Issue #2 | Spring 2023

sustain

AI and the Challenge of Sustainability

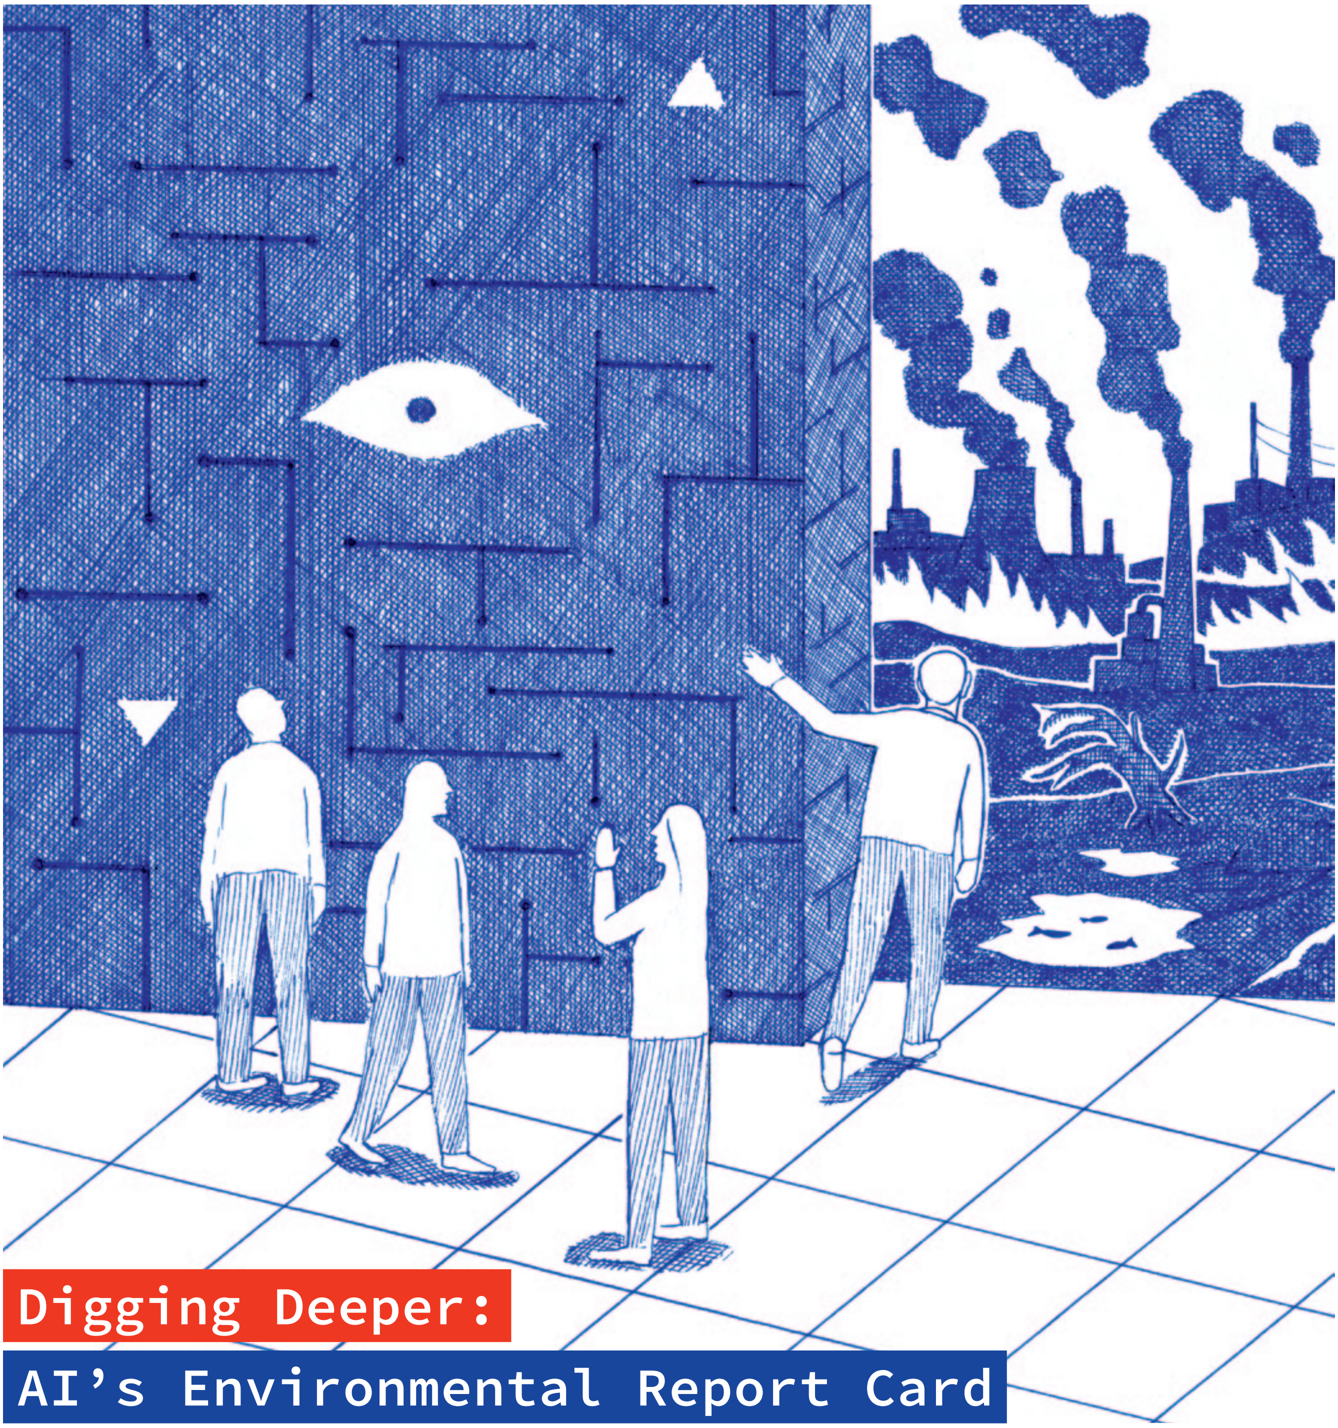

Digging Deeper:
AI's Environmental Report Card

Does Artificial Intelligence Consume More Resources than it Conserves?

sustain
Sustainability Index
for Artificial Intelligence

Project Partners:

ALGORITHM WATCH

iöw
INSTITUT FÜR ÖKOLOGISCHE WIRTSCHAFTSFORSCHUNG

DAI-Labor
TU Berlin

Supported by the Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (BMUV) based on a decision of the German Bundestag

Federal Ministry
for the Environment, Nature Conservation,
Nuclear Safety and Consumer Protection

Dear Reader,

The discussion surrounding Artificial Intelligence (AI) could hardly be more divided. On the one hand, it is becoming increasingly clear that AI systems only work through the exploitation of social and natural resources. On the other hand, AI continues to be seen as a strategic technology, without which we will be unable to solve complex societal reform projects such as the transition to renewable energies and the mobility revolution. These two aspects cannot be separated, because neither the potential of AI nor the dangers and damaging consequences associated with it can be ignored. The discussion, however, must be rooted in fact.

Those who present AI as a solution must also deliver proof of its efficacy. In the second issue of SustAIn Magazine, we show that AI systems in energy supply can indeed increase the use of renewable energies. But their full potential can only be realized if the appropriate infrastructure is in place. One also frequently hears claims that autonomous minibuses are good for the climate. In most cases, though, the fact that significant resources are required for the operation of such minibuses goes unmentioned. Overly sanguine AI fantasies are unhelpful. We have to examine the entire lifecycles of AI systems if we want to assess their sustainability. We must stop only asking about CO_2 emissions when examining the environmental effects of these systems. And we have to undertake a precise, comprehensive and impartial analysis of AI systems if we are sincerely interested in ensuring the sustainability of AI.

Read more on these issues in this edition of SustAIn Magazine.

Dr. Anne Mollen
Project Manager “SustAIn: The Sustainability Index for Artificial Intelligence”

AlgorithmWatch

CONTENTS

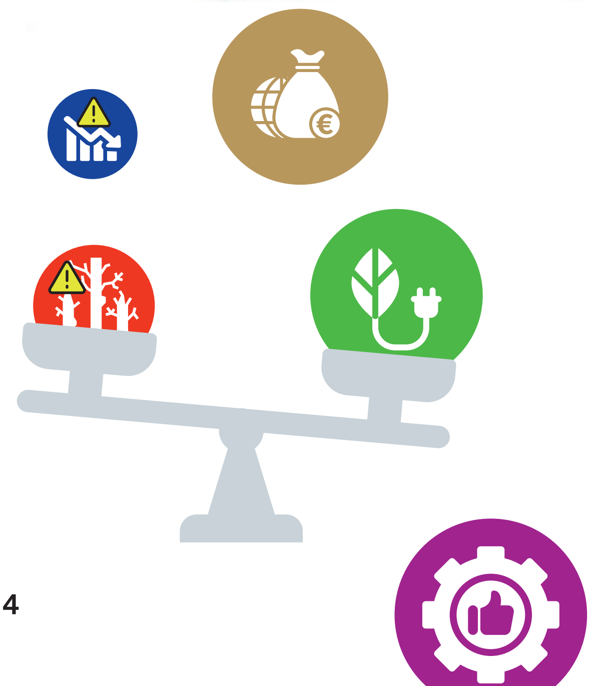

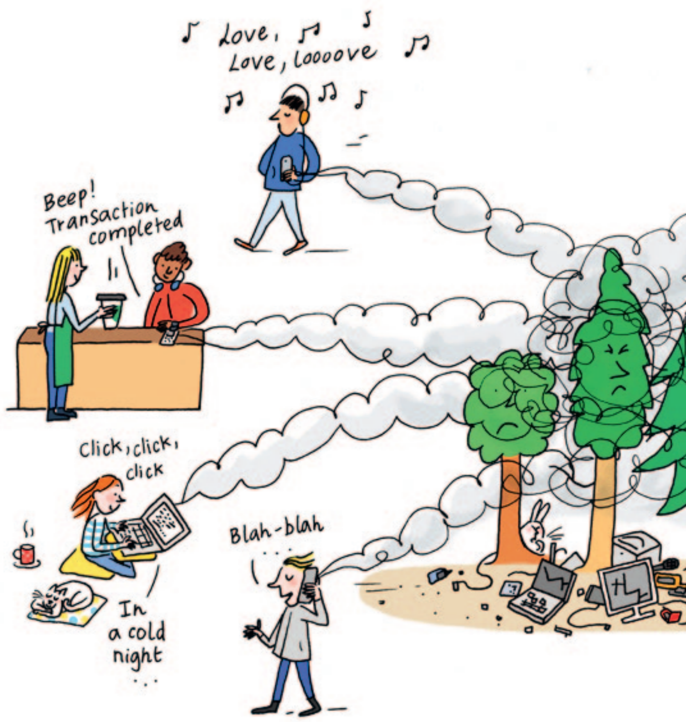
Love, Love, Loooove
Beep! Transaction completed
Click, click, click
Blah-blah
In a cold night

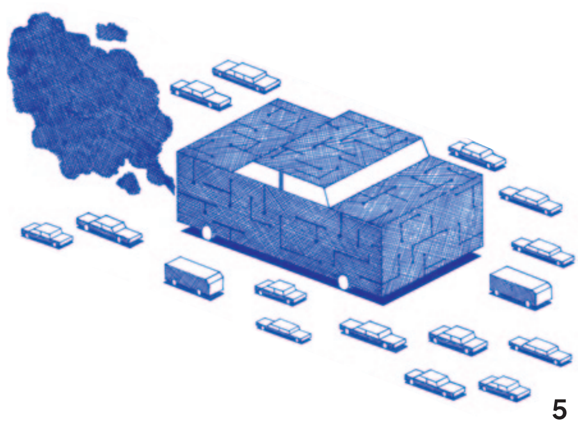

Step by Step Towards Sustainable AI

The sustainability of an AI system depends on many decisions taken during its lifecycle. At the SustAIn project, we have developed criteria that help us illustrate which steps should be heeded.

PHASE 1: PLANNING AND DESIGN

To a significant degree, the sustainability of an AI system is determined during the design and planning stage. As such, careful consideration should be given from the start as to whether the development and deployment of a complex and resource-intensive AI system is even necessary for a particular task. If it is, then additional fundamental questions must be asked: What data must the system process and how? What are the risks? And what security measures will be implemented?

General Requirements Checklist

Code of Conduct: A Code of Conduct sets out values and standards, such as transparency and equity, that provide guidance to all participants during planning and development.

Stakeholder Participation: All stakeholders related to the AI system must be identified and consulted during planning.

Documentation: Important decisions about the system must be documented in conjunction with its functions.

Risk Management: Once potential risks have been identified, appropriate safety measures must be taken.

Responsibilities: Where responsibility lies for the results produced by the AI system must be clearly defined.

Social Requirements Checklist

Transparency and Responsibility: A determination must be made about what information is disclosed and how responsibility for the system is distributed.

Non-Discrimination and Fairness: When humans are affected by an AI system's decisions, the fairness of its application must be determined. In addition, measures to eliminate bias and discriminatory consequences must be defined.

Technical Reliability and Human Supervision: Possible technical risks must be identified. A determination must be made about how harmful system operations can be remedied.

Self-Determination and Data Protection: Appropriate data protection measures must be implemented to ensure that data subjects know how their personal data is used. They must be given the opportunity to withhold personal data.

Inclusive and Participatory Design: The design must be barrier-free and inclusive according to applicable standards.

Cultural Sensitivities: Local knowledge assets (from stakeholders, for example) must be integrated into the development process. Team diversity is essential.

Environmental Requirements Checklist

Energy Consumption: Requirements for the necessary performance of the model and the resource budget must be defined, metrics for capturing energy efficiency established and test procedures for early detection of failing experiments developed.

Greenhouse Gas Emissions: Measures for offsetting CO_2 emissions must be defined and tools for recording emissions must be identified.

Sustainability in Use: Positive and negative potentials must be identified, analyzed and, if possible, quantified. Key performance indicators must be defined to record whether the potentials are being fully utilized.

Indirect Resource Consumption: Certifications and efficiency metrics for data centers must be the basis for ensuring that the hardware and the data centers used meet sustainability requirements.

Economic Requirements Checklist

Working Conditions and Jobs: When deployed in the workplace, impacts of the planned AI system on the workforce, working conditions and potential job losses must be analyzed and, if necessary, minimized.

PHASE 2: DATA

Data is at the core of all machine learning-based AI systems. System sustainability depends on the quality and amount of data to be processed, its interpretation and the way in which the data is obtained.

General Requirements Checklist

Documentation: Data sheets must be implemented to document information about the data collected and used and make it publicly available.

Stakeholder Participation: Individuals whose data will be used during the training or use of the AI systems must be consulted to ensure data is interpreted in an appropriate manner.

Social Requirements Checklist

Non-Discrimination and Fairness: Datasets must be examined for possible bias.

Technical Reliability: It is essential to ensure that current, complete, representative and reliable data is used.

Self-Determination and Data Protection: To protect personal data, principles such as data minimization, encryption, aggregation and anonymization must be adopted.

Environmental Requirements Checklist

Energy Consumption: A determination must be made on how much data is needed to train and operate the systems. If necessary, methods must be developed to minimize the amount of data needed.

Economic Requirements Checklist

Working Conditions and Jobs: When digital crowd-workers label datasets, they must be paid fair wages and be provided with good working conditions.

PHASE 3: DEVELOPMENT

How is an AI model developed? Which model should be selected? How is the model trained? These questions have far-reaching implications, particularly for the social and environmental sustainability of AI systems. Energy consumption must be minimized during the development phase. Risks must also be contained.

General Requirements Checklist

Responsibility: A body set up for this purpose should ensure compliance with the Code of Conduct during the development and evaluation of the models.

Documentation: Model cards must be created to document the experiments performed during development and the corresponding key performance indicators.

Stakeholder Participation: Individuals who may be affected by the AI systems or who are expected to use the AI systems must be involved in development to ensure the prevention of negative outcomes.

Social Requirements Checklist

Non-Discrimination and Fairness: Tools and evaluation criteria for measuring fairness must be implemented.

Technical Reliability and Human Supervision: Model weaknesses must be identified and corrected, and control measures defined.

Inclusive and Participatory Design: Inclusive design principles must be observed, including in the design of user interfaces.

Cultural Sensitivities: Systems must be developed to be applicable in different local contexts – through training with local datasets, for example.

Environmental Requirements Checklist

Energy Consumption: The energy efficiency of systems must be recorded. Preference should be given to pre-trained models and models with lower complexity. Methods should be deployed for optimizing energy efficiency.

Greenhouse Gas Emissions: CO_2 emissions generated in the development process must be rigorously recorded. Selection of appropriate training locations and timing will increase CO_2 efficiency.

Sustainability in Use: A system's resource consumption must be recorded and compared to its sustainability potential.

Indirect Resource Consumption: The resource efficiency of the hardware must be recorded and optimized.

PHASE 4: IMPLEMENTATION

During the deployment and operation of AI systems, as well as during additional training runs, risks to data subjects must be monitored and steps taken to ensure data protection. During the application phase, the consumption of energy and resources must be kept low.

General Requirements Checklist

Stakeholder Participation: Stakeholders must be consulted on new releases.

Risk Management: The potential risks identified must be monitored, and any new risks that may arise must be identified.

Social Requirements Checklist

Transparency and Responsibility: Stakeholders and end users must be informed about the use and operation of the AI system.

Non-Discrimination and Fairness: The AI system's decisions must be evaluated to determine whether they meet previously established fairness standards.

Self-Determination and Data Protection: Persons whose data is used must receive information about that use. Simple consent or revocation options must be guaranteed.

Inclusive and Participatory Design: Barrier-free accessibility as well as access for disadvantaged groups must be ensured.

Environmental Requirements Checklist

Energy Consumption: Energy consumption during deployment must be documented and optimized.

Greenhouse Gas Emissions: CO_2 emissions must be recorded and optimized during use.

Sustainability in Use: It must be determined whether the use of the AI system can be made more sustainable by, for example, conserving resources.

Indirect Resource Consumption: Resource efficiency of the hardware must be recorded and optimized.

Economic Requirements Checklist

Working Conditions and Jobs: When deployed in the workplace, impacts of the system on the workforce, working conditions and potential job losses must be determined and, if need be, minimized.

AI Terribilis

AI systems aren't just harmful to the environment due to the immense amounts of energy and resources they consume. Their production also frequently involves the exploitation of workers. Furthermore, the trend in AI development is all-too-often in the wrong direction: Toward larger and larger systems rather than sustainability.

Ignoring Inference When Calculating Resource Consumption

When discussing the environmental impact of AI systems, the primary focus tends to be on the volume of resources consumed during the development and training phases of Machine Learning models. In most cases, in fact, the figures provided in conjunction with such models tend to refer to just these phases. However, a big question mark hangs over the utilization phase of AI systems. In technical jargon, this is called the "inference" phase. The development and training of AI models are very complex processes and consume a relatively large amount of energy. At the same time, though, the number of processes during these phases is limited, and they are usually completed within a foreseeable timeframe. Each utilization of an AI system during the inference stage, on the other hand, usually consumes relatively little energy. However, inference

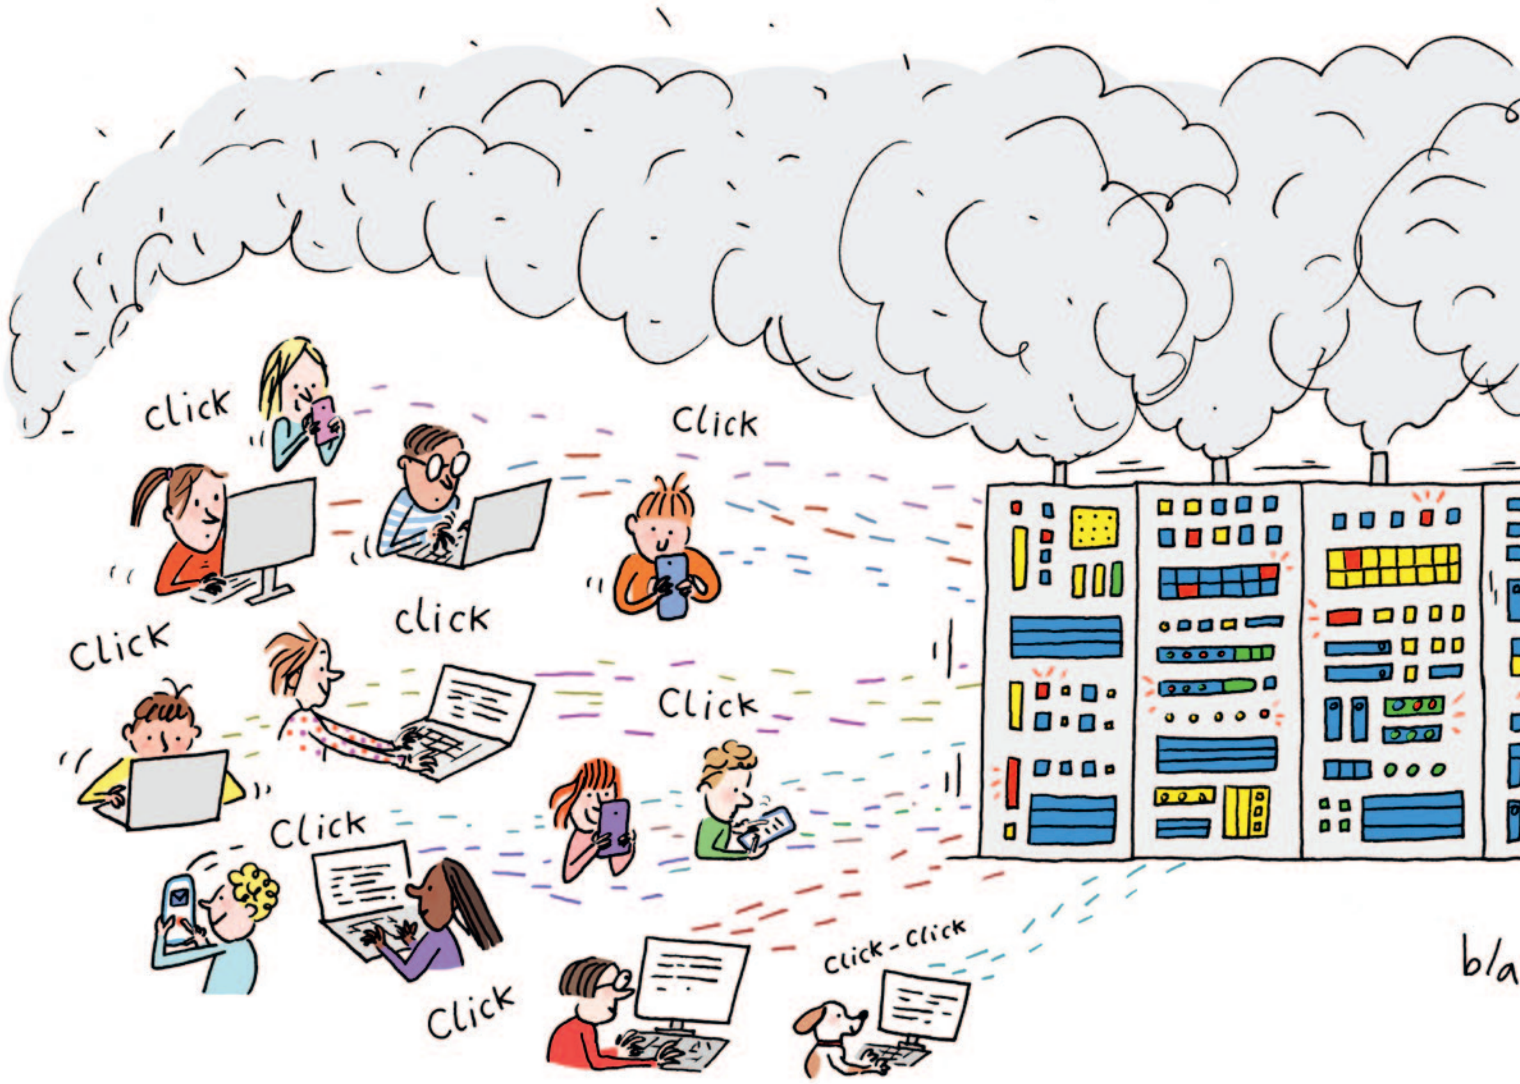

can take place extremely frequently. In late 2022, Facebook AI researchers concluded in a scientific paper that Facebook data centers performed trillions of inference operations each day, a figure they say has doubled in the last three years. The significant growth in inference has also led to an expansion of the infrastructure required to support it, the researchers say. Between the beginning of 2018 and mid-2019, the number of servers devoted specifically to inference at Facebook's data centers increased by 2.5 times, according to the study. At a company like Facebook, this volume of inference comes from things like recommendations and ranking algorithms, for example – algorithms that are used each time Facebook's nearly 3 billion users worldwide access the platform and view content in their newsfeed. Other typical applications that contribute to high inference rates on online platforms include image classification, object recognition in images, and translation and speech recognition services based on large language models.

Even if the amount of energy consumed by each inference operation were minimal, total resource consumption is still likely to be immense due to the sheer volume of operations and the infrastructure they require. The CEO of *Nvidia*, one of the largest processor manufacturers, and executives at *Amazon Web Services* (AWS), one of the largest cloud computing providers, announced back in 2019 that inference is responsible for approximately 90 percent of the costs of the entire Machine Learning process.

Because costs are closely linked to the computing power necessary, scientists have concluded that the emissions produced in the inference phase of AI models are likely to be significantly higher than those produced during the development and training phases. This presumption is supported by internal figures from facebook, which confirm that for in-house systems, resource consumption during the inference phase can be, depending on the application, far higher than during development and training.

As such, it would be negligent to disregard the inference phase when calculating the energy consumption of AI systems. When determining the resource consumption of automobiles, after all, we don't ignore the gasoline consumed while driving.

Moderators: Exploited to Train AI

Online moderators and clickworkers provide data for training AI systems under poor working conditions. Their plight is often overlooked in the AI debate.

Facebook says that Artificial Intelligence technology is a key component in moderating posts on the platform. "AI can detect and remove content that goes against our Community Standards before anyone reports it," the company states on its website. AI is also a central component of ChatGPT's automated text production: Toxic content must be removed to make the ChatBot's output suitable for widespread use. Indeed, all major online platforms likely rely on the support of AI for content moderation. One reason is that deploying AI systems is cheaper than using exclusively human moderation. Moreover, the job is psychologically challenging:

Moderators are constantly exposed to disturbing content circulating on the internet.

But AI moderation systems also rely on human decision-making, and not just for murky moderation issues. Moderators provide training data to the systems, and are thus a mandatory prerequisite for the systems to be developed in the first place.

Nevertheless, moderators are exposed to extremely poor working conditions. Given the psychologically stressful nature of the work, large platforms have an extra obligation to take care of their moderators. But there are instead constant reports of subcontractors of Facebook, TikTok or OpenAI not paying their moderators and clickworkers enough, not offering them adequate psychological support, and exerting extreme pressure on them through constant monitoring and by way of threats aimed at preventing their unionization. The French company *Teleperformance*, which provides moderation services to TikTok, among other platforms, was recently the target of all these accusations. In response to an investigative report, the Colombian Labor Ministry ordered an investigation into working conditions at *Teleperformance* sites in the country. Subsequently, public pressure on *Teleperformance* grew so great that the *UNI Global Union* was able to reach a worldwide agreement with the company in December 2022 to secure greater rights for workers and better protections for their health.

Such steps are urgently needed to commit AI system manufacturers to fair working conditions along the entire value chain. Current European digital policy projects such as the AI Act, the Platform Worker Directive and the Data Act ignore the problem. But if the development and use of AI systems is to adhere to European values, as formulated in the AI Act, then the EU cannot close its eyes to poor working conditions, and not just in the case of online moderators.

P

PaLM: The Wrong Understanding of Efficiency

Google's 540-Billion-Parameter Language Model PaLM

In spring 2022, Google researchers unveiled a new AI language tool called the Pathways Language Model, or PaLM. Its capabilities include the ability to interpret inputted text and produce new, meaningful text segments of its own. Whereas previous models like BERT (110 million parameters) or GPT-3 (175 billion parameters) attracted attention due to their incredible size, PaLM has now set a new record, with 540 billion parameters. Parameters are values that a Machine Learning model learns during the training process, and they form the basis for the outcomes the model then produces.

The number of parameters also determines the number of computing opera-

tions that must be performed and, thus, the amount of energy consumed. It is likely that a model with 540 billion parameters also consumes an extremely high amount of energy – during development, during training and, presumably, also during use.

At the same time, the Google research team claims to have made a breakthrough in terms of training efficiency. This advance has been achieved, the researchers say, by way of newly developed hardware called Tensor Processing Units (TPUs), which enable accelerated computation, and through new strategies in parallel computing. Google says it was able to significantly reduce the amount of time it took to train the vast model, thus saving energy.

A single training run of PaLM at a Google data center in Oklahoma, which obtains 89 percent of its energy requirements from carbon-free energy sources, resulted in 271.43 tons of CO_2 emissions. That is roughly the equivalent to the emissions produced by a fully occupied commercial jet during 1.5 flights across the United States.

Comparative values regarding the emissions produced by previous models during training are mostly based on estimates. As such, one can only assume that around 270 tons of CO_2 emissions for a system as large as PaLM represents, from a relative point of view, a significant improvement. But the question remains as to why a far more efficient hardware innovation and new training methods were only deployed to make models even larger, rather than to improve the energy efficiency of smaller, yet still quite substantial models. That isn't just irresponsible from the perspective of resource conservation. Such vast models also make it more difficult to detect and remove discriminatory, misogynistic and racist content from the data used in training.

Machine Learning research has not yet focused sufficiently on resource conservation – and it's not just Big Tech companies that must be held accountable on the issue. The example of PaLM once again clearly shows that the mentality of "the bigger the better" continues to dominate Machine Learning research, which stands in direct contrast to the urgent need to reduce resource consumption in the entire digital sector, especially in the resource-intensive AI branch.

Furthermore, Google's emphasis on the comparatively low emissions produced during the training of PaLM is misleading. The training of a model never reflects the total amount of emissions generated – indeed, it is often just a fraction of that total. To be able to make a comprehensive determination regarding the resource efficiency of specific AI systems, emissions produced during development and application must be quantified along with the emissions associated with the hardware used. Google could at least have specified the number of training runs it performed during the development phase and how high emissions were in total. That, though, would likely have produced a significantly different image – and the green finish Google has sought to apply to itself would quickly have flaked off.

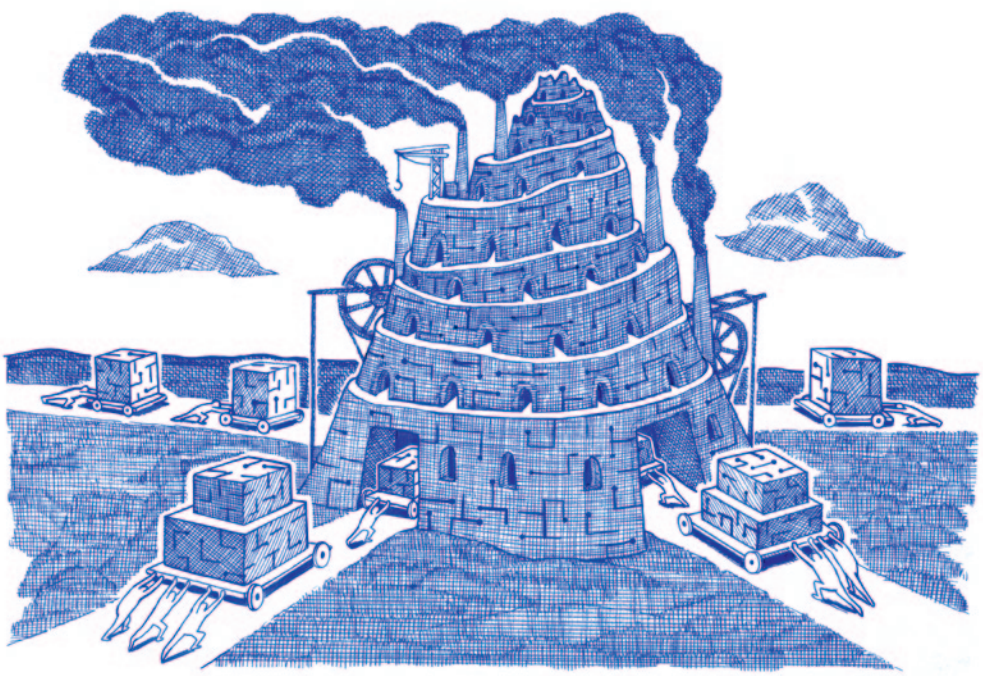

Community Building: Deploying Open Source Against Megalomania

True to her view that the climate crisis is the greatest global challenge we currently face, Machine Learning (ML) researcher Sasha Luccioni focuses her efforts at *Hugging Face* on making Artificial Intelligence (AI) models more sustainable. The AI startup has set out to tackle problems like emissions, bias and discrimination by supporting open-source approaches in the ML community. Luccioni provides insight into measuring the carbon footprint of AI models, while *Hugging Face* Policy Director Irene Solaiman explains how this will help policymakers generate much needed pressure.

According to DALL·E 2, Women and POC Can't Be Authors

Transformer models have become the standard for large language models. Search engines, automated translation services, content moderation systems, speech recognition tools, text-to-image generators and many other applications are based on them. The underlying deep learning models are usually trained on extremely large datasets to deduce intrinsic structures within them. These structures form the basis for the model's automation process while transforming input data into output data – for instance while generating an image based on a text entry. Recently, models were released that can generate either text or images based on text prompts (like GPT3, ChatGPT, Stable Diffusion, DALL·E, etc.). While their capabilities are undoubtedly impressive, they come with serious risks. Usually trained on unfiltered data scraped from the internet, they often create discriminatory, racist, misogynist or otherwise deeply biased content. Researchers at *Hugging Face* have developed a tool to reveal bias inherent in text-to-image generators. The tool enables the generation of prompts for DALL·E and Stable Diffusion by choosing from a list of 150 professions and 20 related adjectives. The Bias Explorer clearly demonstrates how prone to bias these models can be. If you enter the prompt "author," and combine the profession with each of the 20 adjectives available, DALL·E 2 generates 179 images of white men and just one image of a white woman. Stable Diffusion (version 1.4) performs only slightly better, generating 13 images of persons of color out of a total of 180. When it comes to gender representation, Stable Diffusion clearly demonstrates a female bias, generating 140 pictures of female authors out of 180 overall.

Link: https://huggingface.co/spaces/society-ethics/DiffusionBiasExplorer

Sasha, your primary focus at *Hugging Face* is on large language models or transformer models. Language models have become notorious for their immense energy consumption and CO_2 emissions. Would you say that open-source language models are per se more sustainable from an ecological point of view?

Sasha: Open source helps with the recycling of models. Instead of training transformer models once, you can reuse them. All the pretrained models on *Hugging Face* can be fine-tuned for specific use cases. That's definitely more environmentally friendly than creating a model from scratch. Several years ago, the main approach was to accumulate as much data as possible to train a model, which would then not be shared. Now, data-intensive models are shared after training. People can reuse and retune them according to their particular use cases.

Do you also see social or economic advantages?

Sasha: With the size of transformer and AI models growing bigger and bigger, the entry barrier for joining the AI community is becoming correspondingly high, especially for countries that don't have access to the extremely powerful computers being used to create these models. *Hugging Face* has several offers available for such cases – for example, the ability to query a large language model using an API, so you don't need to run it on your own computer. This makes such models more accessible.

Is *Hugging Face* receiving any political support for its work on boosting AI sustainability?

Irene: The regulations pertaining to AI that have been issued in recent years haven't focused particularly on sustainability. Measuring carbon emissions has likewise not been prioritized, but there aren't a lot of tools available to adequately measure them. We find ourselves in a dilemma: There is an urgent need for policymakers to up the pressure, but to do so, they need emissions data. Political regulations, however, do not currently include a requirement to deploy tools for measuring emissions – which means that policymakers don't have the data they need.

What political approaches are currently being pursued to increase the sustainability of AI?

Irene: The European Union's Artificial Intelligence Act is one of the most robust and prominent approaches to regulating AI in the public's interest. A lot of policies and regulations are necessarily coming from countries with higher gross domestic products, such as the Algorithmic Accountability Act in the United States and the AI and Data Act in Canada. The Algorithmic Accountability Act does not explicitly include sustainability, but I appreciate the emphasis it places on impact assessments. Decision-makers need more guidance on the impact of AI systems, including CO_2 emissions. Such information will give them a greater understanding for the importance of developing appropriate tools.

> "How much CO_2 will be emitted through deployment depends on a number of factors, including the hardware you're using and where the computing is being done."

At *Hugging Face*, you have also developed a tool to uncover bias in language models. How does it work? And what kinds of bias does it detect?

Sasha: These models are trained on data scraped from all over the internet. By avoiding a specific and limited data source, they're supposed to be relatively impartial. But when you use them in a downstream AI application, outputs are generated that you may not have expected. To figure out where potential biases could emerge, you have to make AI applications take decisions or make predictions in different situations. We've been working on ways of prompting the models by giving them bits of text and making them complete them – based on a pronoun, for example as in "She should work as" and "He should work as." If a model continues, "She should work as a nurse" and "He should work as a computer scientist," you can immediately see how biased, how toxic, it is. Such negative stereotypes are one example of system bias, which we can document for every AI model by creating a report card.

It is also possible to perform emissions-based searches for models in the *Hugging Face* database. Is this feature used frequently?

Sasha: Most of the emissions numbers we have are from training. We don't have many numbers from deployment. A lot of people are interested in how much CO_2 will be emitted through deployment, but that's extremely complicated, because it depends on a number of factors, including the hardware you're using and where the computing is being done. Without knowing those factors, it's impossible to provide information on the emissions. In order to do so, you would need to evaluate different architectures, different models, different GPUs, etc. Still, a lot of people would find such information extremely useful.

Why is it important to measure ML model emissions, and how can this be achieved?

Sasha: If people start using tools to measure the emissions of their ML models and disclose that information, we can start making decisions about AI models based on facts and figures. Tools like Code Carbon calculate a model's carbon footprint in real-time. It's a program that runs in parallel to any code and will estimate the carbon emissions at the end. We also run a website allowing you to enter information like training hours and the type of hardware used. It then provides an estimate of the system's carbon footprint. It is less precise than Code Carbon, but it still gives you an idea.

How can the use of more sustainable AI systems be promoted?

Sasha: I think that bottom-up approaches work, especially in terms of research. At conferences, we are constantly asked for more information. But there's the issue of reproducibility: A lot

AI Lifecycle and CO_2: The Emissions Never Stop

There is hardly any information available about the energy consumption of AI systems and the CO_2 emissions they produce. This state of affairs makes it more difficult to develop targeted political approaches aimed at reducing these emissions. It is well known that data centers, like the production and operation of all hardware, make a significant contribution to global CO_2 emissions. And they provide the necessary infrastructure for the operation of AI systems. The lack of reliable numbers on the emissions produced by the usage of AI systems is an additional factor.

Sasha Luccioni, Sylvain Viguier and Anne-Laure Ligozat have taken the first step toward closing this information gap. They have produced an estimate for the amount of emissions produced by the language model BLOOM (176 billion parameters) over much of its lifecycle. The result: During the training of BLOOM, around 24.7 tons of CO_2 equivalent were produced, if only direct energy consumption is taken into account. If, however, processes such as hardware manufacture and operational energy consumption are also included in the calculation on a pro-rata basis, the emission values double. Training alone, in other words, is not sufficient as a reference variable when calculating the emissions produced by AI systems. Measurements and methodically stringent calculations must cover their entire lifecycles to sensitize companies, developers and researchers and to initiate targeted political regulations.

Source: "Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model" by Sasha Luccioni, Sylvain Viguier, Anne-Laure Ligozat
https://arxiv.org/pdf/2211.02001.pdf

“I do see cloud providers taking advantage of carbon offsetting, and some are switching to renewable energy sources.”

Sasha: When we look at the infrastructure, there are both positive and negative developments. Hardware development is making rapid progress when it comes to computing efficiency. If you compare a GPU from this year to one built two or three years ago, there’s a significant difference. It’s literally 10 times faster. But with this positive development comes a negative one, because that efficiency leap means that people are doing more computing. It’s a Catch-22. If we kept the size of our models and the amount of computation needed at a constant level, we would definitely be going in the right direction. But since both are growing so fast, it’s hard to say where we might end up. I do see cloud providers taking advantage of carbon offsetting, and some are switching to renewable energy sources. On the other hand, though, the concept of “the bigger the better” in AI modeling is getting out of hand.

of research can’t be reproduced because it is highly contingent upon specific factors. This is something the AI community has been trying to tackle by implementing certain guidelines. If you submit a paper, you have to disclose parameters X, Y and Z. You also have to make your code and data freely available. In terms of sustainability, there have to be similar measures in place pertaining to efficiency or accuracy. Only then can we compare different models. We have to provide a technical procedure that a broader community can adopt.

Irene: A lot of policy conversations I’ve been involved in have focused on lowering the regulatory burden on small and medium-sized enterprises, since these companies have fewer resources than Big Tech. Since smaller companies are less likely to have the infrastructure for analyzing their carbon emissions, we can’t expect them to be responsible for monitoring them.

What could policymakers do to help make the CO_2 emissions of AI models more transparent?

Irene: Something we’ve been working on is documentation. We need more guidance from policy institutions on what, specifically, would be helpful to report over and above the information included on model cards. A lot of governments have been asking the industry for more information about models without specifying what aspects of AI sustainability the industry and developers should report. And we also need to know how to report that information in ways that are understandable to high-level policymakers, who may not have a technical background. Developers definitely need more information if we want them to think about how their systems can become more sustainable.

Nonetheless, we keep building unsustainable AI systems. Are we stuck with an unsustainable AI infrastructure for the coming years?

DR. SASHA LUCCIONI

... is a Research Scientist at *Hugging Face*, where she works on the ethical and societal impacts of Machine Learning models and datasets. She is also a Co-Chair of the Carbon Footprint Working Group within the *Big Science Workshop* and a member of the board of directors of *Women in Machine Learning* (WiML).

IRENE SOLAIMAN

... is an AI safety expert and Policy Director at *Hugging Face*, where she is conducting social impact research and building public policy. She also advises responsible AI initiatives at the *Organization for Economic Cooperation and Development* (OECD) and the *Institute of Electrical and Electronics Engineers* (IEEE). Irene formerly developed AI policy at the *Zillow Group*. Before that, she led public policy at *OpenAI*, where she initiated and led bias and social impact research. Irene holds a master’s degree in public policy from *Harvard University*.

AI Gives You Wings?

The Opportunities and Risks of AI in Energy Supply

It has been said time and again that the transition to a carbon-neutral economy can only succeed with the help of Artificial Intelligence. But it still isn't clear whether AI systems can actually deliver the desired benefits. As part of the research project "SustAIn: The Sustainability Index for Artificial Intelligence," researchers from the Institute for Ecological Economy Research (IÖW) analyzed the diverse interests that shape the discussion about the opportunities and risks associated with AI in energy supply. In doing so, they examined 22 strategy and position papers from the German, European and international context in addition to conducting interviews with experts.

What hopes do stakeholders have for AI when it comes to energy supply?

Josephin Wagner: Our study shows that they primarily expect the use of AI to generate considerable opportunities. AI systems are designed to make energy supplies more efficient – by optimizing processes, for example, processing data in real time or generating forecasts. Automation can also help determine early on whether equipment needs to be repaired. Stakeholders are hoping to reduce costs and maximize profits through the use of AI. At the same time, AI is also expected to drive the transition to clean energies and provide solutions that account for the escalating complexity of the energy system. As the number of producers and consumers in the system grows and the energy sector becomes increasingly digital, the amount of data needing to be processed will expand significantly. Stakeholders believe that AI systems will be particularly adept at managing the mass of data.

Friederike Rohde: Our analysis shows that they expect AI applications to optimize energy consumption through data-driven load and feed-in forecasts and to improve management based on that data. The optimized adjustment of power generation and consumption through better forecasting can reduce the burden on power grids. Stakeholders thus believe that AI systems can make an important contribution to the security of supply and system stability. AI will also be used to better match renewable energy to energy needs and integrate it into the energy system. AI applications are expected to provide the necessary precision needed for optimization services. The automated and adaptive processing of consumption data is seen as a necessity for providing these services at a reasonable cost in the first place.

What risks have stakeholders identified?

Josephin Wagner: The documents we analyzed addressed classic risks – in the field of cybersecurity, for example – but also the energy and resource consumption required to operate AI systems and their infrastructure. In the expert interviews, though, we also addressed the question as to which stakeholders are actually benefiting most from the increasing use of AI. The interviewees identified the inherent risk that especially those who already have a lot of data, such as large transmission

“In the 22 documents we analyzed on the role of AI in the transition to clean energy, the opportunities are given greater emphasis than the risks.”

system operators, could stand to benefit the most. Small players like municipal utilities, on the other hand, would first need to build the skills and data infrastructure necessary to reap the benefits of intelligent optimization. The interviewees were also critical of the fact that financial and human resources are currently being invested disproportionately in the development of AI in the energy sector. They argue that the further development of AI applications in other areas has been stalled by the fixation on the energy sector, even though the benefits of AI applications in energy supply are far from proven.

What attracts greater attention in the discussions: the opportunities or the risks?

Friederike Rohde: In the 22 documents we analyzed on the role of AI in the transition to clean energy, the opportunities are given greater emphasis than the risks. However, the interviews with experts also revealed the crux of this technology: On the one hand, people want AI to help them better deal with systemic complexity. On the other, though, the use of AI also makes the system even more complex, which further increases the risks. But this ambivalence is rarely discussed.

Josephin Wagner: In terms of the sustainability of AI technologies in the energy system, our case study makes it clear: The amount of energy required for training AI models to optimize energy consumption management is negligible. Thus, it makes a lot of sense to rely on smart optimization and AI technologies in the energy system. Nonetheless, the positive expectations people have are exaggerated, because there are still hurdles that AI cannot remove – such as the development of an appropriate regulatory framework or the participation of citizens in the design of the energy system. The reality is that many pilot projects fail due to the regulatory framework. These conditions must be shaped in parallel with the development of the technology so that AI can actually make a positive contribution to the transition to a carbon-neutral economy.

FRIEDERIKE ROHDE

… is a sustainability researcher and technology sociologist at the *Institute for Ecological Economy Research* (IÖW). She is pursuing her PhD at the *TU Berlin* and works on socio-technical futures in the context of digital transformation, social innovations and algorithmic decision-making systems.

JOSEPHIN WAGNER

… is a research associate at the *Institute for Ecological Economy Research*. In the research field of environmental economics and policy, she focuses on digitalization and social change as well as on the economic and institutional analysis of environmental policies.

Don't Believe the Hype: What AI Really Means for the Energy Supply

Does AI help conserve resources, or does it increase resource consumption? This question cannot be answered in general terms – it must be examined on a case-by-case basis. In the energy sector especially, AI is currently being given the benefit of the doubt. Andreas Meyer from the Distributed Artificial Intelligence Laboratory (DAI Lab) at the Technical University of Berlin applied a computer simulation to a regional project to analyze the conditions under which AI can contribute to the reduction of CO_2 emissions and when it cannot.

The WindNODE Project in Berlin

More than 70 partners from business, science and industry have joined forces for the WindNODE joint project to advance the shift to renewable energies. The *DAI Lab* at *TU Berlin* and an experimental quarter in Berlin's Prenzlauer Berg district also participated. The quarter complex includes six buildings with a total of 224 apartments that are heated by a local heating plant. The quarter concept also includes an intelligent energy management system, which enables flexible access to existing renewable energy offerings.

How Can AI Contribute?

While AI applications can be used to more efficiently deploy locally generated renewable energy and thus conserve resources, the AI applications themselves also consume resources. To ensure effective conservation, AI applications must consume less than they save. AI systems are particularly helpful in producing forecasts for how much energy will be consumed in the quarter at a specific time, and forecasting when and how much energy can be produced via photovoltaic facilities. The diagram on the next page depicts the model of the residential quarter that served as the basis for the simulation study. The forecasts produced by AI ensure that the energy saving system (ESS) and the quarter's hot water reservoir can be used most efficiently.

What Was Simulated?

Such forecasting systems can exhibit a low level of complexity if, for example, statistical models are used. If they are based on a Deep Learning approach, however, they can be extremely complex, while a traditional Machine Learning approach results in mid-level complexity. The greater the complexity of a forecasting model, the more energy it consumes in the development, training and usage (inference) phases. Computer simulations can provide insight into whether the elevated resource demands of more complex AI systems pay for themselves in the form of a more efficient usage of renewable energies.

Results of Simulation Studies

The example of the Berlin quarter makes clear that AI systems don't always live up to expectations. The studies completed as part of the SustAIn project show that the resource consumption of AI models aimed at optimizing the feed-in of renewable energies in the quarter was not particularly high. At the same time, however, the savings achieved by those models were relatively small. The largest savings potential was produced by the infrastructure, especially the energy saving system, which allows for the more flexible use of locally produced energy.

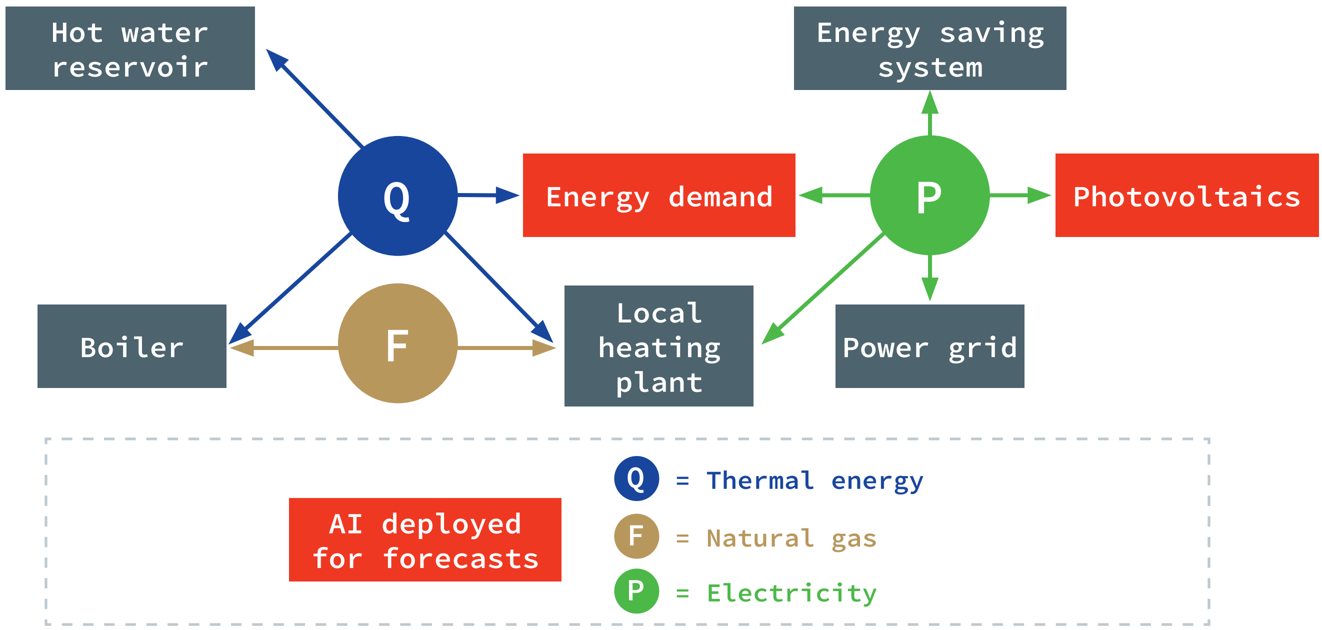

How Much Energy Do Different AI Systems Require?

The simulations were performed for three AI systems of differing complexity, all designed to increase the share of renewable energy in the electrical power supply of the residential quarter in Prenzlauer Berg. As the table below shows, the model XGB (XGBBoost – Extreme Gradient Boosting) consumes the least amount of energy. It is based in traditional methods of Machine Learning and consumes a total of just 0.38 kWh during development, training and usage. That makes its consumption even lower than that of the less complex statistical model SARIMA (Seasonal AutoRegressive Integrated Moving Average). As expected, the two Deep Learning models – ANN (Artificial Neural Network) and LSTM (Long Short-Term Memory) – consumed the most electricity.

Digitalization Scenarios: More Renewable Energy through AI?

Do AI applications increase the share of renewable energy in the quarter's electricity supply? As part of the case study, the period from October to December was simulated on the basis of energy consumption data from the quarter. Three different digitalization scenarios were modeled:

Scenario 1: Low Level of Digitalization (LLD)

An energy management system for the quarter with a low level of digitalization served as the base scenario. Excessive heat from the local heating plant is absorbed by a hot water reservoir. To cover peak consumption periods, a boiler is available. Energy for the quarter

Forecast models	Electricity consumption (kWh)				CO_2-emissions (development only, in kg)
Model	Development	Training	Inference	Total	
Low Complexity					
SARIMA	0.88	0.21	0.000018	1.09	0.38
Medium Complexity					
XGB	0.36	0.03	0.000008	0.39	0.16
High Complexity					
ANN	1.71	0.13	0.000025	1.85	0.73
LSTM	5.3	0.36	0.000079	5.67	2.27

Digitalization Scenario	Energy Consumption (%)
Low Level of Digitalization	
Base scenario	45.13
Base scenario with energy saving system (ESS)	58.32
Medium Level of Digitalization	
Base scenario (ESS) augmented with heating behavior in the quarter	58.73
High Level of Digitalization	
Base scenario (ESS) + SARIMA	58.76
Base scenario (ESS) + HD-XGB	61.47
Base scenario (ESS) + HD-ANN	61.73
Base scenario (ESS) + HD-LSTM	62.39

is produced by a photovoltaic facility. In addition, the effects of the installation of an additional energy saving system on overall energy consumption is analyzed.

Scenario 2: Medium Level of Digitalization (MLD)

Based on Scenario 1, data on the heating behavior of the residents is factored in. This information is used to develop an optimized control strategy for the provision of energy and heat.

Scenario 3: High Level of Digitalization (HLD)

In this scenario, the forecasts from the different AI systems are used to optimize energy management.

The results show that the greatest difference isn't achieved through the deployment of AI, but by using an energy saving system in the low digitalization scenario. By doing so, the share of renewable energy in the quarter's electricity supply is increased from 45 percent to 58 percent. The use of AI models only produces an additional 4 percentage points, to 62 percent. In months with more hours of sunlight, this share will likely be slightly higher.

AI Alone Won't Suffice

In the case of the Berlin quarter, complex AI systems were best at producing forecasts for energy consumption and the possible amounts of photovoltaic power that could be produced. The models used for the simulation were all quite economical, but their benefits were limited. AI applications are able to make the energy sector more sustainable by effectively integrating and distributing renewable energy. But they can only realize their full potential within a modern and intelligent network infrastructure. Furthermore, sufficient sources of locally produced renewable energy must be available and storage technologies are necessary to flexibly incorporate that energy. AI, in other words, is not the sole solution.

ANDREAS MEYER

… is a Research Associate at the *Distributed Artificial Intelligence Lab* at *TU Berlin*, where he is researching applications of Machine Learning methods for load forecasting and the sustainability of AI systems.

"AI applications are able to make the energy sector more sustainable by effectively integrating and distributing renewable energy."

More Hope than Concern

We wanted to get a better idea of the expectations raised by the use of AI in smart grids. To do so, we evaluated documents from different stakeholder groups that comment on opportunities and risks.

- Science & research
- Energy industry
- Consulting
- Network operators
- Civil society
- Politics

RISKS

Societal Risks

Cybersecurity is one of the most important issues related to critical infrastructures and is thus mentioned most frequently. But societal risks are also addressed, such as discrimination, transparency shortcomings and the unknown workings of AI systems.

Ecological Risks

Concerns are frequently expressed that model development and digital infrastructures could lead to increased energy and resource consumption. Some have also noted that rebound effects could reduce potential savings.

Economic Risks

Very few risks have been identified here. The possibility that jobs might be lost is rarely raised.

Economic Advantages

It is often said that AI will reduce costs and increase revenues. Many expect stronger customer loyalty as well.

ADVANTAGES AND EXPECTATIONS

Advancing the Shift to Clean Energy

Many consider AI to be indispensable in the transition to a climate-neutral economy, particularly its usefulness in the integration of renewable energies. AI is also expected to promote societal acceptance of renewables in addition to participation.

Optimization & Increasing Efficiency

In the energy sector, AI is expected to help increase the level of automation, improve system efficiency through better forecasting and optimize maintenance processes.

Technical Advantages

AI is credited with enabling greater precision, better data processing and real-time analytics. Occasionally, the expectation is mentioned that new phenomena could be discovered through the use of AI.

Ecological Advantages

It is expected that AI will advance the transition to a carbon-neutral economy, aid with climate protection and ensure boosts in resource and energy efficiency.

With Google as My Neighbor, Will There Still Be Water?

Artificial Intelligence relies heavily on data infrastructure especially data centers. At a recently organized series of talks, researcher Sebastián Lehuedé and his colleagues collected insights from grassroots groups and activists around the globe who are opposed to data centers operated by Big Tech companies like Google, Meta and Microsoft. They're worried about the harmful environmental impact of data centers, which they fear could result in energy instability and water shortages.

Why are activists fighting data centers and what are their main concerns?

The different grassroots groups are active locally and not part of a broader movement. What they have in common are their environmental concerns – specifically the data centers' energy and water consumption. The management and processing of data requires vast amounts of energy. The data centers built in Ireland, one of the world's main poles for this type of infrastructure, are expected to consume 27 percent of the entire country's electricity by 2029. That's quite a lot. Data centers increase the load on the electricity grid, consuming up to 2 percent of global electricity demand. With data centers becoming increasingly efficient, this might change. At the

same time, though, more and more are being built. And then there is the issue of water consumption. Vast amounts of water are required to cool them, which is why data centers in Europe are usually built in colder areas like Nordic countries, where due to lower outside temperatures cooling down the server is easier. In 2019, however, Google planned on building a data center in central Chile, which is home to an increasingly arid Mediterranean climate. According to the first structural design, the Google data center's cooling system required 169 liters per second, and that in an area where people have been struggling with droughts for years. The average data center uses as much water as three average-sized hospitals. There are also other issues, including air pollution. In the Netherlands, activists worried that the construction of a Microsoft

Lower emissions, more drought: Is Google's data center sustainable?

Google has built a data center in Quilicura, just outside the Chilean capital city of Santiago, which it touts as being one of the most efficient and environmentally friendly in Latin America. In fact, a comparison of Google data centers worldwide based on the company's self-reported data shows that the Santiago facility produces comparatively low emissions. Around 69 percent of its energy supply is carbon-free, and it emits "only" 190 grams of CO_2 equivalent per kilowatt hour. Google encourages customers who are interested in hosting an application via the Google Cloud to choose data centers with low CO_2 emissions. At first glance, the data center in Santiago would seem to be the perfect choice.

But it's not that simple. The information provided by Google is self-reported and hasn't been independently verified by anyone outside the company. Like any other data center, the facility in Chile requires minerals that are extracted in environmentally damaging ways as well as hardware that will eventually need to be disposed of as e-waste. In addition, water consumption at the data center in Quilicura is particularly problematic, given that the region suffers from persistent droughts. Looking at emissions alone is hardly sufficient for assessing the sustainability of data centers. This example in Chile underscores once again how Big Tech companies greenwash their business operations by providing selective information, pretending to offer sustainable choices, when they are in fact aggravating existing environmental crises.

data center would affect their agriculture and, beyond that, ruin the landscape.

Are activists concerned that data centers located near their homes will have an immediate impact on their everyday lives?

Patrick Brodie from the University College Dublin thinks that new partnerships between Big Tech and renewable energy companies could lead to the exclusive use of renewable energy for data centers, meaning that ordinary citizens are then denied access to clean energy because of green data centers. In Santiago de Chile, people were uncertain about whether there would still be enough water for them with the Google data center being operated there. There were no reliable studies conducted beforehand to estimate the data center's impact on the local environment and the local communities. In such a constellation, the shift toward sustainable energy becomes a social problem.

Who is responsible for ensuring better protection for local populations – the Big Tech companies, local authorities or national regulatory authorities?

Primarily the Big Tech companies, of course, because they build data centers to profit from them. But we need the authorities to ensure that the Big Tech companies comply with regulatory standards. Alphabet, Meta and Microsoft all issue reports on their electricity and water use, but most of the time, those reports aren't fact-checked. That's why we need third-party actors who keep an eye on how Big Tech enterprises affect the environment and local communities. If this could be done on a global level, then perfect. The problem is that international organizations like UN agencies seem to be quite biased when it comes to the deployment of technologies like Artificial Intelligence. The UN sees AI as a potential solution to the climate crisis without taking into account that the use of AI itself produces carbon emissions that aggravate this very crisis.

> "The problem is that international organizations like UN agencies seem to be quite biased when it comes to the deployment of technologies like Artificial Intelligence."

Are the activists you talk to optimistic about what they can achieve? Where do they see limits to negotiating with Big Tech companies at the local level?

In Chile, after years of sustained activism, the people of Cerrillos managed to negotiate with Google, after which the company decided to use a less water-intensive technology. That was a big achievement. But at the same time, lithium is being extracted in the north of

Chile, which causes considerable harm. We need to take a general look at the entire AI lifecycle process to identify problems. But it is difficult to make an assessment because these companies are so opaque. They can't even say where the minerals that they use to build their technologies come from – probably because they don't know. Under which labor conditions do people assemble the device? How does the training of the algorithms work? Given that activists and scientists haven't yet figured out the extent to which the production of digital technologies damages the environment, I would say that activists are rather pessimistic at the moment.

“This is in line with Western Europe's historical hypocrisy. Europe is infamous for defending values such as democracy at home while not showing too much concern about human rights in the rest of the world.”

The EU is seeking to establish the AI Act to regulate high-risk AI systems, but the legislation isn't really addressing the environmental harm caused by data centers and their energy and carbon costs. The EU emphasizes that it wants to formulate AI rules based on European values, so why is it neglecting the environmental impact of AI systems?

This is in line with Western Europe's historical hypocrisy. Europe is infamous for defending values such as democracy at home while not showing too much concern about human rights in the rest of the world. There is this idea of green technologies, the question is: Who defines what is green? If you want to build electric cars, the flagship product of green technologies, you need to extract lithium. But indigenous communities in Chile are standing up against the extraction of lithium because the environmental and social harms are so staggering. If the environmental problem isn't present on European ground, then it's not surprising that the issue isn't addressed in European legislation. To a large extent, the European lifestyle depends on the global exploitation of minerals and other resources, and digital technologies are no exception.

Have the activists you know identified any way to reclaim the resources from Big Tech companies and use them for their own good?

Not really. The Chilean activists in Cerrillos didn't want to address Big Tech companies like that. This is partly a strategic decision because they thought that demanding water justice would be the most effective approach. By seeing Google just as an actor that demands lots of water, they didn't have to get involved with privacy and other issues. They weren't even aware of how problematic the situation is on a global level, which is one of the reasons we organized the Data Territories conference at the Centre of Governance and Human Rights in Cambridge. In the case of Santiago de Chile, the local population initially tended to be sympathetic with Google. Due to the successful PR work companies such as Google keep doing, people tend to think: “Oh, that's great, we'll get more innovation, and more jobs,” when the construction of a data center in their home area is announced. But through their experiences, I do think that digital rights activists are becoming more and more aware of the nature of these actors. And I'm optimistic that through the actions and the educational work of environmental activists, people will realize just how harmful the influence of their new Big Tech neighbors is. And then we can talk about alternative forms of organizing and alternative technologies. That hasn't really happened yet, but it might in the future.

DR. SEBASTIÁN LEHUEDÉ

... is a postdoctoral scholar at the *Centre of Governance and Human Rights* at Cambridge University. His research focuses on the regulation of digital technologies, with an approach influenced by Latin American critical thought that seeks to decolonialize bodies of knowledge. Lehuedé's current project explores the nexus of digital and environmental rights in Latin America. His research has been published in several peer-reviewed journals, including *Information, Communication & Society*. He also writes for *Open Democracy and Progressive International.*

Can Self-Driving Minibuses Reduce Our Carbon Footprint?

AI and Mobility in Rural Regions

If Germany wants to achieve its climate goals, the transportation sector must make a considerable contribution. In 2021, the sector produced around a fifth of all greenhouse gas emissions in Germany. There is plenty of room for cutting those emissions. According to the 2022 Climate Action Report, it is likely that the goal of reducing transportation sector emissions to a maximum of 85 million tons of CO_2 equivalent by 2030 will be missed by approximately 41 million tons of CO_2 equivalent – if there are no changes to current trends. One reason for the transportation sector's poor emissions performance is the ongoing dominance of cars, i.e., motorized individual passenger vehicles. Private vehicles accounted for 74 percent of all passenger traffic in 2019, yet bus travel, rail and cycling all produce lower emissions than cars, assuming average occupancy. From the perspective of climate protection, it is necessary to significantly increase their share.

The think tank *Agora Verkehrswende* points out that shifting to public transportation is particularly challenging for the roughly 30 million people who live in rural regions of Germany, whose share of total passenger-vehicle traffic is around 37 percent. Population numbers in rural areas are shrinking, and that has led to lower public transportation occupancy rates. Financially, it has grown more difficult to maintain public transport networks in such regions, and the options available have declined. For many people, that has translated into a longer distance to the next bus or train stop, and thus a higher hurdle in the way of doing without a private vehicle. Those who have no driver's license or their own car face limited mobility in rural regions. The result is a lack of equal access to mobility, which makes equal participation in society more difficult.

Digital networks simplify access to car-sharing options and to public transportation, which is why they are frequently touted as a driving force behind the mobility transition. But whereas digital options in densely populated urban areas are a lucrative business for the mobility industry, sparsely populated rural regions with low user rates remain unattractive. AI-supported mobility could change this. Autonomous, networked minibuses can lower labor expenses and be flexibly deployed in rural areas. Such an affordable and efficient development in sparsely populated areas should encourage more people to give up their cars. The *Association of German Transport Companies* believes that autonomous buses have great potential, and in its immediate plan for the mobility revolution, it has formulated the goal of advancing automation in public transport by 2030. The development of autonomous and networked vehicles is also being pushed at the political level. Since 2016, for example, the *German Ministry for Digital and Transport* has invested around 256 million euros into 72 research projects in the field, with one of the goals that of achieving Germany's climate goals in the transportation sector. But will such efforts be enough to meet the urgent challenges of implementing a climate-friendly mobility strategy? And what about the sustainability of the AI systems that will be deployed and of the infrastructure they require? The SustAIn team sought answers to these questions by taking a closer look at the example of autonomous buses in rural regions.

JOSEPHIN WAGNER

... is a Research Associate at the *Institute for Ecological Economy Research*. In the research field of environmental economics and policy, she focuses on digitalization and social change as well as on the economic and institutional analysis of environmental policies.

Autonomous Buses in Rural Regions

Currently, 61 projects in Germany focused on autonomous buses in public transport are in various stages of development, most of them in urban areas. Sixteen are based in rural areas.

These pilot projects are testing automation in public transport and setting the course for the regular operation of autonomous buses. Most of these projects receive the bulk of their funding from the federal government, but also from the states of Thuringia, Baden-Württemberg and Saxony-Anhalt as well as from the European Union. The majority of the 16 projects address the challenges of sustainable mobility in rural areas. All of the projects are focused on further developing the technology and boosting acceptance.

1 **Bad Birnbach Shuttle / Line 7015**
https://www.badbirnbach.de/geschichten/autonomer-kleinbus

2 **EASY (Electric Autonomous Shuttle For You)**
https://www.probefahrt-zukunft.de

3 **“Emil” NAF-BUS**
https://www.naf-bus.de/

4 **RABus**
https://www.projekt-rabus.de/

5 **EMMA – Autonomous Driving in Gera**
https://nuts.one/emma-automatisiertes-fahren-in-gera-emma-in-the-city/

6 **SMO – Shuttle Model Region Oberfranken (Hof)**
https://www.shuttle-modellregion-oberfranken.de/strecken-smo/hof-smo#Hof

7 **“AutoNom” NAF-Bus**
https://www.naf-bus.de/

8 **SMO – Shuttle Model Region Oberfranken (Kronach)**
https://www.shuttle-modellregion-oberfranken.de/strecken-smo/kronach-smo#Kronach

9 **Lahr Shuttle**
https://vm.baden-wuerttemberg.de/de/service/presse/pressemitteilung/pid/erste-autonom-fahrende-bus-im-oeffentlichen-strassenverkehr-rollt-in-lahr/

10 **SAM**
https://www.sam-unterwegs.de/

11 **“HFM” NAF-Bus**
https://www.naf-bus.de/

12 **Hambach Shuttle**
https://www.hambach-shuttle.de/

13 **SMO – Shuttle Model Region Oberfranken (Rehau)**
https://www.shuttle-modellregion-oberfranken.de/strecken-smo/rehau-smo#Rehau

14 **Ride4All**
https://ride4all.nrw/

15 **AS-NaSA**
https://www.as-nasa.ovgu.de/

16 **AutoNV_OPR**
https://www.bmvi.de/SharedDocs/DE/Artikel/DG/AVF-projekte/autoNV_OPR.html

Will AI Make Rural Mobility More Sustainable?

Far away from the cities, population numbers are falling even as the average age of rural residents is rising. That has consequences for public mobility: Because it is expensive to serve expansive, sparsely populated areas, public transportation coverage is shrinking.

At the Federal Level:

German politicians are promoting the technological developments necessary for autonomous vehicles. They are also trying to create greater public acceptance for driverless vehicles. But little progress has been made in rural mobility.

At the State Level:

The expectations for autonomous driving outlined in the mobility and digitalization strategies developed by Germany's states are vague. The hope is that autonomous vehicles can expand public transport offerings and promote innovative business models. Sustainability considerations, however, are hardly mentioned in connection with rural mobility.

"In rural areas, pilot projects aimed at ensuring mobility through greater automatization are to be developed."

Digitalization Master Plan Lower Saxony

"Introducing automated driving is an opportunity to provide mobility services, particularly given demographic change and rural development needs."

State Transportation Plan Saxony

A Work in Progress

To effectuate the mobility revolution in rural areas, public transportation and cycling must be strengthened relative to automobile traffic. Available offers must be widely available, inclusive, climate neutral and economically viable. At best, there are only a few hints as to how and when autonomous minibuses might be able to make a contribution.

The Goals of Current Pilot Projects

The test projects presented on page 31 aim to make progress towards achieving federal political targets: Technical systems for autonomous minibuses are being tested and further developed. Of the 16 projects in rural areas, nine focus on the typical challenges rural areas face in the mobility revolution and in providing flexible services to meet demand. Studies are less likely to focus on economic feasibility (three projects), social inclusivity (one project) and the potential for reducing reliance on private cars (one project). Sustainability criteria for the development of the necessary AI systems are not explicitly mentioned in any of the projects.

Note: Though we use the term "autonomous," the minibuses currently in use tend only to be "highly automated." Usually, there is a safety driver in the vehicle to monitor operations and intervene if necessary. The ultimate goal, however, is for the buses to drive without human intervention.

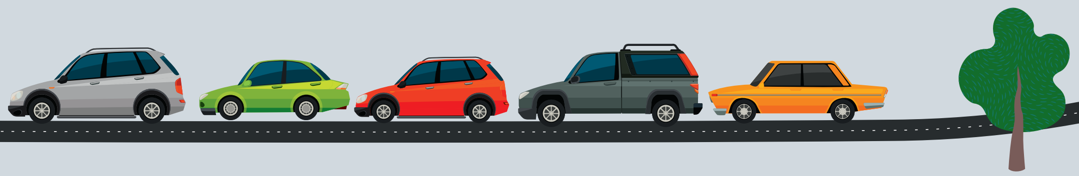

CO_2 Avoidance or Resource Depletion:

How Ecological Is Autonomous Driving?

Self-driving minibuses are seen by many as a possible solution for ensuring rural mobility and keeping people in the countryside connected in an environmentally friendly way – thus enabling rural residents to be part of the transition to a carbon-neutral economy. However, we need to make a realistic assessment of the true extent to which the AI built into these vehicles, and the resources they consume, are actually beneficial to the environment.

Sensors in Autonomous Vehicles

Autonomous vehicles rely on a range of sensors to detect their surroundings and make decisions based on that information:

CAMERAS Object detection and classification, scene understanding, localization and other functions rely on cameras. Typically, image sensors capture the environment, producing data that is then processed by computer vision algorithms.

ULTRASONIC SENSORS These sensors use high-frequency sound waves to detect objects and determine how close they are. They are usually inexpensive and compact, and are used for parking assistance, object detection and obstacle avoidance.

RADAR (Radio Detection and Ranging): This sensor uses radio waves to detect objects and their distance in addition to measuring speeds and angles. The radar function is reliable even in poor weather conditions. It can be used for obstacle detection and tracking, lane detection or vehicle tracking.

LIDAR (Light Detection and Ranging): Remote sensing technology uses laser light for obstacle detection and navigation to quickly and accurately measure distances and create a high-resolution 3D map of the environment.

GPS The satellite-based navigation system provides accurate location and time information. In autonomous vehicles, it enables global localization for navigation and mapping.

ENCODERS These sensors measure the rotational position of the wheels to provide information about the vehicle's movement and position.

INERTIAL MEASUREMENT UNITS (IMUS) IMUs are sensors that measure acceleration and angular velocity. They are used in self-driving vehicles to detect and control movement.

Components	Number	Weight (kg)	Rated Power (W)	Total Rated Power (W)
180° Mono-Layer LiDARS	6	6.6	8.0	48.0
360° Multi-Layer LiDARS	2	3.0	12.0	24.0
Computer	2	20.0	80.0	160.0
Inertial Measurement Devices	1	0.0	0.2	0.2
Router	1	2.5	25.5	25.5
Front/Rear Camera	4	1.5	1.0	4.0
GNSS Radio Module	1	0.2	0.2	0.2
GNSS Module	1	0.0	5.6	5.6
Wheel Encoder	4	0.2	0.15	0.6
Steering Wheel Encoder	2	0.2	0.6	1.2
Touchscreen	1	3.0	15.0	15.0
3G & Ethernet Router	2	1.0	6.0	12.0
4G Antenna	1	0.25	5.0	5.0
GPS Antenna	2	1.2	1.6	3.2
Total		**39.65**		**304.5**

Source: Huber, D., Viere, T., Nemoto, E. H., Jaroudi, I., Korbee, D., & Fournier, G. (2022). "Climate and environmental impacts of automated minibuses in future public transportation". Transportation Research Part D: Transport and Environment, 102, 103160.

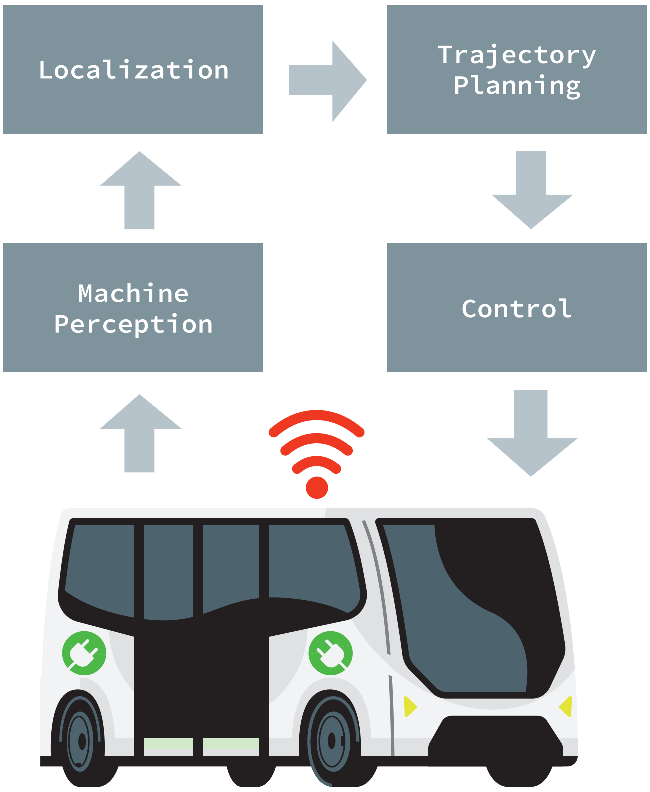

How Much Hardware Is Needed?

In addition to sensors, self-driving minibuses also require additional hardware. The average power consumption of a bus is around 550 Wh/km. The share of hardware components in that consumption is only around 5 percent. However, the production of the components accounts for more than a quarter of all the CO_2 emissions of autonomous minibuses.

AI in Autonomous Vehicles

AI systems are used in machine perception, localization, trajectory planning and control.

Machine Perception

For a self-driving minibus to navigate safely, sensors must scan its surroundings. With the help of Machine Learning, the large amount of sensory data can be processed in real time. Algorithms are trained on large datasets of annotated imagery, LiDAR point clouds and radar data to detect and classify objects. Convolutional neural networks (CNNs) can identify vehicles, other road users, road signs, lanes and traffic lights.

> “The planned trajectory still needs to be executed safely and efficiently. In the control phase, the position, speed and orientation of the vehicle must be determined using sensor data.”

The vehicle uses the high-resolution map of the surroundings created from the data collected by its sensors to plan its route.

Localization

Simultaneous localization and mapping (SLAM) algorithms are often used. Data from the different sensors is processed in real time to create a map of the surroundings, on which the estimated position and orientation of the vehicle is recorded. Deep Learning algorithms such as Convolutional Neural Networks and Recurrent Neural Networks (RNNs) are increasingly being relied on for localization.

Trajectory Planning

To ensure that an autonomous vehicle follows a safe and efficient driving route, algorithms dynamically process information about the surroundings based on mathematical modeling. Machine Learning algorithms are increasingly being used to improve trajectory planning. Deep learning algorithms such as CNNs and RNNs can recognize patterns in sensor data such as camera images and point clouds from LiDAR sensors to estimate the position and speed of other road users.

Control

The planned trajectory still needs to be executed safely and efficiently. In the control phase, the position, speed and orientation of the vehicle must be determined using sensor data. In trajectory following, control algorithms ensure that the vehicle follows the trajectory as closely as possible, taking into account the dynamics of the environment. Actuators such as the steering, throttles and brakes control the vehicle’s movement. For autonomous control, a model of the vehicle is often used to identify important process variables and the dynamic relationships between them (“model predictive control”).

How Resource-Intensive Must the Mobility Transition Be?

Autonomous driving is only possible with the use of many sensors and algorithms. Self-driving minibuses cannot be sold as environmentally friendly lodestars of the mobility revolution without calculating the resource consumption of all their components. Will they contribute to a reduction in emissions by encouraging rural residents to rely less on their cars and more on convenient public transportation options? Or will the resource consumption of minibuses actually produce even more emissions? Might there be simpler means available for advancing the mobility revolution in rural areas? Answers to these questions are not yet available because there hasn’t been an honest debate so far about emissions reductions and the resource consumption of the technology required by autonomous minibuses.

ANDREAS MEYER

... is a Research Associate at the *Distributed Artificial Intelligence Lab* at *TU Berlin*, where he is researching applications of Machine Learning methods for load forecasting and the sustainability of AI systems.

The AI Mobility Revolution in the Countryside:

Optimism versus Reality

The *Association of German Transport Companies* is predicting that the introduction of the first autonomous public transport services in 2025 will be the start of a "robo-shuttle revolution." Politicians have likewise shown increasing support for automating rural mobility with the help of AI. Yet it is by no means certain that doing so will result in decreased usage of private automobiles, which are particularly damaging to the climate. Analyses of the mobility and digital strategies developed by German state governments show that political targets with regards to sustainable transportation are either imprecise or nonexistent.

AI systems for autonomous public transportation are widely considered to be a driver of the economy. In the mobility strategy for the German state of Hesse, a mobility company executive forecasts that the costs of public transportation will drop thanks to AI. The *Economy Ministry for the state of Lower Saxony* even claims to have reached "concrete agreements" with the automobile industry on digitalization. One of the aims of the new business models being developed for public transportation is that of improving rural connections. The hope is that by eliminating the cost of drivers, route frequency will improve, which would increase flexibility for people who do not possess their own cars. Whether the new offerings will be affordable is largely dependent on the net savings produced through driverless operations. The level of those savings, however, cannot yet be assessed on the basis of the analysis of current projects, in which autonomous minibuses are being tested in rural areas. Indeed, efforts are primarily focused on "getting it to work at all somehow [...] and you're happy if you find any solution that's technically feasible," according to an interviewee who researches and applies AI solutions in the processing of sensor data.

> "Whether the new offerings will be affordable is largely dependent on the net savings produced through driverless operations."

The SustAIn analyses show that far too little attention is paid to environmental issues – both by politicians and in the development of the technology. One test project employee with many years of experience in the transport sector says: *"I don't think we will experience a completely automated public transport system in my lifetime."* So far, there are only a few examples of autonomous minibuses ferrying passengers between tourist attractions, university buildings and clinics. As a rule, they travel between 15 and 18 km/h, and have a safety driver on board.

Currently, there is little to indicate that autonomous minibuses will provide a sustainable solution to the deep structural problems facing rural mobility. Political expectations are high but there is a lack of sound evidence. Relying on rail or large buses remains – with almost no exceptions – more climate friendly than the use of smaller vehicles, even if they are electric. Instead of focusing on the question of how we can replace people who drive public transport vehicles, politicians and those working in research and development should be searching for answers to the more fundamental question as to how we can make the shift to sustainable transport.

FRIEDER SCHMELZLE

... is a Researcher at the *Institute for Ecological Economy Research.* His focus is on the social and technological conditions necessary for a sustainable transformation in addition to governance in the areas of digital technologies and energy systems.

Where Digital Rights and Climate Justice Converge

We are living in a climate crisis. The science is clear, and we need to rapidly change to become a more just and sustainable society. This includes investing in a sustainable internet, which is currently the world's largest coal-powered machine. More than building sustainable internet infrastructure, we also need to examine how and where the internet aligns with the movements for climate and environmental justice – and where it works against them. This includes all the environmental effects inherent in the hardware and software systems we use, including the mining that produces the rare earth metals in our phones, the carbon emissions of massive machine learning models and the bias of algorithms and AI systems that may propagate climate misinformation.

Often considered separately, the environment and the internet share much in common. Both are global in scope, they are linked to the exercise and erosion of human rights and they require international cooperation and coordination for their successful continuance. The repercussions of one can also extend to the other: from water rights disputes between data centers and local residents (see page 24 of this issue), to rampant greenwashing misinformation by fossil fuel companies, the internet's ecological consequences are just some of the many complex problems at the intersection of climate justice and technology.

Despite this nexus between the internet and environment, we are only in the early stages of integrating philanthropic funding strategies across these intersections. Through a series of reports supported by the *Ford Foundation*, the *Mozilla Foundation* and *Ariadne*, we can now present research on the implications and intersections of climate justice for digital rights (see links below). While the primary audience for this research is digital rights funders or adjacent technology funders, we believe the work can be useful for other funders and organizations working across issues, given that the climate crisis and technology touch other fields related to human rights, from migration to land tenure and indigenous rights.

> "Enviroment and the internet are global in scope, they are linked to the exercise and erosion of human rights."

The Engine Room is a group whose mission is to support civil society organizations in using technology and data in strategic, effective and responsible ways. The NGO authored a landscape report, *"At the confluence of digital rights and climate & environmental justice,"* which provides an accessible and thoughtful overview of the climate and environmental justice issues that emerge from technological innovation. The report provides an analysis of the environmental toll of digital infrastructures and informs about climate disinformation, open data and climate monitoring, migration justice as well as the increased surveillance of environmental activists and land defenders. It also details the issues and challenges where the climate action and digital rights movements disagree. Finally, it offers recommendations to digital rights funders on how to center the intersections of climate justice and technology in their work.

Four key takeaways from the report *"At the confluence of digital rights and climate & environmental justice"*:

1. **Climate and tech movements can learn from and support each other:** There is a lot of space for relationship building and collaboration between movements centered on climate and environmental justice and those focused on technology justice. Finding the best strategic moments to connect the two movements will be crucial, as will leveraging opportunities for shared learning and deep engagement on specific topics. Learning from other funders about intersectional and trust-based funding approaches can provide a path forward.

2. **The global north must follow leadership from the global south:** Digital rights organizations in the global south have long been working on connecting the impacts of extractive industries and digital technology. Groups in the global north must not only learn from their work, but also follow suit. Increasingly, funders have highlighted the importance of shifting power to those most impacted, believing the peo-

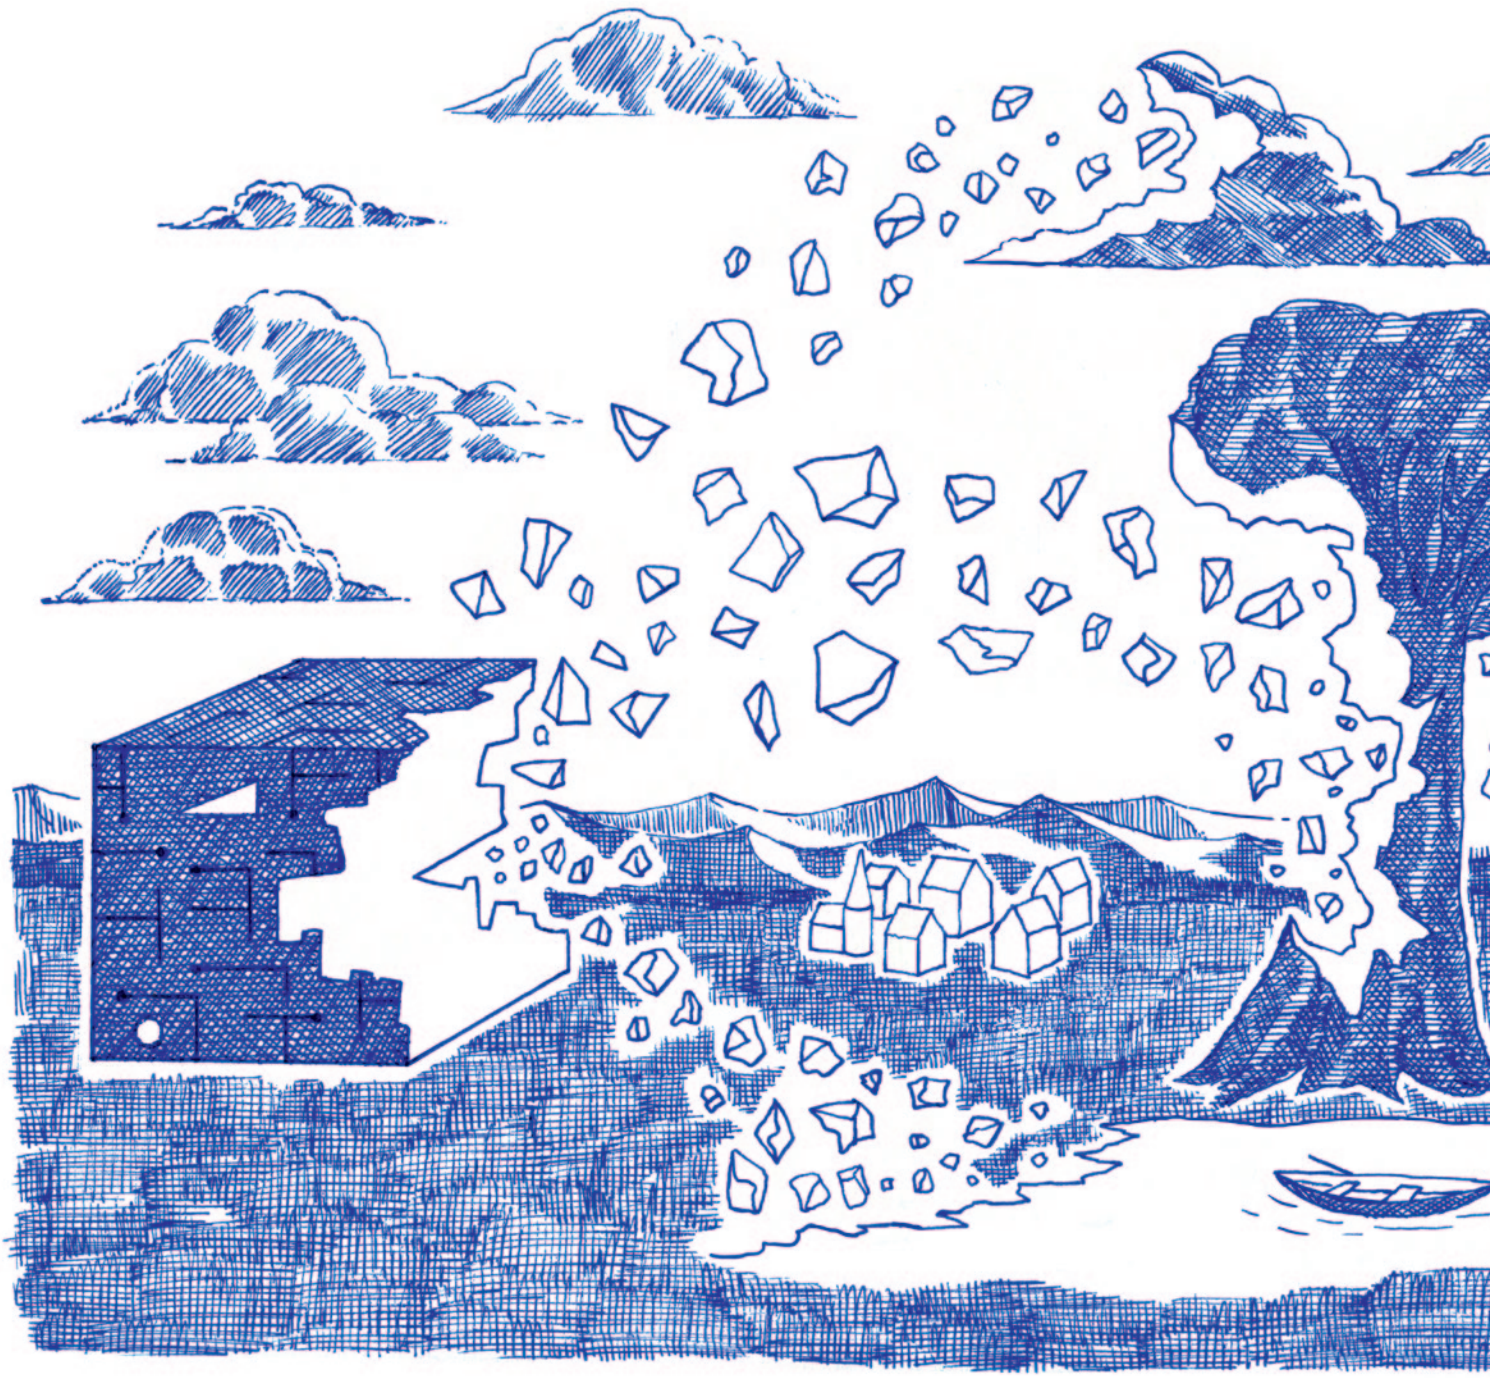

> "The carbon footprint of the internet is an important issue that deserves more attention."

ple who are closest to the problem are also closest to the solution. As we reflect on investment in this intersection, there is an opportunity to "walk the talk" and bring in resource groups in the global south who have clearly demonstrated experience and creativity in developing impactful responses to the climate crisis.

3. Carbon is just the beginning: The carbon footprint of the internet is an important issue that deserves more attention, especially from funders. But there are many more areas that also require urgent change, such as tech companies' inaction against the spread of climate misinformation or the use of advanced cyberweapons to target and harm climate and environmental activists.

4. Data is at the heart of many of the problems and can also be part of the solution: Access to reliable climate and environmental data is key for remedying misinformation, driving policy agendas as well as influencing public understanding and opinion. Tech companies are currently withholding important information related to critical issues like the water and energy usage of data centers and the efficacy of initiatives addressing climate misinformation. At the same time, large data models being developed by the tech sector are a huge driver of emissions and are increasingly central to Big Tech and their efforts to grow. Examining the intersection of climate, environment and data is critical.

Links:

Business for Social Responsibility (BSR): Building a High-Quality Climate Science Information Environment: The Role of Social Media

Association for Progressive Communication (APC): At the Interstice of Digital Rights and Environmental Justice: Four Issue Briefs to Inform Funding

Open Environmental Data Project & Open Climate: Opportunities for the Digital Rights Space / Environmental Justice, Climate Justice and the Space of Digital Rights

DR. MICHAEL BRENNAN

... is a Senior Program Officer on the Technology and Society team at the *Ford Foundation*. He oversees a portfolio of grantees that globally addresses open internet issues through a technical lens and also helps to develop and manage a technology fellows program at the foundation. Michael earned a PhD in computer science from Drexel University.

MAYA RICHMAN

... has been working in the international human rights and technology space for over 10 years. She has facilitated retreats, conferences and workshops all over the world with social justice organizations, foundations and technology companies centered on well-being, digital safety and technology strategy. Her current work is with the *Green Screen climate justice and digital rights coalition*, convening a burgeoning intersection of practitioners and foundations working to build a just and sustainable internet for the planet.

Managing Green Digitalization

Is digitalization destined to worsen climate change, the most challenging crisis of our times? Or can it contribute to solving the crisis? The answer depends on us. Currently, information and communication technologies produce an estimated 2 to 4 percent of all greenhouse gas emissions worldwide, for a total at least as high as Germany's entire share. Global energy consumption from server farms, dataflows and private end devices is rising so dramatically that we must formulate standards and conditions so that the positive effects of digitalization on the climate and the environment outweigh the negative. We can decide whether digitalization will have a negative impact on the environment like burning coal or whether it will become an innovative field comparable to renewable energies. We must now take advantage of the opportunity to push forward environmental and climate protection with the help of digital solutions.

Climate-neutral data centers are justifiably a focus of the current debate about making digitalization more sustainable. But we cannot limit our attention solely to that issue. In the effort to improve sustainability, it would also make sense to limit the amount of unnecessarily generated dataflows. For example, hardly any sustainability criteria have been formulated for software development. Inefficient programming is usually compensated for with faster processors or more powerful hardware components. "Green coding," though, would enable us to significantly improve energy efficiency. Measures to cut down on the use of data, resources and energy must become the new imperatives of software development. We can establish incentives, anchor sustainability in university curricula and develop training programs.

We need European standards for the consumption of energy and resources by software and hardware. Big platforms in particular consume vast quantities of electricity through the exorbitant collection of our personal data for the adver-

“We must make it cheaper to repair a defective device than to buy a new one.”

tising industry. We in the Green Party are pushing in the European Parliament for standards to promote business models that require less data and for greater transparency for consumers. In the future, they should have the ability to compare browsers, search engines, digital marketplaces and social networks based on their electricity and energy consumption. Only then will consumers be able to proactively choose to use, for example, a more sustainable browser. The large digital platforms are currently profiting from internet services that users are not paying for with money: The currently dominant, data-thirsty advertising model is generating enormous sums for large digital companies. And that is precisely where competitors could get their foot in the door through the deployment of more sustainable practices, such as transparent energy savings through data efficiency.

I am also pushing in the European Parliament for the complete melding of the Green Deal and digitalization. I am calling for sustainability criteria for all laws currently being developed at the EU level. The German government must do all it can inside the European Council to ensure strict rules in the Data Act and the AI Act. Because sufficient data has not yet been compiled for the energy and resource consumption of AI systems, clear transparency rules must be included in the AI Act to promote efficient technologies. The risk evaluation framework included in the AI Act must be expanded to include the risks that AI systems may pose to the environment. In addition, the European Union must establish a framework for measuring the environmental impact of AI systems.

The Greens are also committed to ensuring that no more electronic waste from Europe is dumped in poorer regions of the world, exposing local populations to health and safety risks. Digitalization must no longer be rooted in the exploitation of people and the environment. We need improved conditions for the extraction of raw materials around the world by establishing standards for supply chains to Europe. At the same time, we must make it cheaper to repair a defective device than to buy a new one. We can achieve this by making repairs easier to perform, using standardized parts, making replacement parts available for a longer period, through the preparation of repair manuals and the extension of warranty coverage. With the implementation of such binding standards, we can reduce the quantity of electronic waste and save on costs related to disposal and recycling. Binding sustainability labels (which could, for example, provide information on a product’s repairability) or digital product passports would enable more sustainable purchasing decisions in favor of products that are easier to recycle.

How can we prevent a situation in which efficiency improvements made through digitalization are negated by additional consumption – such as when improved data transfer results in the increased use of digital services? To prevent such “rebound effects,” we need control instruments and absolute limits on resource consumption in the digital transformation. Fiscal policy must eliminate environmentally damaging subsidies and shift the tax burden away from labor and toward the consumption of resources.

The systematic reorientation of digitalization toward sustainability will open up new opportunities for European companies. Thus far, a handful of giant corporations have dominated the industry and cemented their position in the market with opaque data-use practices that undermine our democracy, exacerbate the climate crisis and stifle competition. So, it’s worth fighting for a digital Green Deal!

ALEXANDRA GEESE

... has been a member of the European Parliament since 2019 and is the digital expert for the *Greens/EFA* parliamentary group. She was elected Vice President of the group in 2022. She was involved in the negotiations for the Digital Services Act, which regulates digital platforms and social networks. Her main focus areas are democracy in the digital age, sustainable digitalization and gender equality.

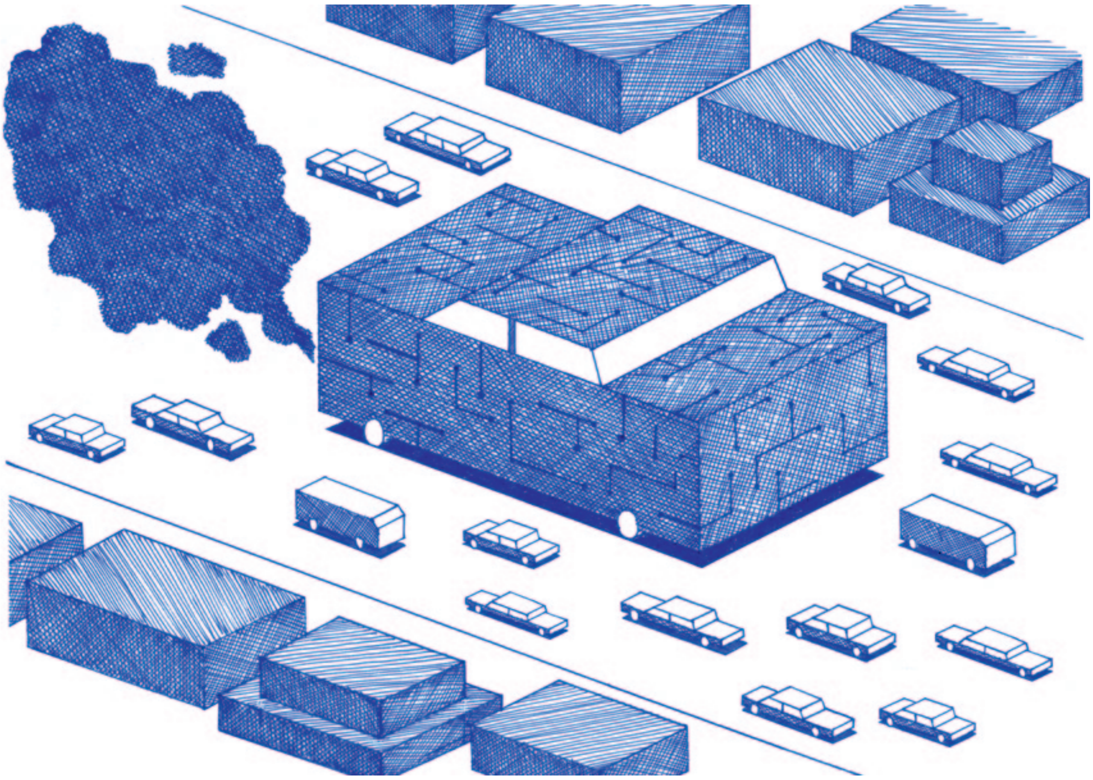

Coming in Our Next Issue of

sustain

In our next issue, we will be exploring the question of who could take responsibility for steering AI development in the right direction. Topics include:

- Online advertising using AI
- Tools aimed at boosting the sustainability of AI
- How politicians should regulate AI to promote sustainability

In a case study, we will demonstrate how opaque marketing AI can produce addictive behavior in online shoppers. We will also be introducing digital evaluation tools that we have created to help AI developers design their systems in ways ensuring that they are more sustainable.

Imprint

Sustain Magazine #2
March 2023
Available online at https://algorithmwatch.org/en/sustain-magazine-march-2023/

Publishers
AW AlgorithmWatch gGmbH
Responsibility (according to § 5 TMG and § 55 RStV): Matthias Spielkamp
Linienstr. 13, 10178 Berlin, Germany

Content Lead
Anne Mollen

Editorial Lead
Waldemar Kesler

Content Support
Andreas Meyer (DAI Laboratory, TU Berlin)
Friederike Rohde (Institute for Ecological Economy Research)
Frieder Schmelzle (Institute for Ecological Economy Research)
Josephin Wagner (Institute for Ecological Economy Research)

Illustrations
Juliette Baily, Kevin Lucbert

Media
Marc Thümmler

Design / Layout
Beate Autering

Translation and Copy Editing
Dacha Media GbR: Daryl Lindsey, Charles Hawley

Output from the project
“SustAIn: The Sustainability Index for Artificial Intelligence”

sustain
Sustainability Index
for Artificial Intelligence

Project partners:

ALGORITHM WATCH

i|ö|w
INSTITUT FÜR ÖKOLOGISCHE WIRTSCHAFTSFORSCHUNG

DAI-Labor
TU Berlin

Supported by:

Federal Ministry
for the Environment, Nature Conservation,
Nuclear Safety and Consumer Protection

based on a decision of
the German Bundestag

sustain

Sustainability Index
for Artificial Intelligence

Project Partners:

AlgorithmWatch

AlgorithmWatch is a non-profit research and advocacy organization whose aim is to monitor and analyze automated decision-making (ADM) systems and their impact on society. One pillar of our work is addressing ADM systems and sustainability. https://algorithmwatch.org/en/

ALGORITHM WATCH

Institute for Ecological Economy Research (IÖW)

IÖW is a leading scientific institute in the field of practice-oriented sustainability research. We devise strategies and approaches for viable, long-term economic activity – for an economy which enables a good life and preserves natural resources. We have been dealing with issues of the future for more than 30 years, consistently finding new and frequently unusual answers. https://www.ioew.de/en/

i ö w

INSTITUT FÜR ÖKOLOGISCHE WIRTSCHAFTSFORSCHUNG

Distributed Artificial Intelligence Laboratory

The DAI-Labor at the Technical University of Berlin views itself as a mediator between university-driven research and industrial applications. With our interdisciplinary team, we create innovations and translate university research into applications for everyday life in close cooperation with other scientific and industrial institutions.
https://dai-labor.de/en/

DAI-Labor
TU Berlin