big scale economical ----- small scale economical *[Ad.M3]* special materials ----- standardized materials	*[Ad.P1]* big scale economical ----- small scale economical *[Ad.P2]* special machines ----- everyday tools *[Ad.P3]* special conditions ----- everywhere possible
BIO-INTERACTION *How does it interact with living organisms?*	>< **environmental impacts of ML:** environmental impacts of ML throughout supply chain	>< **environmental impacts of ML:** environmental impacts of ML throughout supply chain
	[B.M1] set of environmental impacts illness/death ----- supports health deteriorating soil ----- improving soil water-polluting ----- improving water quality air-polluting ----- supports clean air violent ----- nonviolent hazardous potential ----- safety proven and tested toxic waste ----- biodegradable	*[B.P1]* set of environmental impacts illness/death ----- supports health deteriorating soil ----- improving soil water-polluting ----- improving water quality air-polluting ----- supports clean air violent ----- nonviolent hazardous potential ----- safety proven and tested toxic waste ----- biodegradable

panies such as Amazon and Google leads to a centralized rather than a distributed infrastructure *[Ad.I3]*.

Bio-interaction: Real-time monitoring requires a constant running time of all devices involved in the algorithm. This accentuates the energy consumption of this kind of ML system (highlighted in the "environmental impacts of ML" section above) and thereby further contradicts the bio-interaction dimension *[B.I1]*.

Discussion and outlook

The analysis showed that some fundamental aspects of ML-based PdM are in contradiction with the concept of conviviality (I argue below that these findings are relevant for ML as a whole). If a degrowth perspective would therefore only consider convivial technologies to be appropriate, then ML would not have a place within the movement's vision.

However, if we go back to Illich's conception of a convivial society, we note that he did not argue that all tools in a certain society had to be convivial for the society as a whole to be convivial. Rather, he argues that a convivial society should find a healthy balance between convivial and non-convivial tools and should strive for conviviality as an underlying value (Illich 1973, p. 24). Embracing this argument, a degrowth perspective would instead remain open to some technologies with limitations to conviviality, while centering conviviality as a value to aim towards. For ML, it can do so by turning its identified conviviality shortcomings into the following suggestions for its development: 1. dealing with technical opacity by promoting peer-learning, open-specified aspect of ML), and containing the subset of Vetter's antagonist terms that drove the identification of these conviviality limitations. Note: Each antagonist term is associated with a unique code (e.g., *[R.M1]*) for referencing.

USE *procuring the task it was built for*	**INFRASTRUCTURE** *needed environment for using*
>< **high complexity of ML:** technical opacity of ML algorithms >< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure	>< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure
[R.U1] preconfigured only ----- allows creativity *[R.U2]* alienating from own body ----- useful body enhancement *[R.U3]* heteronomy ----- self-determination	*[R.I1]* distance-creating ----- connects with eco processes
>< **high complexity of ML:** technical opacity of ML algorithms >< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure	>< **(size of the) infrastructure needed:** : real-time monitoring requires big data infrastructure
[A.U1] usable by an elite ----- usable by anyone *[A.U2]* need of foreign expert ----- use of local knowledge *[A.U3]* abstract ----- comprehensible	*[A.I1]* usable by an elite ----- usable by anyone *[A.I2]* cost intensive ----- low cost
>< **high complexity of ML:** technical opacity of ML algorithms >< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure	>< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure
[Ad.U1] fixed once finished ----- permanently changeable *[Ad.U2]* repairable by experts ----- repairable by skilled *[Ad.U3]* infrastructure needed ----- independent use possible	*[Ad.I1]* repairable by experts ----- repairable by skilled *[Ad.I2]* operable only from distance ----- locally operable *[Ad.I3]* centralized ----- distributed
>< **environmental impacts of ML:** environmental impacts of ML throughout supply chain >< **environmental impacts of ML:** eco-innovation rather than ecosystem beneficial	>< **(size of the) infrastructure needed:** real-time monitoring requires big data infrastructure
[B.U1] set of environmental impacts illness/death ----- supports health deteriorating soil ----- improving soil water-polluting ----- improving water quality air-polluting ----- supports clean air violent ----- nonviolent hazardous potential ----- safety proven and tested toxic waste ----- biodegradable *[B.U2]* suppresses organic processes ----- allows co-productivity	*[B.I1]* set of environmental impacts illness/death ----- supports health deteriorating soil ----- improving soil water-polluting ----- improving water quality air-polluting ----- supports clean air violent ----- nonviolent hazardous potential ----- safety proven and tested toxic waste ----- biodegradable

source ML and prioritizing explainability, 2. reducing environmental impacts of ML through recycling of electronics and algorithm design striving for low energy usage, 3. turning private cloud services into community-run infrastructures, and 4. instigating democratic deliberations in the workplace on the use of ML systems (Meyers 2023).

On whether these findings can be extended beyond PdM, I argue that the three limitations found apply to all ML algorithms but that the degree of intensity of the limitations varies with the local context: 1. the high complexity of ML is shared by all algorithms but varies with the choice of model and size, 2. the environmental impacts of ML are shared between all applications because they rely on the same set of electronic devices but the energy usage can vary with the model and size, and 3. the big data characteristic of PdM does not apply to all applications but many aspects of the infrastructure, such as the computing resources, do.

While the use of the MCT bridges the gap between the degrowth and AI discourses, questions remain on its applicability to digital technologies. Indeed, some elements of the matrix contradict most digital technologies, suggesting either that no digital technology is convivial or that we have to adapt the concept of conviviality to accommodate their unique features. For example, it could be argued that all smartphones contradict conviviality because their inner workings are hidden behind a screen, thereby making them abstract rather than comprehensible *[A.U3]*. Beinsteiner (2021) and Christiaens (2022) have however argued that it is the transparency and explainability of digital tools which determine their conviviality. Their perspective then allows us to compare the conviviality of complex neural network models with >

that of smartphone applications or simple ML models based on their level of explainability and transparency. Hence, I suggest that while all elements of the MCT might not be directly applicable to digital technologies, it is a useful framework for evaluating and comparing different technologies as well as identifying avenues for change.

Many questions remain however unanswered. Future work should for example critically assess whether the environmentally beneficial applications of ML are actually relevant for degrowth-aligned infrastructures and sectors or whether they are only beneficial within today's unsustainable infrastructures.

Acknowledgements: The author would like to thank the *GMSPA* jury and an anonymous reviewer for their helpful comments, as well as *Irmi Seidl, Lorenz Hilty* and *Andrea Vetter* for their insightful feedback on this research.
Funding: This work received no external funding.
Competing interests: The author declares no competing interests.

References

Abdelbasir, S. M., S. S. M. Hassan, A. H. Kamel, R. S. El-Nasr. 2018. Status of electronic waste recycling techniques: A review. *Environmental Science and Pollution Research* 25/17: 16533–16547. https://doi.org/10.1007/s11356-018-2136-6.

Abidi, M. H., M. K. Mohammed, H. Alkhalefah. 2022. Predictive maintenance planning for industry 4.0 using machine learning for sustainable manufacturing. *Sustainability* 14/6: 3387. https://doi.org/10.3390/su14063387.

Althaf, S., C. W. Babbitt. 2021. Disruption risks to material supply chains in the electronics sector. *Resources, Conservation and Recycling* 167: 105248. https://doi.org/10.1016/j.resconrec.2020.105248.

Beinsteiner, A. 2021. Conviviality, the internet, and AI: Ivan Illich, Bernard Stiegler, and the question concerning information-technological self-limitation. *Open Cultural Studies* 4/1: 131–142. https://doi.org/10.1515/culture-2020-0013.

Bieser, J. C. T., R. Hintemann, L. M. Hilty, S. Beucker. 2023. A review of assessments of the greenhouse gas footprint and abatement potential of information and communication technology. *Environmental Impact Assessment Review* 99: 107033. https://doi.org/10.1016/j.eiar.2022.107033.

Bobulescu, R., A. Fritscheova. 2021. Convivial innovation in sustainable communities: Four cases in France. *Ecological Economics* 181: 106932. https://doi.org/10.1016/j.ecolecon.2020.106932.

Burrell, J. 2016. How the machine "thinks": Understanding opacity in machine learning algorithms. *Big Data and Society* 3/1. https://doi.org/10.1177/2053951715622512.

Carvalho, T. P., F. A. A. M.N. Soares, R. Vita, R. da P. Francisco, J. P. Basto, S. G. S. Alcalá. 2019. A systematic literature review of machine learning methods applied to predictive maintenance. *Computers and Industrial Engineering* 137: 106024.https://doi.org/10.1016/j.cie.2019.106024.

Christiaens, T. 2022. Convivial autonomy in platform capitalism. In: *New interdisciplinary perspectives on and beyond autonomy.* Edited by C. Watkin, O. Davis. New York: Routledge. 69–82. https://doi.org/10.4324/9781003331780-7.

Çinar, Z. M., A. A. Nuhu, Q. Zeeshan, O. Korhan, M. Asmael, B. Safaei. 2020. Machine learning in predictive maintenance towards sustainable smart manufacturing in industry 4.0. *Sustainability* 12/19: 8211. https://doi.org/10.3390/su12198211.

De Filippi, P., S. McCarthy. 2012. Cloud computing: Centralization and data sovereignty. *European Journal of Law and Technology* 3/2.

Hashem, I. A. T., I. Yaqoob, N. B. Anuar, S. Mokhtar, A. Gani, S. Ullah Khan. 2015. The rise of "big data" on cloud computing: Review and open research issues. *Information Systems* 47: 98–115. https://doi.org/10.1016/j.is.2014.07.006.

Illich, I. 1973. *Tools for conviviality.* World perspectives 47. New York: Harper and Row.

Jerez, B., I. Garcés, R. Torres. 2021. Lithium extractivism and water injustices in the Salar de Atacama, Chile: The colonial shadow of green electromobility. *Political Geography* 87: 102382. https://doi.org/10.1016/j.polgeo.2021.102382.

Kerschner, C., P. Wächter, L. Nierling, M. H. Ehlers. 2018. Degrowth and technology: Towards feasible, viable, appropriate and convivial imaginaries. *Journal of Cleaner Production* 197: 1619–1636. https://doi.org/10.1016/j.jclepro.2018.07.147.

McKie, R. 2021. Child labour, toxic leaks: The price we could pay for a greener future. *Guardian,* January 3, 2021. www.theguardian.com/environment/2021/jan/03/child-labour-toxic-leaks-the-price-we-could-pay-for-a-greener-future (accessed March 22, 2024).

Meyers, M. 2023. *A degrowth perspective on artificial intelligence: Analysing the appropriateness of machine learning to a degrowth context.* Master's Thesis. ETH Zurich. https://doi.org/10.3929/ethz-b-000622669.

Miller, C. 2022. *Chip war: The fight for the world's most critical technology.* New York: Simon and Schuster.

Mortensen, N. H. 2019. Forced labour behind European electronics. *Danwatch,* June 28, 2019. https://danwatch.dk/en/undersoegelse/forced-labour-behind-european-electronics (accessed March 22, 2024).

Pansera, M. 2011. The origins and purpose of eco-innovation. *Global Environment* 4/7–8: 128–155. https://doi.org/10.3197/ge.2011.040706.

Pansera, M., M. Fressoli. 2021. Innovation without growth: Frameworks for understanding technological change in a post-growth era. *Organization* 28/3: 380–404. https://doi.org/10.1177/1350508420973631.

Parrique, T. et al. 2019. *Decoupling debunked: Evidence and arguments against green growth as a sole strategy for sustainability.* Brussels: European Environmental Bureau.

Pollex, J., A. Lenschow. 2016. Surrending to growth? The European Union's goals for research and technology in the *Horizon 2020* framework. *Journal of Cleaner Production* 197: 1863–1871. https://doi.org/10.1016/j.jclepro.2016.10.195.

Priavolou, C., K. Troullaki, N. Tsiouris, C. Giotitsas, V. Kostakis. 2022. Tracing sustainable production from a degrowth and localisation perspective: A case of 3D printers. *Journal of Cleaner Production* 376: 134291. https://doi.org/10.1016/j.jclepro.2022.134291.

Raid. 2021. *The road to ruin? Electric vehicles and workers' rights abuses at DR Congo's industrial cobalt mines.* www.raid-uk.org/sites/default/files/report_road_to_ruin_evs_cobalt_workers_nov_2021.pdf (accessed March 22, 2024).

Ralph, N. 2021. A conceptual merging of circular economy, degrowth and conviviality design approaches applied to renewable energy technology. *Journal of Cleaner Production* 319: 128549. https://doi.org/10.1016/j.jclepro.2021.128549.

Rolnick, D. et al. 2023. Tackling climate change with machine learning. *ACM Computing Surveys* 55/2: 1–96. https://doi.org/10.1145/3485128.

Saha, S. 2024. *Global AI consulting services market outlook (2023 to 2033).* Future Market Insights. www.futuremarketinsights.com/reports/ai-consulting-services-market (accessed March 25, 2024).

Strubell, E., A. Ganesh, A. McCallum. 2019. Energy and policy considerations for deep learning in NLP. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Edited by A. Korhonen, D. Traum, L. Màrquez. Florence: Association for Computational Linguistics. 3645–3650. https://doi.org/10.18653/v1/P19-1355.

Vetter, A. 2018. The Matrix of Convivial Technology: Assessing technologies for degrowth. *Journal of Cleaner Production* 197: 1778–1786. https://doi.org/10.1016/j.jclepro.2017.02.195.

Vredenburgh, K. 2022. Freedom at work: Understanding, alienation, and the AI-driven workplace. *Canadian Journal of Philosophy* 52/1: 78–92. https://doi.org/10.1017/can.2021.39.

Zoellick, J. C., A. Bisht. 2018. It's not (all) about efficiency: Powering and organizing technology from a degrowth perspective. *Journal of Cleaner Production* 197: 1787–1799. https://doi.org/10.1016/j.jclepro.2017.03.234.

Marion Meyers
BA in data science from Maastricht University, NL, followed by an MSc in Science, Technology and Policy from ETH Zurich, CH. Research interests: investigate and discuss the ethical, social, and environmental issues of digital technologies and their associated infrastructures.