

sustain

AI and the Challenge of Sustainability



A Different Take on AI:

We Decide What AI Has To Do for Us

Case Study:

Tracking and Personalization

Resource consumption by AI in online marketing

Regulation:

Reducing Negative Impacts

We need political solutions

Standpoint:

Thirsty Artificial Intelligence

AI's water consumption has to be reduced



Project Partners:



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



Dear Reader,

When flyers and advertising brochures pile up in the garbage, it is clear to everyone this is not good for the environment. But what effect do personalized ads on the internet have on the planet? With every visit to a website, hundreds of cookies are planted to feed an invisible yet extremely energy-intensive advertising infrastructure with personalized user data.

Artificial Intelligence devours lots of resources and frequently has deleterious consequences for our society and the environment. The sustainable and sensible use of AI is everyone's responsibility, and we should decide together where to deploy AI – and where not.

Whether AI systems that serve us all will be developed in the future is also dependent on an advantageous policy framework. The first Europe-wide obligation to make AI products more sustainable may soon be introduced in the form of the European Union's AI Act. EU member state governments, however, have not shown much interest in including a requirement that the environmental impacts of AI technologies be measured. Omitting that obligation would be negligent and a failure on the part of European leaders to live up to their political responsibilities. It is only possible to know how environmentally damaging these technologies are if comprehensive measurements are undertaken. The problem will not simply go away if we ignore it. Meaningful data would help us better understand the problem and enable us to put more pressure on policymakers. That's why we need more of it. And that's why we need to start measuring.

This magazine is an invitation to think more about what concrete steps could be taken to better regulate the development and implementation of AI technologies – through greater transparency regarding their consumption of energy and water, through prohibitions of harmful uses of AI and stronger incentives to improve efficiency.

Dr. Anne Mollen
SustAIn Project Leader

Self-Assessment Tool:


How Sustainable Is My AI?

Organizations can now test their AI's sustainability independently. 6

Glossary:

Keeping an Eye on AI

Don't leave it to tech providers to stop AI's aberrations. 10



Key Outputs:

SustAI for Reading and Listening

Our work so far 13



Case Study:

Personalized Online Marketing: AI Technology Gone Astray

How does personalized online advertising work? 14

Analysis:


What Effect Does Personalized Online Marketing Have on Energy Consumption?


Simulations provide data about risks 19

Survey:

Artificial Intelligence and Marketing: A Concerning Combination?

What companies and civil society organizations think... 21





Policy Brief:

**Just Measure It:
The Environmental Impact of AI**

Developing and applying AI consumes more and more resources. What to do? 26

Standpoint:

**The Costly Backbone
of Consumerism**

A closer look at digital ad auctions, the advertising industry's cash cow30



Standpoint:


Water Is the New CO2


We need to focus on AI systems' water footprints32

Looking Ahead:

The SustAI Project Is Ending, Work on Sustainable AI Is Just Beginning

We need to talk about AI systems' sustainability34





How Sustainable Is My AI?

As part of the SustAIIn project, we have done pioneering work in compiling comprehensive indicators that can be used to assess the social, environmental and economic sustainability of AI systems. Using a Self-Assessment Tool, organizations that develop AI in-house or purchase it externally can now test the sustainability of their AI systems with a digital app.

Everyone is talking about Artificial Intelligence, and not just because of the European Union’s planned AI regulation. But this also raises the increasingly pressing question of how sustainable the AI systems used by companies and organizations are. The resource and especially energy consumption of AI can be immense, and in times of an energy and climate crisis, we cannot sweep its ecological impact under the rug. The systems also leave much to be desired in terms of social and economic sustainability. Within the AI industry, market power is concentrated among a few large companies, leading to barriers to entry for smaller firms. Furthermore, exploitative working conditions prevail along the entire AI value chain, and decisions automated with the help of AI can lead to discrimination or to a cultural dominance of Western values, which AI systems latently propagate around the world.

With our project “SustAIIn: The Sustainability Index for Artificial Intelligence,” we have presented the very first comprehensive blueprint with which the sustainability of AI systems can be assessed and improved. Through it, we would like to make a contribution to ensuring that sustainability is put into practice in the development and use of AI. With our Self-Assessment Tool, we are now providing organizations with a questionnaire that can help determine how sustainable their AI systems are. Our traffic-light color scheme helps companies rank their answers, and we also provide recommendations on how to make systems more sustainable.

The Criteria Catalog

The basis for the questionnaire used in the Self-Assessment Tool comes from the criteria and indicators developed as part of the SustAIIn project, all of which reflect the current state of discussion on the social, environmental and economic sustainability of AI systems. We have identified 13 overarching criteria for assessing the sustainability of AI systems. These criteria have been broken down into over 40 indicators, which have been deployed in a way that makes them practical to use.

An Example

To provide one example: The social sustainability criterion “self-determination and data protection” includes the indicator “ensuring informational self-determination.” This self-determination can be implemented by letting those affected by AI systems know how their personal data will be used by the systems. They should be given control over the use of this data with, for example, *opt-in* or *opt-out* functions. In the current discussions about the mass use of protected data to train large language models and chatbots like ChatGPT, it has become clear time and again how important such approaches are for preventing personal data from being misused or copyrights from being infringed upon.

The 13 Sustainability Criteria for AI Systems

-  **Transparency and Accountability**
-  **Non-Discrimination and Fairness**
-  **Technical Reliability and Human Supervision**
-  **Self-Determination and Data Protection**
-  **Inclusive and Participatory Design**
-  **Cultural Sensitivity**
-  **Market Diversity and Exploitation of Innovation Potential**
-  **Distribution Effect in Target Markets**
-  **Working Conditions and Jobs**
-  **Energy Consumption**
-  **CO₂ and Greenhouse Gas Emissions**
-  **Sustainability Potential in Application**
-  **Indirect Resource Consumption**

The Questionnaire

The questionnaire for the Self-Assessment Tool was developed based on this set of criteria, indicators and deployments. Not all of the questions included in it are relevant to all organizations that are developing or using AI. If, for example, an organization produces AI systems that do not make direct decisions about people but are used only in industrial processes,

Selbstbewertungstool zur Nachhaltigkeit von KI

Der Fragebogen zur Selbstbewertung gibt Ihnen Orientierung, wie nachhaltig Ihre KI-Systeme sind und wo Sie nachbessern können.

Das Bewertungstool ist für Organisationen ausgelegt, die KI selbst entwickeln oder in ihrer Organisation einsetzen. Es basiert auf den im Projekt „SustAln“ entwickelten Kriterien zur sozialen, ökologischen und ökonomischen Nachhaltigkeit von KI. Nach Beantworten des Fragebogens erhalten Sie ein Auswertungsdokument. Ihre Antworten werden nicht bei uns gespeichert.

Das Tool zur Bewertung der Nachhaltigkeit von KI

Schritt 1 von 11

Organisation

Was trifft auf Sie und Ihre Organisation zu? (erforderlich)

- Wir entwickeln intern KI-Systeme für den Einsatz in unserer Organisation
- Wir setzen KI in unserer Organisation ein, die wir extern entwickeln lassen oder einkaufen
- Wir sind eine Organisation, die KI als Dienstleister für Kund*innen entwickelt

Transparenz und Verantwortungsübernahme

Liegt in Ihrer Organisation ein Code of Conduct (wie z. B. ein AI Code of Ethics) vor, in dem grundlegende Werte und Normen beschrieben werden, die bei der Implementierung und Nutzung von KI-Systemen beachtet werden sollen? (erforderlich)

- Ja
- Nein

Gibt es in Ihrer Organisation Ansprechpartner*innen oder Organisationseinheiten, die für ethische Belange in Bezug auf die Entwicklung und den Einsatz von Künstlicher Intelligenz in Ihrer Organisation verantwortlich sind und an die sich Entwickler*innen, Anwender*innen, Endnutzer*innen, Auditor*innen etc. ggf. wenden können? (erforderlich)

- Ja
- Nein

Stellen Sie detaillierte Informationen zu einzelnen, mehreren oder allen KI-Systemen bereit? (erforderlich)

- Nein, zu keinen Systemen
- Ja, zu einzelnen Systemen
- Ja, zu den meisten Systemen
- Ja, zu allen Systemen

Zu welchen Aspekten stellen Sie in Ihrer Organisation in der Regel Informationen bereit?

- Zur Zielsetzung des Systems
- Zum Anwendungsbereich
- Zu den Nutzer*innen
- Zu den verwendeten Daten
- Zur Art des Modells
- Zum Reflexgrad des Modells
- Zur Anzahl und Auswahl der Parameter und Features
- Zu den genutzten Inputs
- Zu den Testverfahren
- Zu den Metriken
- Zu Risiko-Analysen
- Sonstige Informationen

Werden Komplexität oder Verständlichkeit zur Entscheidungsgrundlage in Entwicklungs- oder Einkaufsprozess, indem z. B. weniger komplexe Systeme bevorzugt werden?

- Ja
- Nein

Welche Methoden zur Erhöhung der Transparenz und Erklärbarkeit werden in der Regel eingesetzt?

- Einsatz von transparenten Machine Learning Modellen (z.B. Linears/Logistische Regression, Entscheidungsbaum, Rule-based learning)
- Vereinfachung der Modelle
- Erhebung der Feature Importance
- Visualisierung
- Sonstige Methoden

Speichern und später fortfahren Zurück Weiter

questions about their fairness become unnecessary. Organizations must answer between 48 and a maximum of 66 questions in total to obtain an assessment of the sustainability of their AI systems. They must indicate, for example, whether they follow a Code of Conduct that lists fundamental values and standards in the implementation and use of AI systems, or whether sustainability certifications are available for the data centers and hardware they use. Since the questionnaire is based on a broad understanding of sustainability, it covers a wide variety of areas. Presumably, the questions will have to be answered by multiple departments in the organization.

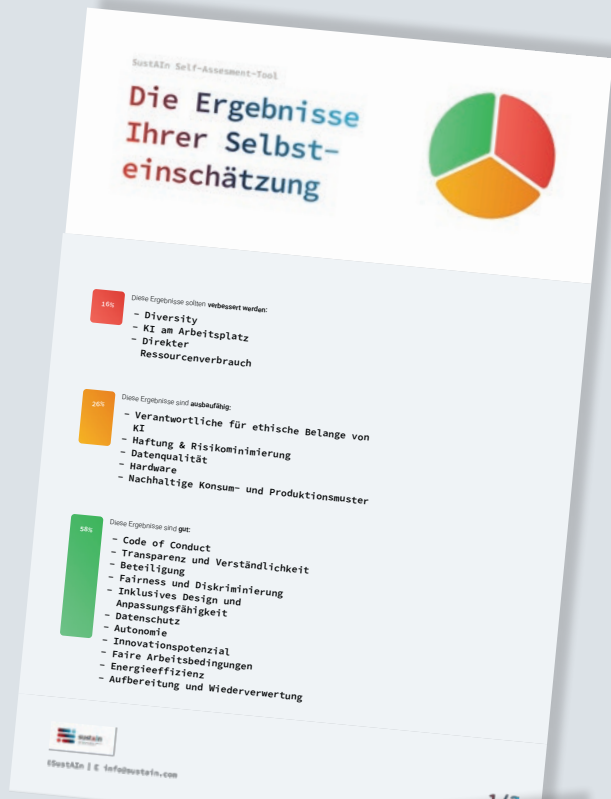
The Appraisal

The results of the Self-Assessment Tool are made available to organizations as a PDF download at the end of the questionnaire. They are not visible to third parties and are not stored. Our graphics are intended to provide guidance to organizations on which areas of their AI system are faring well in terms of sustainability and where there is room for improvement. The questionnaire is evaluated based on a point system. Due to the complexity and multilayered nature of the topic, we have decided against calculating a sustainability score from the responses. Instead, we provide organizations with a granular, differentiated assessment that does not aim to provide AI products with a “green” label, instead seeking to make the organizations that develop or use such systems aware of their sustainability.

The Recommendations

Our online Self-Assessment Tool also provides organizations with concrete recommendations for action. There is no ready-made recipe for the sustainability of AI, and decisions on how to make them more sustainable must be made on a case-by-case basis. Beyond simple technical measures, numerous micro decisions must also be made in the planning and development process. The most important factor is a cultural shift at the organizational

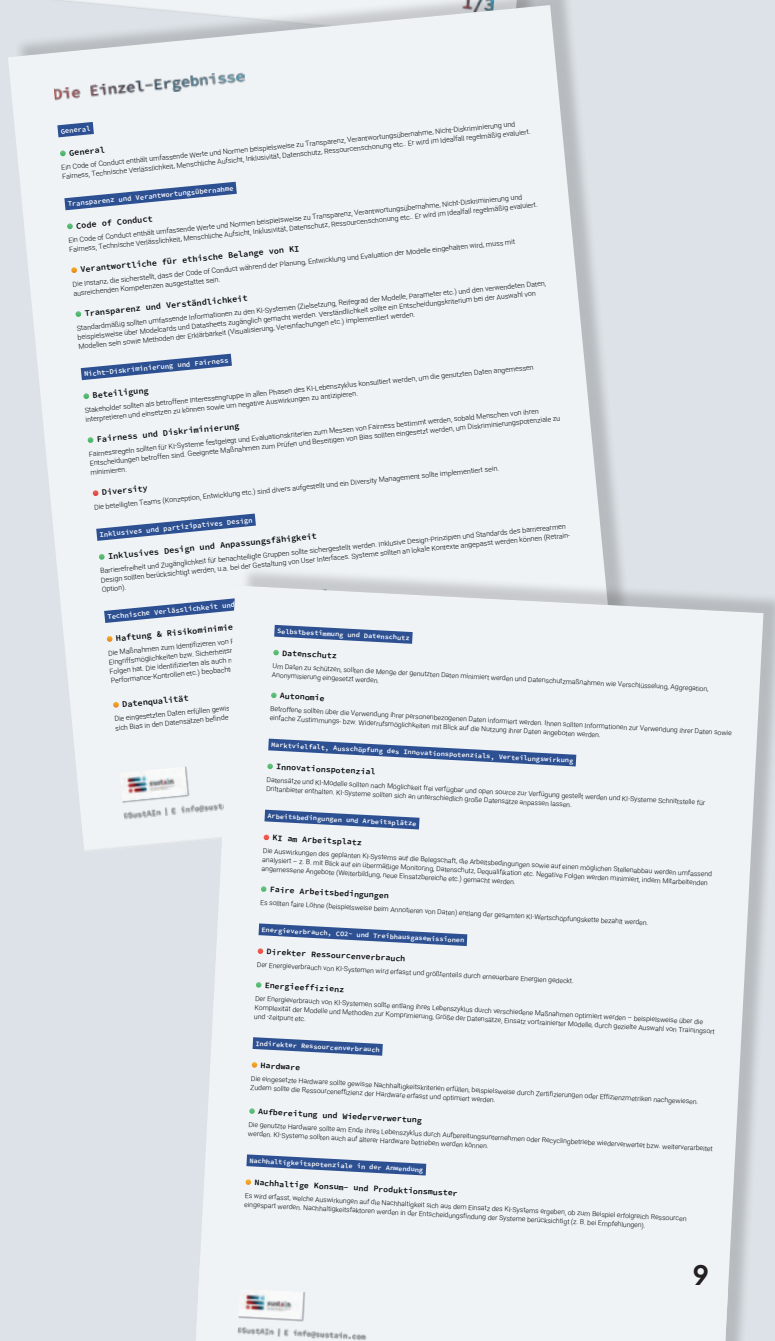
The results of the Self-Assessment Tool are made available to organizations as a PDF download at the end of the questionnaire. Our graphics are intended to provide guidance to organizations on which areas of their AI system are faring well in terms of sustainability and where there is room for improvement.



level, through which sustainability becomes a guiding principle for all decisions that are made. Such a shift must be planned, moderated, executed, evaluated and technically implemented. Our tool is designed to pave the way for that to happen.

What Comes Next?

Through the AI Act, the EU is pushing for sustainable AI systems. But even apart from such regulatory approaches, all organizations should address the sustainability of their AI systems. Indeed, we need to act now before unsustainable AI ecosystems have become established to the point that they are difficult to modify. In expectation of further regulatory action, various industry standards are already being renegotiated against the backdrop of demands for AI system sustainability. We hope to bolster that development with the sustainability approach conceived in the SustAln project.



Keeping an Eye on AI

We can't leave it to the technology providers to stop AI's worst aberrations - we must also be proactive.

The Rebound Effect

When a specific technology is developed to boost efficiency, yet ultimately leads to increased energy consumption, the phenomenon is referred to as the Rebound Effect, sometimes also called the Jevons Paradox or, more prosaically, the boomerang effect. Significant efficiency improvements can produce the opposite of the desired effect if our use patterns change as a consequence. When consumers become convinced, for example, that certain products or services are particularly efficient, they might feel better about using them - which then leads to greater use. This psychological effect can be seen, for example, with more fuel-efficient automobiles, which are then driven more frequently than their inefficient predecessors. A similar phenomenon appears when the use of technological devices becomes more attractive through efficiency improvements. The processors in our laptops and smartphones have, for example, become far more efficient over the years, but partly as a result, we use more devices more frequently and for longer periods of time than ever before - in


addition to loading them with a greater number of programs and apps. The result is that they (we) are consuming more total energy.

However, the Rebound Effect doesn't just apply to end users. It can be seen in all areas where growing demand cancels out efficiency gains. Take, for example, a company that successfully integrates AI to optimize its transportation systems and sink costs. That increased efficiency may motivate the company to actually increase its overall transport volume to take advantage of the savings.

There are also situations in which putative energy savings generated through the implementation of AI systems are actually illusory - for example, when they are counteracted by higher energy costs for the server infrastructure. Similarly, longer battery life for tablets may be the result of shifting much of the energy consumption to cloud services and the data centers that power them. Energy-intensive AI models can even greatly increase the energy and resource consumption of certain apps without end users even realizing it.

The use of AI applications can also produce what's called the "Spare Time Rebound Effect." If, for example, a fully automated robotic vacuum mop takes over the cleaning, people have more time to do other things. But if they then fill that time with energy-intensive activities, it can lead to greater total energy consumption.

Does that mean that technological advancements in AI energy efficiency are necessarily destined to be cancelled out by the Rebound Effect? Not necessarily. Rebound Effects show that we need political and regulatory tools to eliminate them when they appear. Only then will we be able to sustainably reduce energy consumption.





C

Cultural Hegemony: How Generative AI Systems Reinforce Existing Power Structures

In their development and use, AI systems are always embedded in a specific social context. That also applies to large text- or image-generating models, the output of which suggests that we are being presented with facts or realistic images. Ultimately, though, that output is little more than realistic-appearing content that disseminates specific cultural codes.

Large AI text-to-image generators like Stable Diffusion, DALL-E 2 or Midjourney are trained with vast quantities of data. They analyze frequently appearing patterns, such as the typical proportions of a face, or what landscape pictures tend to look like. When they then produce their own images of faces or landscapes, they may reflect biases that were present in the training data (such as distortions of human facial features stemming from racist caricatures)


or false representations (like including typically Western architecture in an image of a city from a completely different part of the world with a radically different skyline).

Biased training data, though, isn't the only problem. Many visual models create far less realistic images of Black women than they do of white women. They more often contain distortions and outright errors, as the [artists Stephanie Dinkins and Minne Atairu](#) discovered. Some image generator providers have reacted to such potentially harmful outputs (due to the latent racism they reveal) by blocking specific prompts (the requests made to the AI system to generate a specific image). Artist Auriea Harvey discovered, for example, that some image generator tools block prompts like "slave" or "slave ship." But such blocks do more to conceal the problem than to actually solve it. Indeed, such blocks amount to essentially stifling historical realities, which can also magnify cultural dominance by suppressing the perspectives and experiences of minorities.

Such cultural dominance doesn't necessarily have to manifest itself through

discrimination. Western norms are often subtly forced onto the users of image generators – for example in the way people [smile in AI-generated pictures](#). Even such a primeval form of expression like the smile may be perceived and reacted to divergently by people from different cultures.

The same risk – that of propagating a hegemonic monoculture – is also present with text generators like ChatGPT. Different languages describe human experiences in their own way. But the [cornucopia of smaller languages is at risk](#) of being lost in the algorithmic hegemony because of the vast amount of data from books, magazines, newspapers and online content necessary for training – a volume that smaller languages are simply unable to provide, particularly those languages that are only spoken. It's no secret that English is the dominant language of technology and that many smaller languages are [falling to the wayside](#) in AI applications. The team behind Stable Diffusion even notes in its own [model card](#) that the vast majority of the training data is in English, and that entries in other languages don't work (as) well.



R

Regulating Data Centers

In 2020, data centers in Germany consumed a total of 16 billion kilowatt hours of electricity. That is the equivalent of around 3 percent of Germany's total energy consumption. The waste heat generated by the data centers still goes

largely unused, despite the fact that it could contribute significantly to reducing CO₂ emissions in the heat supply.

To improve this less-than-ideal situation, the German government intends to introduce a law that will require more efficient use of energy in data centers. But relative to the promises made by the three parties in Germany's current coal-

ition government, the plan has been significantly watered down. By 2030, electricity consumption for cooling, energy distribution and energy storage in data centers must account for no more than 30 percent of the electricity required for the actual computing power. This target, however, is usually achieved in large data centers anyway. In addition, the draft law does not provide for sanctions if the requirements are not met.

Only the very largest data centers are affected by the law (probably fewer than 1 percent of all those in Germany), and the requirements for waste heat utilization (such as mandatory feed-in to heat grids) are lax. As proposed, the current draft legislation fails to take into account the approximately 40,000 smaller German data centers. Yet they, too, produce waste heat that could be used for heating. The federal government is also failing to promote needed resource-conserving designs for data centers that go beyond mere energy consumption.

The claim can often be heard that setting overly ambitious efficiency requirements creates a competitive disadvantage. But that seems unlikely. Other countries, like France, have also enacted a legal obligation to publish the environmental impact of digital services and data centers. The French law aims to raise awareness of the environmental impact of digital technologies by requiring regular disclosure of key environmental impact indicators throughout the life cycle, such as CO₂ emissions or energy and water consumption.

Since data center infrastructures cannot be changed overnight, our society will have to live with unsustainable infrastructures for decades to come. As such, we need legislative proposals that – in contrast to the one put forward by the German government – really put digitization in Germany on a more sustainable path.

SustAIIn for Reading and Listening

From 2020 to 2023, the Environment Ministry-funded AI flagship project “SustAIIn: The Sustainability Index for Artificial Intelligence” established a baseline for the ongoing discussion of AI sustainability. The project’s most important findings can be found on the [project homepage](#). Additional key outputs can be found here:



1) A discussion paper on sustainability criteria for Artificial Intelligence

“[Sustainability Criteria for Artificial Intelligence: Developing a Set of Criteria and Indicators to Evaluate the Sustainability of AI Systems Throughout Their Life Cycles](#)” (German only)

By Friederike Rohde, Josephin Wagner, Philipp Reinhard, Ulrich Petschow, Andreas Meyer, Marcus Voss and Anne Mollen

2) A video recording of the digital policy discussion at the *Bits und Bäume* 2022 Conference

“[Political Responses to the Sustainability Costs of AI](#)” (German only)

With Alexandra Geese, Tabea Rössner, Marina Köhn

3) An English-language publication looking at sustainability criteria for Artificial Intelligence

“[Broadening the Perspective for Sustainable AI: Comprehensive Sustainability Criteria and Indicators for AI Systems](#)”

By Friederike Rohde, Josephin Wagner, Andreas Meyer, Philipp Reinhard, Marcus Voss und Ulrich Petschow

4) A short study on sustainable approaches to AI systems




“[Sustainable AI and Digital Self-Determination: Requirements for a Sovereign and Sustainable Approach to AI in Our Daily Lives](#)” (German only)

By Anne Mollen

5) A video recording of the digital policy discussion at the launch of the first issue of SustAIIn Magazine

“[Political Support for Sustainable AI?](#)” (German only)

With Lynn Kaack, Sergey Lagodinsky, Pascal König, Marcel Dickow



Personalized Online Marketing: AI Technology Gone Astray

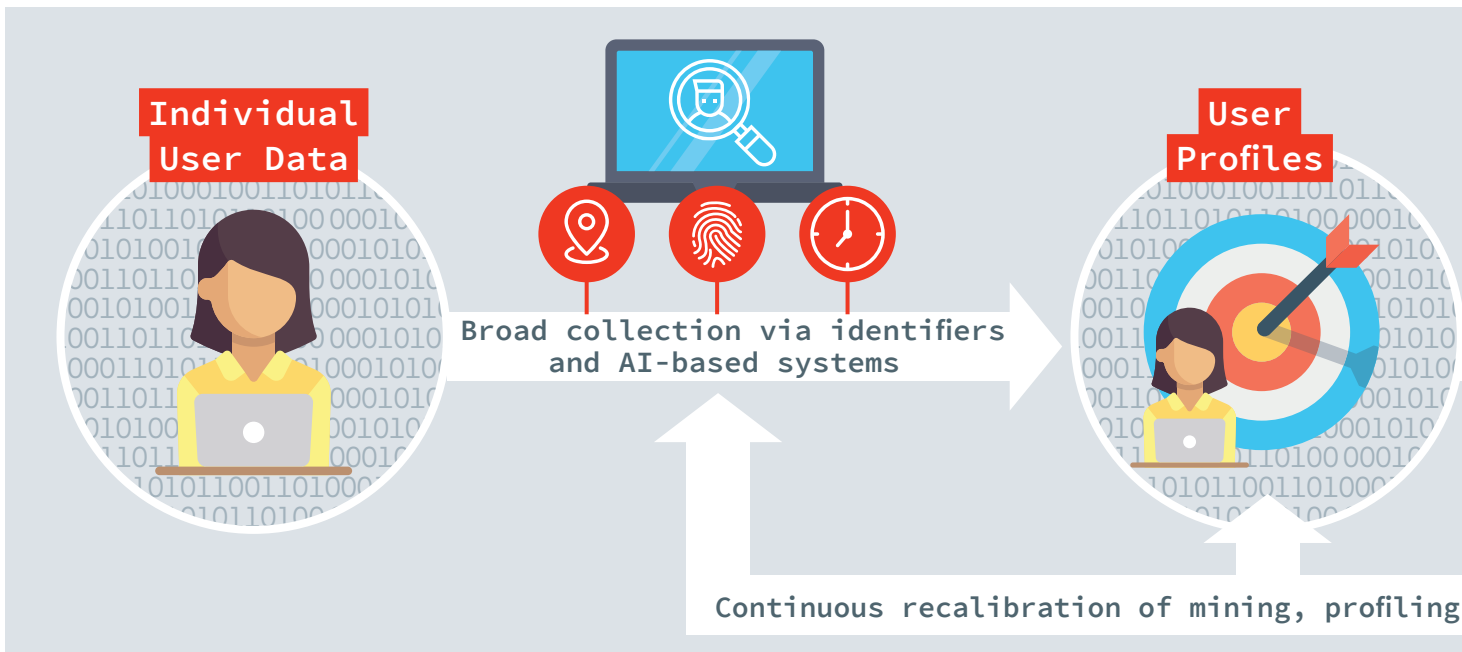
Nine out of 10 internet users worldwide use the Google search engine. The social network Facebook has 3 billion users, Instagram has 2 billion. In Germany, more than 70 percent of the population have a Facebook account, and over 60 percent use Instagram. But if using these platforms is free of charge, how have their operators – Meta and Alphabet – grown to be two of the largest companies in the world, making billions of dollars in profit every year?

The short answer: Their business models are based on generating revenue from advertising and using personal data to optimize the targeting of ads. This enables them to provide “free” content or services to users. Google, for example, generates 78 percent of its revenue through advertising.


Yet even though users don’t have to pay money to use such platforms, ad-funded content and services come at a

price: Users pay with their personal data and by exposing themselves to advertising. We encounter advertising primarily as paid hits in search results or as ads on websites and social media channels – for example in the form of the banners or videos, which are omnipresent on most of our screens. Their tremendous value does not stem solely from their presence online. Rather, advertisers spend substantial amounts of money on ads that are published exactly where the

CIRCULATION OF DATA AND ADS



PROCESS STEPS & ACTOR ROLES



generation of clicks, views and sales is maximized – this is called personalized advertising. Digital targeting mechanisms have inflated online advertising markets. In Germany, the size of the market has doubled over the last six years. Furthermore, online marketing is making significant advances through Artificial Intelligence (AI) technologies. Marketing has become one of the most important fields for AI applications.

How Does Personalized Online Advertising Work?

The journey of a personalized ad can be quite complex: Starting from an advertiser, it passes through several processes managed by intermediary actors in the ad tech sector before it finally appears in a publisher’s advertising space on the screen of a targeted online user. The basis for the whole personalization cycle is personal data, which is most

commonly collected, analyzed and managed by the intermediaries.

All steps are based increasingly on machine learning techniques (ML), which form the most common foundation for AI systems. To increase the likelihood that ads will lead to purchases, marketing strategies are usually designed to target specific online users as precisely as possible. Intermediaries use data mining to capture personal data. This data can include: online (and, increasingly, offline) behavior; geolocation data and movement trajectories; device or user identifiers (mobile IDs or advertising IDs, for example) or information from user profiles, which may include sensitive demographic data such as age, gender, ethnicity, sexual orientation, political or religious beliefs, education level, employment status or income. Even if this data is not directly available, ML models can infer such information

with surprisingly high accuracy. Cookies – small data files containing browsing histories (website visits) and background connections – are widely used to feed AI systems.

Intermediaries then use this personal data to create user profiles with the help of digital identifiers. Using a method of online tracking known as fingerprinting, they capture identifiers by combining user attributes with data generally provided in network requests (such as IP address, web browser, operating system and hardware specifications) and compress them into a single “digital fingerprint.” Further sources of information include payment histories or cross-platform verifications using, for example, a unique phone number. It is also possible to obtain relevant personal data from third parties. By segmenting such digital fingerprints into various sorts of groups, advertising can now target individually configured profiles of users who are most likely to follow up with a purchase. The profiles are fed into ML models, which are trained and fine-tuned to “predict” the success of the particular ad as accurately as possible. These predictions are then used in real-time micro auctions that sell advertising space to the highest bidder. These real-time auctions take place at the moment a user accesses a website, opens an email or launches a mobile app – invisibly and in the blink of an eye.

In practice, advertisers are not just companies advertising their products but media agencies that plan, implement and evaluate advertising campaigns for companies (Figure 1). Intermediaries link demand-side platforms (operated by advertisers) with supply-side platforms (operated by publishers) to organize the advertising spaces that ultimately deliver ads to the end user. Those spaces can be on feeds, apps, search engines, websites, or other interfaces, and they are constantly monitored to maximize ad performance. The success of spe-

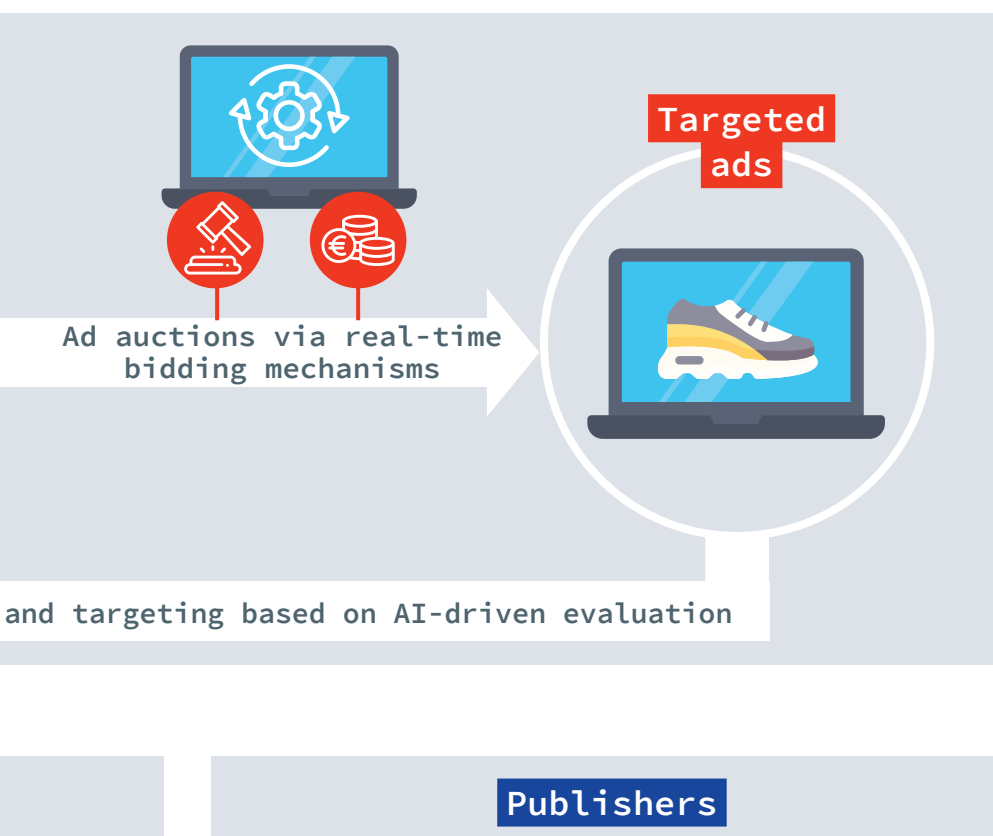
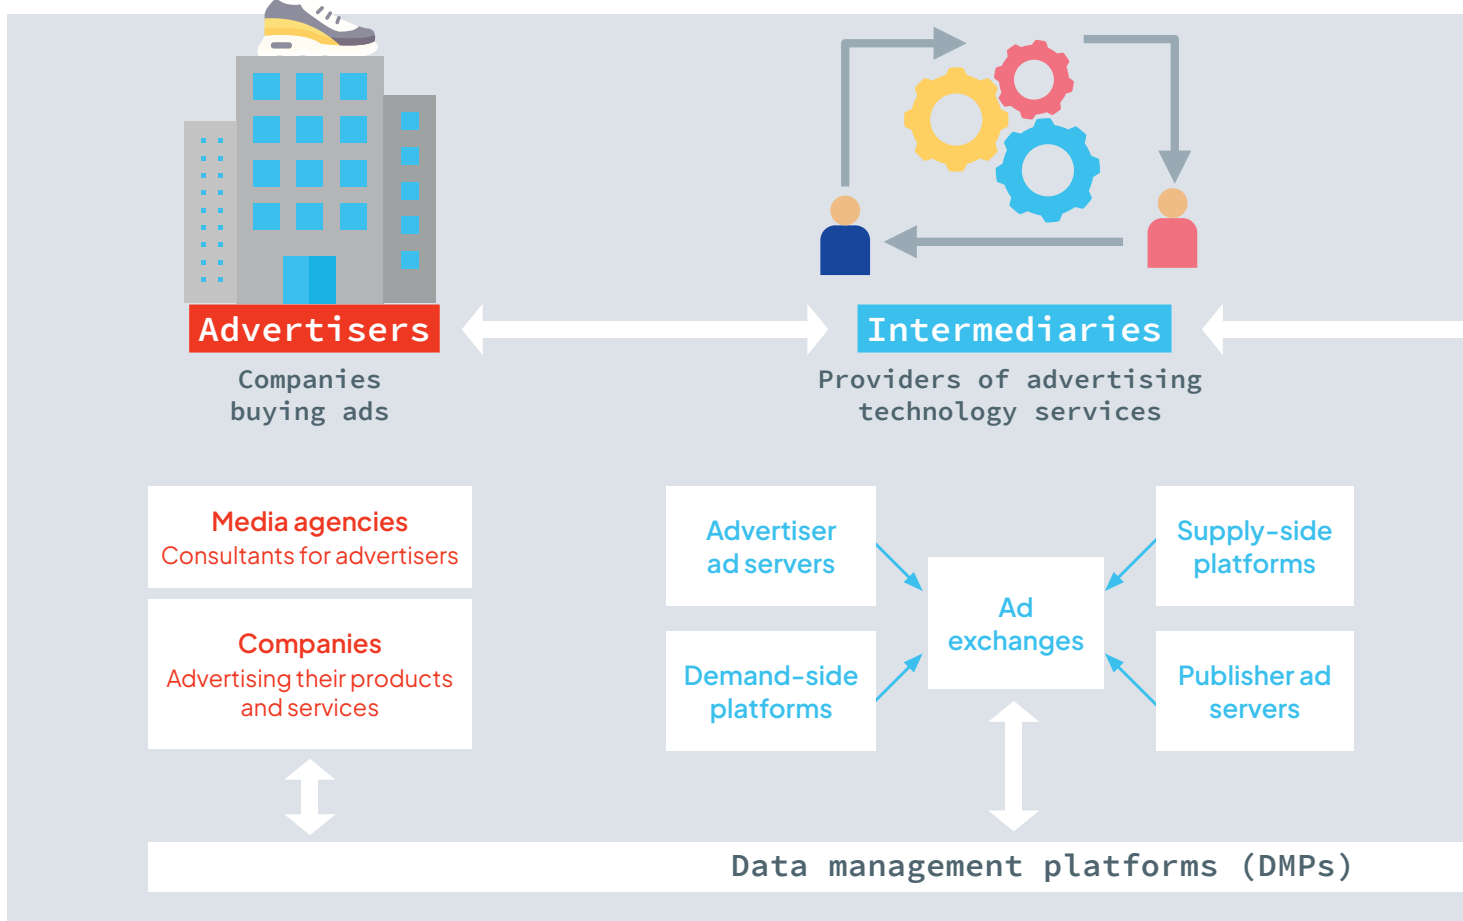


Figure 1

ACTOR MAPPING



Actors:

Advertisers want to increase the visibility of products and increase revenues through online advertising. They develop advertisements and purchase advertising space.

Intermediaries are interposed between advertisers and publishers, forming a complex network of actors that offer technology services.

Publishers provide the online space that is considered to be most suitable (profitable) for a specific ad to reach a specific user. Publishers might be search engines, social media sites or media and video platforms. They provide ad space in the form of banners, video clips or search results. Meta and Alphabet act as publishers, for example, as do providers of various apps and website operators such as online news outlets.

Users interact with their devices, operating systems, application software or online services. Through these interactions, they see advertising, but they also directly or indirectly provide personal data, which fuels the further personalization of the ads they see.

cific ads can be adjusted by varying ad content, the frequency with which they are served, their size and shape, and other factors. Buying and selling ad space is taken care of on ad exchange platforms. Data can be stored, enriched, analyzed and segmented on data management platforms. They are also used to create usage profiles.

Finally, media agencies are tasked by advertisers with measuring and predicting the success of placed ads. They use various criteria to do so: Engagement rates indicate how many people click or swipe a specific ad; impressions represent the number of times an ad is served; viewability shows whether an ad was actually seen; reach describes the number of unique users who view the ad; frequency indicates how many different times a user sees a specific

ad; and conversion describes the extent to which an ad leads people to take a particular action (such as buying a product). These analyses incorporate ML not only to target users, but also to evaluate performance and adjust future ad campaign strategies. AI thus plays a significant role along the entire life cycle of personalized ads, supporting data acquisition, data analysis, user targeting and the continuous adaptation of all these processes.

The Dangers of Market Concentration

The online advertising market is dominated by two companies: Alphabet, better known for its subsidiaries Google and YouTube, and Meta, which owns Facebook and Instagram. Not only are they the most important ad publishers, they also

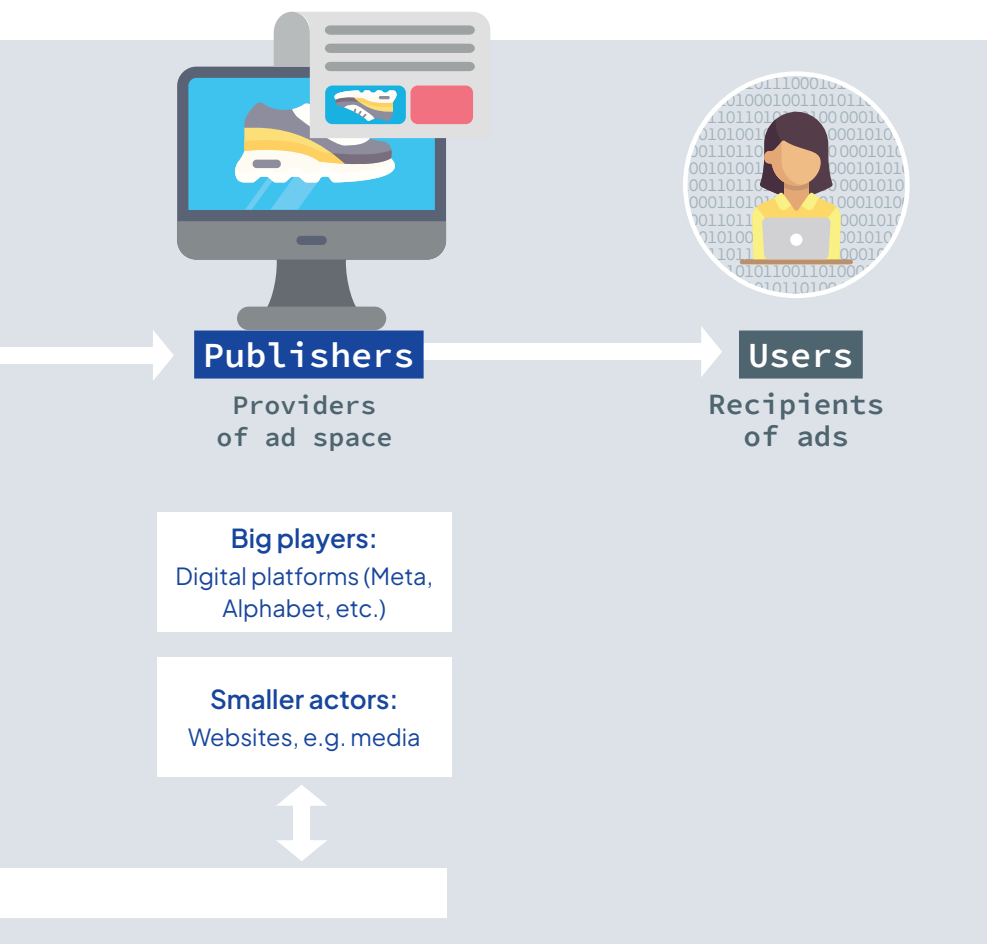


Figure 2: based on Armitage et al., 2022

employ some of the biggest intermediaries, which broker advertising space and use personal data to improve ad targeting. This places the three main activities of the industry – advertising, mediating and publishing – within one and the same company, meaning that just a few players control the entire journey of an ad.

Alphabet and Meta provide the technical infrastructure in addition to much of the advertising space, and they are also in possession of vast amounts of user data. The collection of this data both within and outside of their platforms gives them a significant competitive advantage, making it almost impossible for other companies to compete. This domination of the online advertising market is frequently referred to as a “duopoly” – the monopoly of two. Many of their services (search engines, social

networks, navigation and office apps, cloud storage, translation tools, entertainment platforms, development tools, news feeds, etc.) have become indispensable. With a lack of alternatives that provide a similar level of utility, users often have no choice but to opt for their services. This concentration is not new to the digital economy, where the five companies known as GAFAM (Google, Amazon, Facebook, Apple and Microsoft) are dominant. All of them were among the eight companies with the highest market value worldwide in 2022. In the advertising sector, this centralization of power is particularly extreme as the value and success of advertising is heavily dependent on the volume and accuracy of data.

Due to the lack of competition in the marketing ecosystem, Alphabet and

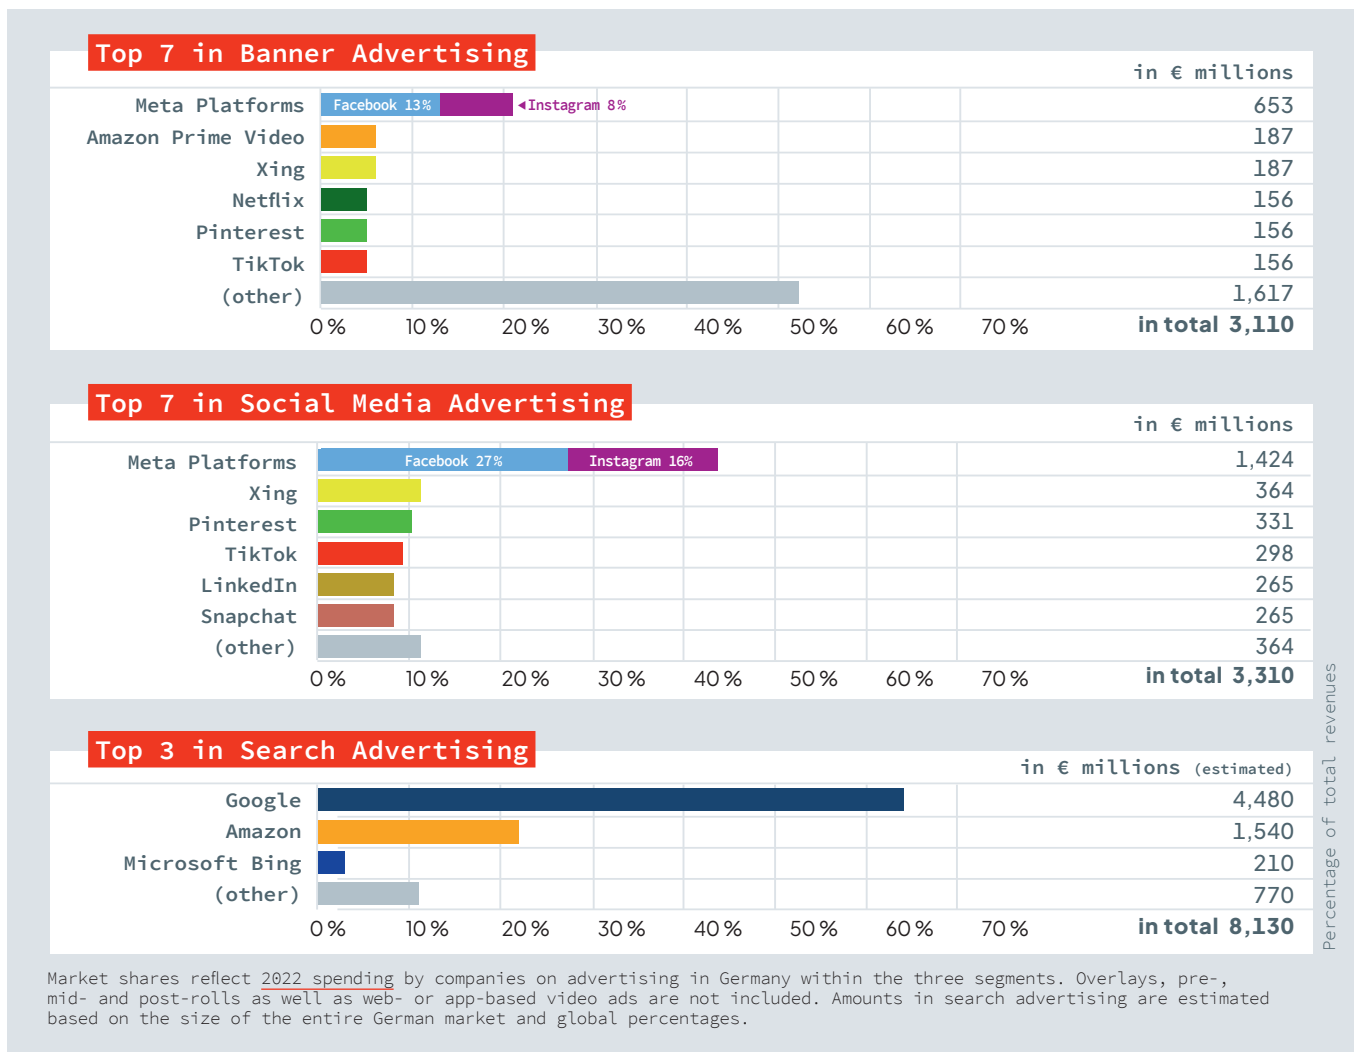
Meta are far more influential – when it comes to pricing, practices and technical standards. This makes it difficult for advertisers, who are dependent on the technological infrastructure offered by, for example, Google when it comes to cloud systems and AI algorithms. Users, meanwhile, hardly have a choice regarding making their personal data available and who is involved in data collection when visiting a website – an information asymmetry that is intentionally reinforced by website design. And governments have a hard time regulating the activities of corporations and keeping pace with technological developments. The upshot is that the current advertising market caters largely to the needs of a tiny handful of companies, while other actors are unable to compete.

People and the Planet Are Paying the Price

The current situation poses a threat to individuals and society. Not only does the extensive collection and analysis of personal data undermine privacy and data protection, resulting in a surveillance economy. Power imbalances and information asymmetries between companies, individuals and countries are also being used to influence political processes. The Cambridge Analytica scandal has clearly demonstrated the extent to which micro-targeting and the dissemination of misinformation endanger the formation of public opinion and thus democracy as a whole.

On top of this, the marketing activities described above consume huge amounts of resources. About 15 percent of the network activity triggered by the loading of a news website comes from ad-related content. Ad exchange servers are running continuously to manage advertising all across the commercial web. The servers consume energy, produce carbon dioxide emissions and trigger additional purchases

RISKS OF MONOPOLIES: SELECTED ONLINE ADVERTISING MARKET SHARES IN GERMANY



of resource-intensive consumer goods and services.

To make the online marketing industry sustainable, it might not be enough to simply open the market to new private players. Genuine competition alone would not solve the problem. Rather, the infrastructure of online business and communication networks must move away from the principle of personal advertising. Solutions may lie in the regulation of tracking and the mass gathering of data. We need to look at alternatives so that public data infrastructures are no longer based on the exploitation of personal data.

Full Study:

Frick, V., Marken, G., Schmelzle, F. & Meyer, A. (2023). "The (Un-)Sustainability of Artificial Intelligence in Online Marketing. A Case Study on the Environmental, Social and Economic Impacts of Personalised Advertising." IÖW Schriftenreihe 228/2023. ISBN 978-3-940920-33-1.

Further reading:

Armitage, et al. (2023) "Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers" (EU report).



**VIVIAN FRICK,
GESA MARKEN
AND FRIEDER SCHMELZLE**

... are Research Associates at the *Institute for Ecological Economy Research (IÖW)* in Berlin. They examine the role of digitalization in the social-ecological transformation.

How Does Personalized Online Marketing Affect Energy Consumption?

An analysis

With the internet and smartphones having become ubiquitous, technology companies now have a greater amount of detailed customer information at their fingertips than ever before. They know how users surf the web, what products they are likely to buy, how they behave on social media platforms and even their location. This information enables advertisers to tailor their ads to extremely precise target groups. And that has revolutionized online advertising. Instead of developing broad advertising strategies aimed at reaching as many people as possible, companies have begun personalizing their ads and placing them only where they will have the greatest effect. To achieve the most accurate targeting possible, AI systems are used to analyze huge amounts of user data and produce detailed user profiles. Based on those profiles, customers are divided up into target groups, and for each of those groups, demand forecasts are generated, which will determine the advertising they are shown.

This practice has generated significant debate, with most of the discussion focusing on data protection and ethical concerns. The potential ecological risks, however, are less frequently addressed. According to estimates, internet use is responsible for the consumption of over 400 terawatt hours (TWh) of electricity

each year. Current trends suggest that this total will continue to grow strongly in the coming years. Little research has been done into how the personalization of advertising using data analysis methods has impacted energy consumption. It seems fair to assume, however, that it has amplified existing trends.

The user data fed into such analyses comes from a number of different sources, including websites, social media platforms and mobile applications. Network infrastructure and data centers are required for the transfer of the data, all of which consume energy. And the data collected must be stored and managed for extended periods in data centers and on hard drives – and here too, energy is consumed. The next step involves using the data to train the machine learning models that are the backbone of personalized advertising campaigns, a process which requires (energy intensive) high-performance computer infrastructure with networked servers. And finally, servers must also be cooled during operation, which accounts for a significant share of a data center's energy consumption.

But that's not all: The determination as to how and where the advertisements are ultimately placed is made in a process called "real-time bidding." Advertis-

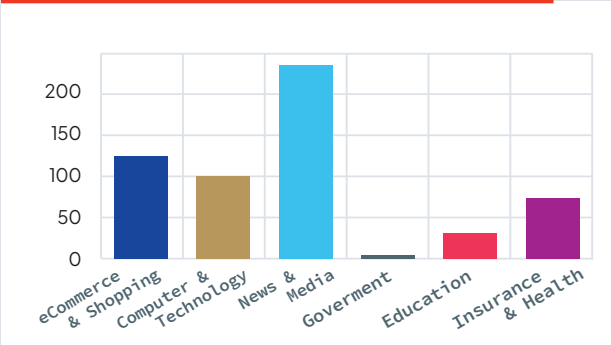
ers bid in real-time auctions for ad real estate on websites that seem relevant to the target groups they are trying to reach. These auctions require extremely quick data processing so that advertisers, on the one hand, and those with advertising real estate, on the other, communicate with each other. The rendering of content, such as the images and video that generally accompany advertisements when they are presented on end devices, also requires energy.

Simulation Studies


Simulations generate data about the amount of energy consumed by the deployment of personalized online advertising. But such studies tend to focus primarily on the energy consumed by the end device to which advertising data is transmitted. It is difficult to estimate, by contrast, the amount of electricity consumed by the AI systems and data analysis procedures deployed in the personalization process or by the storage of the necessary data. Advertising companies are not always willing to share such information.

For our study, we used the data-protection tool OpenWPM to collect data about website visits, and we examined around 200 of the most-visited German internet domains. We installed an

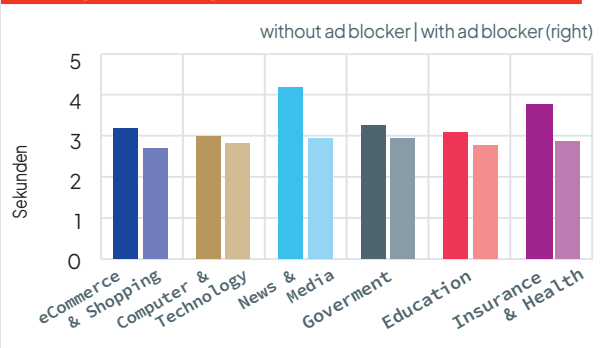
Average number of third-party cookies*



Average amount of data transferred by third parties*



Average loading times per website visit*



*by website category

automated crawler on a laptop and had it repeatedly open websites while we collected data on the amount of energy used by the computer's CPU and information about the cookies transmitted. We conducted the simulation both with and without an ad blocker to determine the effect that ads have on end device energy consumption.

The time it took to open the websites was 14 percent lower when the ads were blocked, which translated to 10 percent less energy being consumed by the end device's CPU. Depending on the complexity of the website visited, the end

of cookie-related data transferred by 75 percent. The greatest savings were achieved with websites that fell into the "news and media" category, which are often financed through the serving of ads.

The amount of energy consumed by a single website visit and the related transfer of cookies is, to be sure, quite small. But just the 200 German websites analyzed for this simulation alone are visited 4.5 billion times per month. In our experiment, we only visited the homepages of the sites we included in the study, but generally, visitors also

It is estimated that the internet is responsible for 15 percent of total global electricity consumption – and rising quickly. Online advertising and the collection of user data are both big contributors.

device specifications and the browser used, opening a website generally resulted in energy consumption of 0.01 to 1 watt hours (Wh). The rendering of an advertisement by the graphic card produced average energy consumption of 0.005 Wh. An average of 155 cookies, with a mean size of 138 bytes, were transferred during each website visit, 87 percent of which came from third parties. During each website visit, an average of 0.2 MB worth of data was transmitted in the form of cookies. The rejection of non-essential cookies reduced the amount

navigate to subsites as well. Ads are frequently served from those pages too, which translates to the additional transmission of data.

Our results show that the risks of personalized online marketing are not limited to privacy protection issues. It is estimated that the internet is responsible for 15 percent of total global electricity consumption – and rising quickly. Online advertising and the collection of user data are both big contributors. The flow and processing of data that is produced by online advertising can be limited through the technology deployed ("Privacy by Design") and the settings chosen ("Privacy by Default"). Companies that use AI-based data analysis procedures to analyze user data for the purposes of personalized online marketing should face stricter requirements to take steps to reduce energy consumption and make available data pertaining to the energy efficiency of their systems so as to gain a better understanding of their ecological impact.



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.

Artificial Intelligence and Marketing: A Concerning Combination?

What Companies and Civil Society Think

The advertising industry has high hopes for the use of AI in online marketing. It also promises benefits for its customers. However, advertisers are primarily interested in boosting sales. As long as there are no rules in place for regulating the way large companies use AI for tracking user behavior and for producing personalized, custom-designed ads, AI will remain a sustainability risk. The environmental and social burdens are rising, while smaller companies and end users are struggling to find a promising way forward.

AI-supported personalized online advertising is expanding rapidly. Companies expect the technology to allow them to serve their ads precisely to those people who are most likely to buy their products. Such targeted advertising is vastly superior to running ads in printed media, putting up billboards or buying television spots, none of which are as effective in targeting potential buyers. Business consultants forecast that the efficient deployment of AI in marketing will produce significantly higher returns for companies. And such personalization, many say, is also in the consumers' interests since it means they will primarily be shown ads for products that actually interest them. It has even been argued that personalized marketing is more environmentally sustainable than traditional advertising since it allegedly cuts back on energy and resource waste by reducing misguided purchases and the returns thus necessitated.

But what do companies and civil society organizations think about these promises? We conducted a number of interviews with experts from companies that rely on ads, from marketing agencies and from civil society organizations that focus on internet policy. Beyond that, we also carried out two representative online surveys in summer 2023 in which we asked 2,000 people from private households in Germany and decision-makers from around 500 companies about the opportunities and risks they see in the use of AI-supported personalized advertising.

Very Few Find Personalized Advertising Helpful

The advantages frequently cited for the use of AI in marketing are based on the assumption that all people online are potential consumers. Companies plac-

ing ads are interested in transforming potential consumers into real customers – a process known in marketing as “conversion” – by making them aware of their products and convincing them to make a purchase. The companies placing the ads are interested in motivating their potential customers to become real customers – by making those potential customers aware of their products and convincing them to make a purchase. From the perspective of these companies, whether users consent to the use of their personal data or agree to receive ads is of secondary concern. They can, of course, use adblockers and adjust their cookie settings to protect themselves. But they are often unable to decide for themselves the degree of access to their personalized data they are willing to allow.

Our survey showed that only between 13 and 22 percent of users find person-

HOW DO YOU VIEW PERSONALIZED ONLINE ADVERTISING?



Figure 1

Around half of those surveyed take measures to protect themselves from the collection of personalized data. Some 50 percent of respondents use ad blockers and avoid (at least to a certain extent) digital services that intensively collect data.

alized advertising to be “interesting,” “helpful,” “welcome” and “trustworthy.” By contrast, around 45 percent find it to be “uninteresting,” “annoying,” “unwanted” and “manipulative” (Figure 1). In short, only a small minority of users have a positive view of personalized advertising.

Around half of those surveyed take measures to protect themselves from the collection of personalized data (Fig-

ure 2). Some 50 percent of respondents use ad blockers and avoid (at least to a certain extent) digital services that intensively collect data. Less than half of respondents use specialized data-saving search engines and browsers. The majority of those survey participants who do not use such data-saving options said they weren't aware they existed, don't know how to use them, find them too complicated or doubt they are beneficial. In sum, around half

of all survey participants report taking no measures to protect themselves against tracking and data collection even though they would like to have such protection. Which means that, at a minimum, methods of ensuring such protection must be made more accessible, more effective and more visible – or companies could forego the collection of personalized data.

Data Protection Is Important to Companies, But They Feel They Have Little Choice

On the flip side of the coin, the company representatives interviewed for this study say they have limited room for maneuver if they want to take advantage of the personalized advertisements made available by the big players in the industry. “If you want to participate in the

The technology sector's energy consumption is attracting critical attention from many directions, but such concerns generally don't play a significant role in the day-to-day of online marketing.

market, there are certain things that you simply cannot avoid. [...] For example, one of our most important [revenue] sources is Google as an advertising medium. And there, we have no influence. We simply have to go along with the market, otherwise we wouldn't be present to the degree we currently are if we didn't participate," says the head of marketing for an online mail-order retailer.

The companies interviewed for this study are primarily focused on issues pertaining to data protection because legal mandates, such as the EU's General Data Protection Regulation (GDPR) require them to do so, and they would like to avoid violations. It's not just 80

percent of the households surveyed that believe the use of AI for the personalization of online advertising produces at least some disadvantages for consumers when it comes to data protection and privacy – a similar share of companies do as well (Figure 3).

Those companies that expressed a greater focus on sustainability during their interviews for the study noted their intention to approach customer data with a greater degree of responsibility. A representative of one such company said during the interview that they only collect data that is directly relevant for marketing purposes and only serve ads sparingly and in a targeted manner. Such

strategies, though, are not widespread. "I would tend to see the advantages [of personalization] or the opportunities it opens up. For users, I don't see much of a danger for the time being," said the marketing director of a smaller online marketplace.

Environmental Effects: The Blindspot for AI in Marketing?

The technology sector's energy consumption is attracting critical attention from many directions, but such concerns generally don't play a significant role in the day-to-day of online marketing. That lack of attention to sustainability is further encouraged by the fact that using technologies based on AI and on the findings of data science are constantly becoming cheaper and easier to use. That in turn has led to the industry's rising energy consumption. "As the technology for computing becomes cheaper, there is a tendency for being less mindful of the performance

HOW STRONGLY DO CONSUMERS PROTECT THEIR DATA?

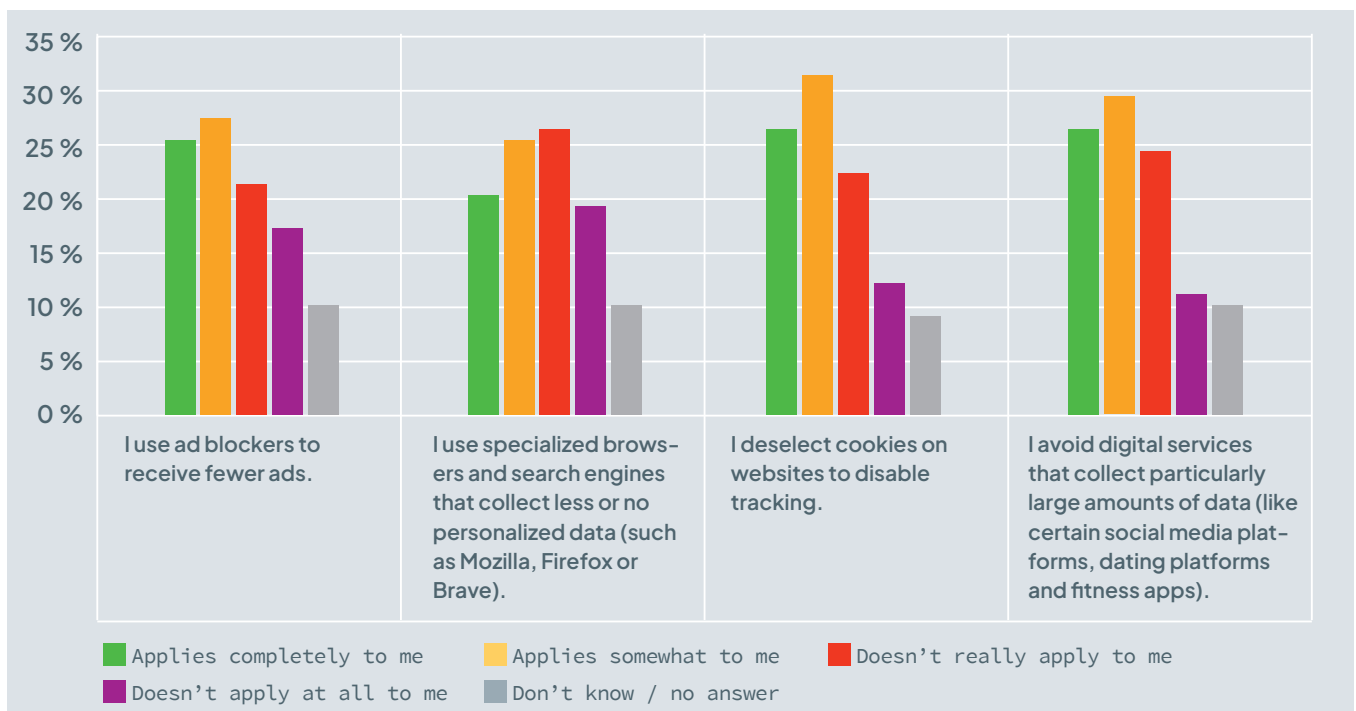


Figure 2

HOW DO HOUSEHOLDS AND COMPANIES VIEW THE OPPORTUNITIES AND RISKS OF AI-SUPPORTED PERSONALIZED ADVERTISING?

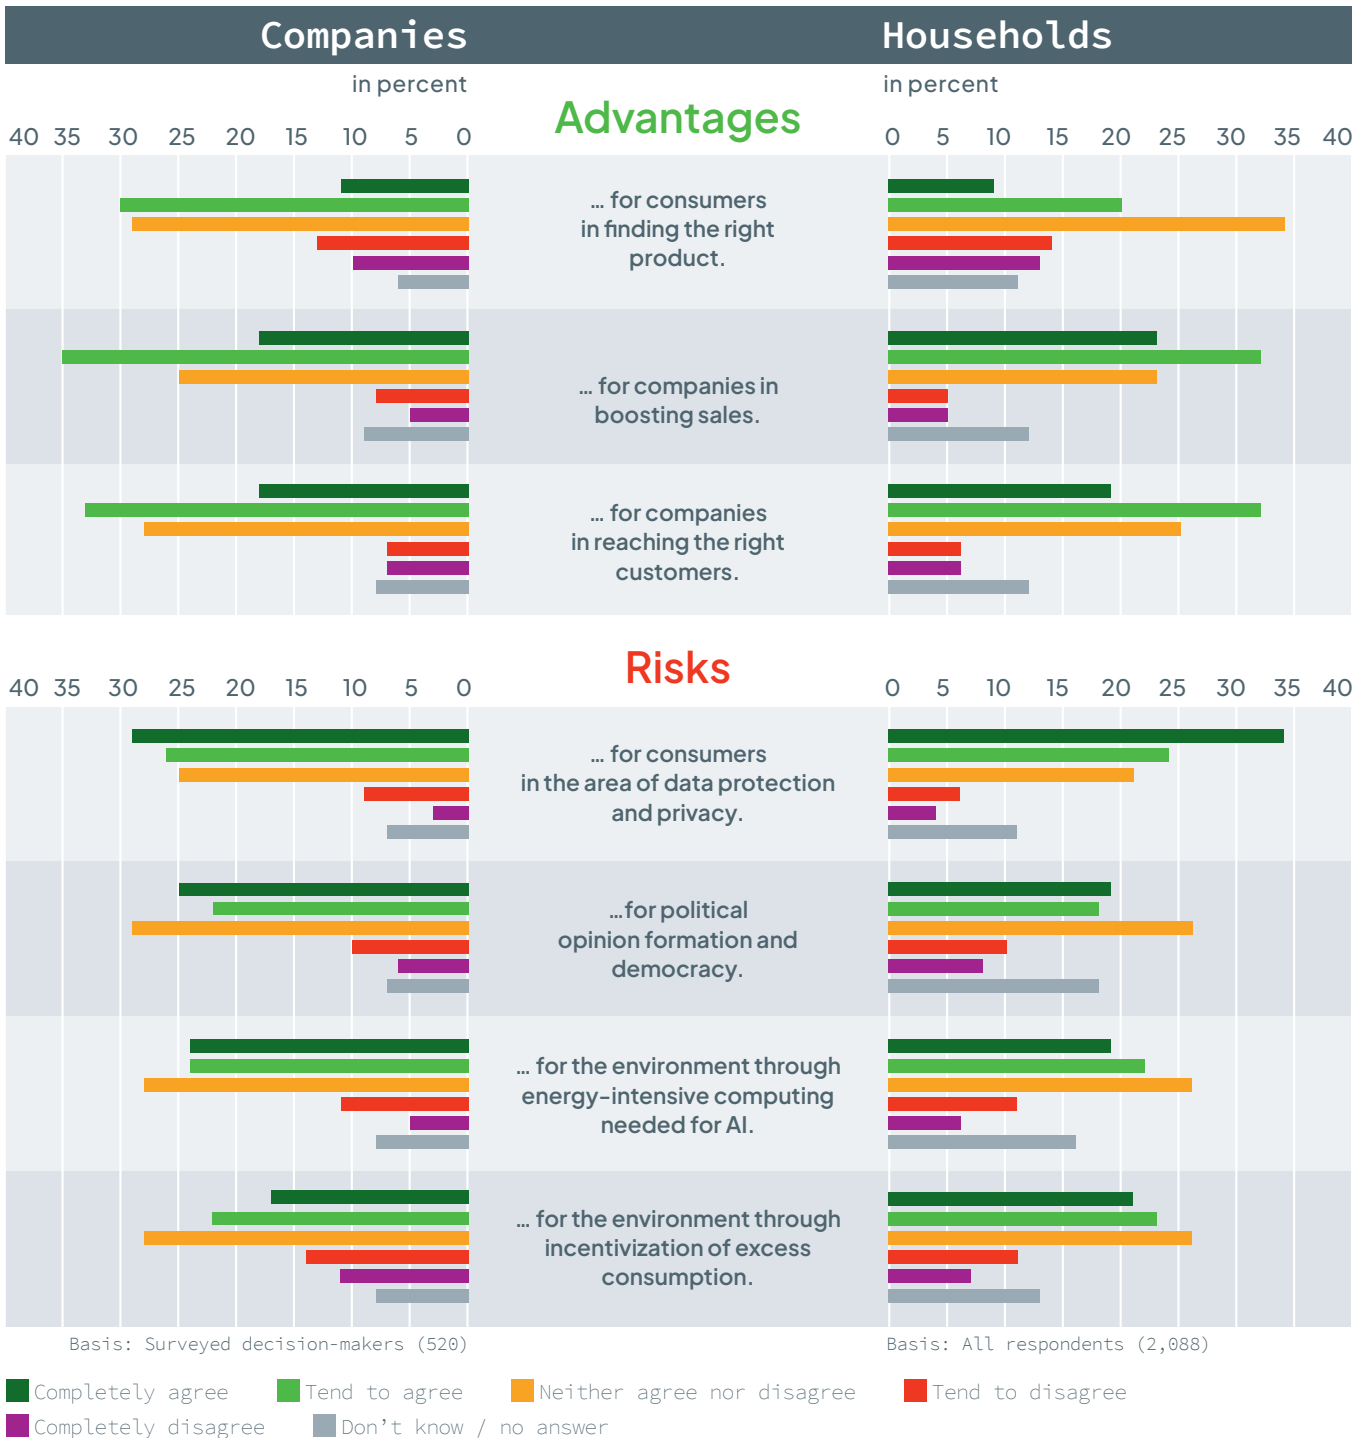


Figure 3

of certain algorithms. Cloud providers make it very easy for you to scale up the resources if you need to. So then, if you want, you can have a very energy-intensive algorithm and just run it, and it will run fast," a data scientist from a

marketing agency said. Effectiveness, speed and costs are more important than environmental efficiency, he continued. "It's no secret how energy-intensive machine learning is. But I hardly think that there are attempts to improve

energy efficiency to benefit the environment, not even from those who are environmentally aware."

Both small and large service providers do offer to calculate carbon footprints.

But our interview partners doubt the effectiveness of such measures. “Google has some services that allow you to calculate the CO₂ spend that you have while running the services. [...] There are some data sources where they put a little green leaf that says renewable energy has been used to produce them. But again, it’s not very transparent and it’s very easy to greenwash.” Furthermore, it is difficult to force market leaders like Google, Microsoft and Amazon to improve transparency and become more sustainable. That limits the options of companies that rely on advertising but would like more environmentally friendly options.

Further environmentally damaging effects can be produced when advertising results in the sale of more products and, by extension, the consumption of more energy and resources. It is difficult to quantify the degree to which this effect impacts the environment. But companies naturally hope that advertising will boost their sales and they don’t critically examine that goal. The immense amounts of money spent on personalized advertising suggest that such advertising strategies have the potential to significantly increase consumption. One marketing-team member at a mid-sized

”Needs are met that 20 years ago we didn’t even know we might have one day.”

mail-order company notes: “[Through marketing], needs are met that 20 years ago we didn’t even know we might have one day. For me, that is evidence that marketing produces demand.” The use of AI for the personalization of advertising serves to effectively convert the interests and preferences of users into purchases by activating new needs.

Over half of company representatives and survey respondents at large are aware of the dangers posed to data and privacy protection.

A Lack of Knowledge among Users and Companies

When asked if they were aware that AI is being used to personalize advertisements on the internet, only 55 percent of survey participants from private households responded with “yes.” Only 29 percent feel that they are well informed about the kinds of personal data collected, when it is collected and what companies are doing the collecting. Over half of company representatives and survey respondents are aware of the dangers posed to data and privacy protection. But many companies and private citizens are unsure about the risks posed by AI to the environment or to the formation of political opinions and to democracy (Figure 3).

This lack of knowledge, combined with the lack of transparency in marketing practices, explains why actors in both civil society and in science are demanding that data protection not be reduced to individual decisions, but that it be defended as a legally protected fundamental right. Organizations like Germany’s Digitalcourage, the Chaos Computer Club (CCC) and Netzpolitik, in addition to companies like Mozilla and political actors like the European Commission have for years been pointing to the opacity, surveillance, manipulation and discrimination of the advertising industry. “The biggest risk is the possibility for manipulation. [...] The more detailed the information collected about people [...], the easier it is to manipulate them or evoke a certain behavior,” says a representative of an association that promotes digital data protection.

Many experts from the worlds of civil society, academia and politics are thus demanding the labeling of Artificial Intelligence in advertising campaigns or a general ban on the use of personalized data in online advertising. Personalized online marketing must be regulated and restrictions on advertising discussed to protect both our democratic societies and the environment.

Complete Study:

Frick, V., Marken, G., Schmelzle, F. & Meyer, A. (2023). “The (Un-)Sustainability of Artificial Intelligence in Online Marketing. A Case Study on the Environmental, Social and Economic Impacts of Personalized Advertising.” IÖW Schriftenreihe 228/2023. ISBN 978-3-940920-33-1.



**GESA MARKEN
AND VIVIAN FRICK**

... are Scientific Research Assistants at the *Institute for Ecological Economy Research (IÖW)* in Berlin. They examine the role of digitalization in the social-ecological transformation.

The Environmental Impact of AI: Just Measure It

With more and more resources being expended on developing and applying Artificial Intelligence technologies, it is also increasingly important to understand the impact these technologies are having on the environment and climate.


AI's underlying infrastructure must be environmentally sustainable and respectful of our planet's limits. At the same time, the discussion about the relationship between the benefits of AI systems and their environmental costs must be grounded in facts and figures. Currently, however, the developers and operators of these systems are not providing the data necessary. This absence of publicly available information hampers the development and enactment of effective policies.

The European Union's AI Act, which is currently being drafted, could for the first time require companies to measure and disclose information on the environmental impact of certain AI systems. The European Parliament has proposed requiring companies to measure the energy and resource consumption of foundation models and high-risk systems. This requires that data collection methods be integrated into these systems.

Critics often claim that the obligation of measuring the environmental impacts of AI systems is too complicated and places too great a burden on small and medium-sized enterprises in particular – and ultimately hinders innovation. But easy-to-use measurement methods already exist for monitoring energy consumption, CO₂-equivalent emissions, water consumption, the use of minerals for hardware and the generation of electronic waste and thus assessing the sustainability of AI systems.

Getting the Whole Picture

Without a comprehensive life cycle analysis, we can't adequately capture the environmental footprint of AI models. Providers of large language models (LLMs), in particular, tend to disclose only the direct energy consumption and emissions for a single training cycle. The result is an incomplete picture.



Consider, for example, the training of the BLOOM model. Energy consumption during the training phase corresponds to the emission of around 24.7 tons of CO₂ equivalents. But if you factor in hardware production and operational energy, the emissions value doubles. And this does not yet include the continuous emissions produced during the application of the model. Reliable figures from this inference phase are lacking, but first indicators suggest that emissions could be immense – both in the production of the necessary hardware for the application and during operation. That's why we need to measure how AI systems impact the climate throughout their entire life cycle – from raw material extraction and CO₂ emissions to pollution and water consumption – so that we can enable informed decisions and targeted policies based on solid knowledge.

Logging Relevant Data

Companies can already automatically log and report much of the data needed to assess the sustainability of AI systems, such as operational data from computer systems – i.e., how




often calculations are performed and how long these processes take. When such metadata is stored in a spreadsheet, it can be used to generate efficiency metrics. Metrics such as “Power Usage Effectiveness” (PUE), for example, show how much energy a data center uses for computing in relation to its overall energy consumption. This parameter makes it possible to compare the energy efficiency of data centers. By examining power consumption, the energy mix of the data center, the carbon intensity of the energy grid and the percentage of CO₂

the provider is potentially compensating, emissions can in turn be calculated.

During system development and training alone, the data listed in the table below should be recorded to comprehensively assess and compare the energy consumption of AI systems. Similar requirements can be formulated for all other environmental impacts, such as emissions, water consumption, mineral extraction and hardware disposal.

STAGES OF THE AI LIFE CYCLE ASSESSMENT



Source: Luccioni et al. 2022

There are already many measurement methods available for assessing the environmental impact during system development and training, material extraction, hardware manufacturing and disposal as well as different ones for tracking carbon.

There are already many measurement methods available for assessing the environmental impact during system development and training, material extraction, hardware manufacturing and disposal as well as different ones for tracking carbon. Some hardware manufacturers already report emissions levels for a handful of their products. Other approaches for assessing environmental impact during system deployment still need to be developed – reliable metrics and comparable units of measurement for assessing emissions during the application phase, for example.

To be as accurate as possible about the environmental impact during the deployment phase, we propose that AI providers define various standard usage scenarios prior to market launch.

Detailed and Standardized Reporting Needed

The life cycle approach demonstrates that various stakeholders need to provide accurate measurements. For instance, hardware manufacturers such as Nvidia should disclose environmental data on products that are widely used in the development and application of AI models.

Measuring Environmental Impacts During System Deployment

Procedures and methods for assessing environmental impact during system deployment have not yet been established. Developers of AI systems can record their energy consumption during training. However, under the requirements formulated within the AI Act to document energy consumption, this

ENERGY CONSUMPTION DURING SYSTEM DEVELOPMENT AND TRAINING

Impact	To Report	Process	Purpose	Source
Energy	Hardware used (e.g., number of GPU models)	Documentation	- Calculation energy use & emissions - Resource consumption manufacturing	Provider / Data Center
	Number of FLOPs	Documentation	Calculation energy use & emissions	
	Computing time	Documentation	Calculation energy use & emissions	
	GPU hours (equivalent depending on hardware)	Documentation	Calculation energy use & emissions	
	Energy used	Documentation	Calculation emissions	
	Power Usage Effectiveness of data center (PuE)	Documentation	Calculation energy use & emissions	
	Hardware energy consumption: - Infrastructure consumption (consumption without computing) - Idle consumption (consumption during computing standby) - Dynamic consumption (consumption with computing running)	Documentation/ information request	Calculation energy use & emissions	

OVERVIEW OF METHODS ALONG THE AI LIFE CYCLE

	Raw Material Extraction	Material Manufacturing	Equipment Manufacturing	System Development and Training	System Deployment	Disposal / End of Life
Energy			●	●	●	
Emissions (CO ₂ e)	●	●	●	●	●	
Water	●	●	●	●	●	
Minerals	●	●				●
Electronic Waste						●

will most likely not be feasible during inference. Thus, energy consumption during the application phase and the extent of emissions generated during this phase must be estimated. To that end, we are proposing two basic options, perhaps in combination:

- Before an AI product is released on the market, different standard use scenarios (low-, middle-, high-use) should be evaluated based on test runs or, preferably, simulations.
- After market introduction, the de-facto average energy consumption over a certain period of time should be calculated. This would allow the estimated standard use scenarios to be evaluated and adjusted if they deviate significantly from the actual value.

Greater Transparency Is Feasible - and Overdue

There is no lack of technical means for measuring the environmental impact of AI systems. There is, however, still a lack of political will to make AI more sustainable. This is all the more irresponsible considering that AI is a resource- and energy-intensive technology that is becoming increasingly pervasive in all areas of life. The European Parliament has taken some important steps in the right direction to ensure that AI does not further harm the environment, the climate, people and the planet. Nevertheless, data on environmental impacts of AI systems is indispensable. Clear and comprehensive requirements must be introduced for publicly available reports that include

such data. This could make AI systems more environmentally sustainable, while at the same time distributing their risks, harmful consequences and benefits more equitably around the world. If the EU is serious about aligning the use of AI with the common good, then it should put all people, and not just Europeans, at the center of its focus. Whatever form the AI Act eventually takes: People will only be truly protected from the negative consequences of AI systems if their impact on the environment is effectively monitored.



ANNE MOLLEN

... is a post-doc researcher at the University of Münster and Senior Research Associate at *Algorithm-Watch*. She works on the sustainability of automated decision-making (ADM) systems and questions of global justice in relation to ADM.



KILIAN VIETH-DITLMANN

... is Deputy Team Lead for policy and advocacy at *AlgorithmWatch*. His research and advocacy work focuses on the use of algorithmic decision-making systems in the public sector and on the sustainability of AI.

The Costly Backbone of Consumerism

As is true of the internet in its entirety, digital advertising is boosting carbon emissions. Currently, there are nearly 2 billion websites online, and the number of advertisements on them is ballooning rapidly.



The digital world's carbon footprint might even be larger than that of the aviation industry.¹ Still, a [recent study](#) revealed that six out of 10 consumers are unaware that digital ads and internet browsing produce carbon emissions.

The online tracking-based advertising ecosystem is fueling the climate crisis in three ways:

First, the technical processes on which the collection of data and the building of user profiles are based increasingly rely on complex machine learning systems, which produce additional carbon emissions. Globally, the internet accounts for 4 percent of CO₂ emissions, with the energy required for digital advertising making a significant contribution to that total. Emissions generated by digital services are [predicted to double by 2025](#).

Second, advertisements are designed to encourage us to buy things we don't necessarily need. A recent [study](#) by Purpose Disruptors, a network of advertising and marketing insiders pushing the industry to take responsibility for the emissions it drives, estimated that personalized advertisements caused the UK's per capita emissions to increase by 32 percent in 2022.

Third, tracking-based advertising helps spread misinformation, which leads to social disruption. The same algorithms that profile users for advertising purposes can also be used as tools for disseminating climate disinformation, fostering alternative reality bubbles and creating conditions that increasingly

threaten democracy itself, all of which is incompatible with the development of fair and ambitious climate policy. Since this aspect has already received extensive coverage in other publications, this article focuses on the carbon impact of advertising technology and on the consumption it drives.

Wasting Energy on the Search for the Highest Bid

A single digital advertisement only consumes a relatively small amount of energy, but the prevalence of such advertising has a huge global impact. Ryan Cochrane, chief operating officer of the global AdTech platform Good-Loop, [estimates](#) that the advertising ecosystem handles 2,000 times more bids than the 8 trillion transactions undertaken by the New York Stock Exchange on any given day. The Irish Council of Civil Liberties [projects](#) that the industry's real-time bidding system places 178 trillion advertisements in the U.S. and Europe each year.

The vast majority of that 1 percent is effectively wasted energy that leads to zero ads being placed due to the underlying auction system.

[New data](#) shows that the energy required for a single ad impression is equivalent to approximately one gram of carbon dioxide emitted. This value includes the energy consumption of user devices, data servers and ad servers. It includes activities such as data collection, processing, storage

¹ Lancaster University, "Emissions from computing and ICT could be worse than previously thought," 2021



and management in addition to the creation of detailed user profiles. Many of the complex machine learning algorithms that transform raw data into actionable and valuable insights for advertisers are also extremely energy intensive. According to a [recent report](#) by Global Action Plan, online advertising accounts for about 1 percent of total global energy consumption.

The vast majority of that 1 percent is effectively wasted energy that leads to zero ads being placed due to the underlying auction system. For each of the trillions of advertisements placed each day, there are several prospective bidders – sometimes even thousands – for the slot in question. The business rationale behind this system ensures that the highest price is paid for each slot, but it comes at a cost: the AdTech ecosystem’s exploding emissions.

An additional aspect is the fact that much online advertising is fraudulent. Far from being a minor problem, the World Federation of Advertisers estimates that fraudulent advertising will soon be worth over \$50 billion per year. Indeed, it is organized crime’s [second largest source of income](#) after the illicit drug trade. The AdTech industry estimates that [the value of fraudulent advertising is around 4 percent of its own business value](#), while independent researchers have [come up with numbers well above 25 percent](#). Given the vast number of online ads served every day, the total is significant no matter which estimate is correct.

Buy!

Advertising is an industry generating hundreds of billions in revenue every year, which, as mentioned above, significantly

boosts the carbon footprint of consumers by encouraging consumption. The New Weather Institute [found](#) that there was “sound empirical evidence” that advertising is indirectly responsible for “climate and ecological degradation through its encouragement of materialistic values and goals.” Indeed, those working in advertising must face the uncomfortable truth that [“the better you do your job, the more damage you cause.”](#)

Happily, an alternative is already available: To reduce the climate impact of AdTech, we must switch to contextual advertising systems. Such systems use the context of the webpage content and post contextually relevant advertisements to the webpage. Such context-based ads aren’t matched to the multitude of a website’s users, but only to the single webpage in question, thus decreasing the number of computing processes necessary and, by extension, the amount of energy consumed. Not only does contextual targeting obviate the need for collecting masses of personal data and building profiles. It also decreases the number of ad impressions up for bid, thus decreasing the amount of energy consumed by the online auction system.

We can’t expect the advertising industry to voluntarily move away from what has proven to be a cash cow.

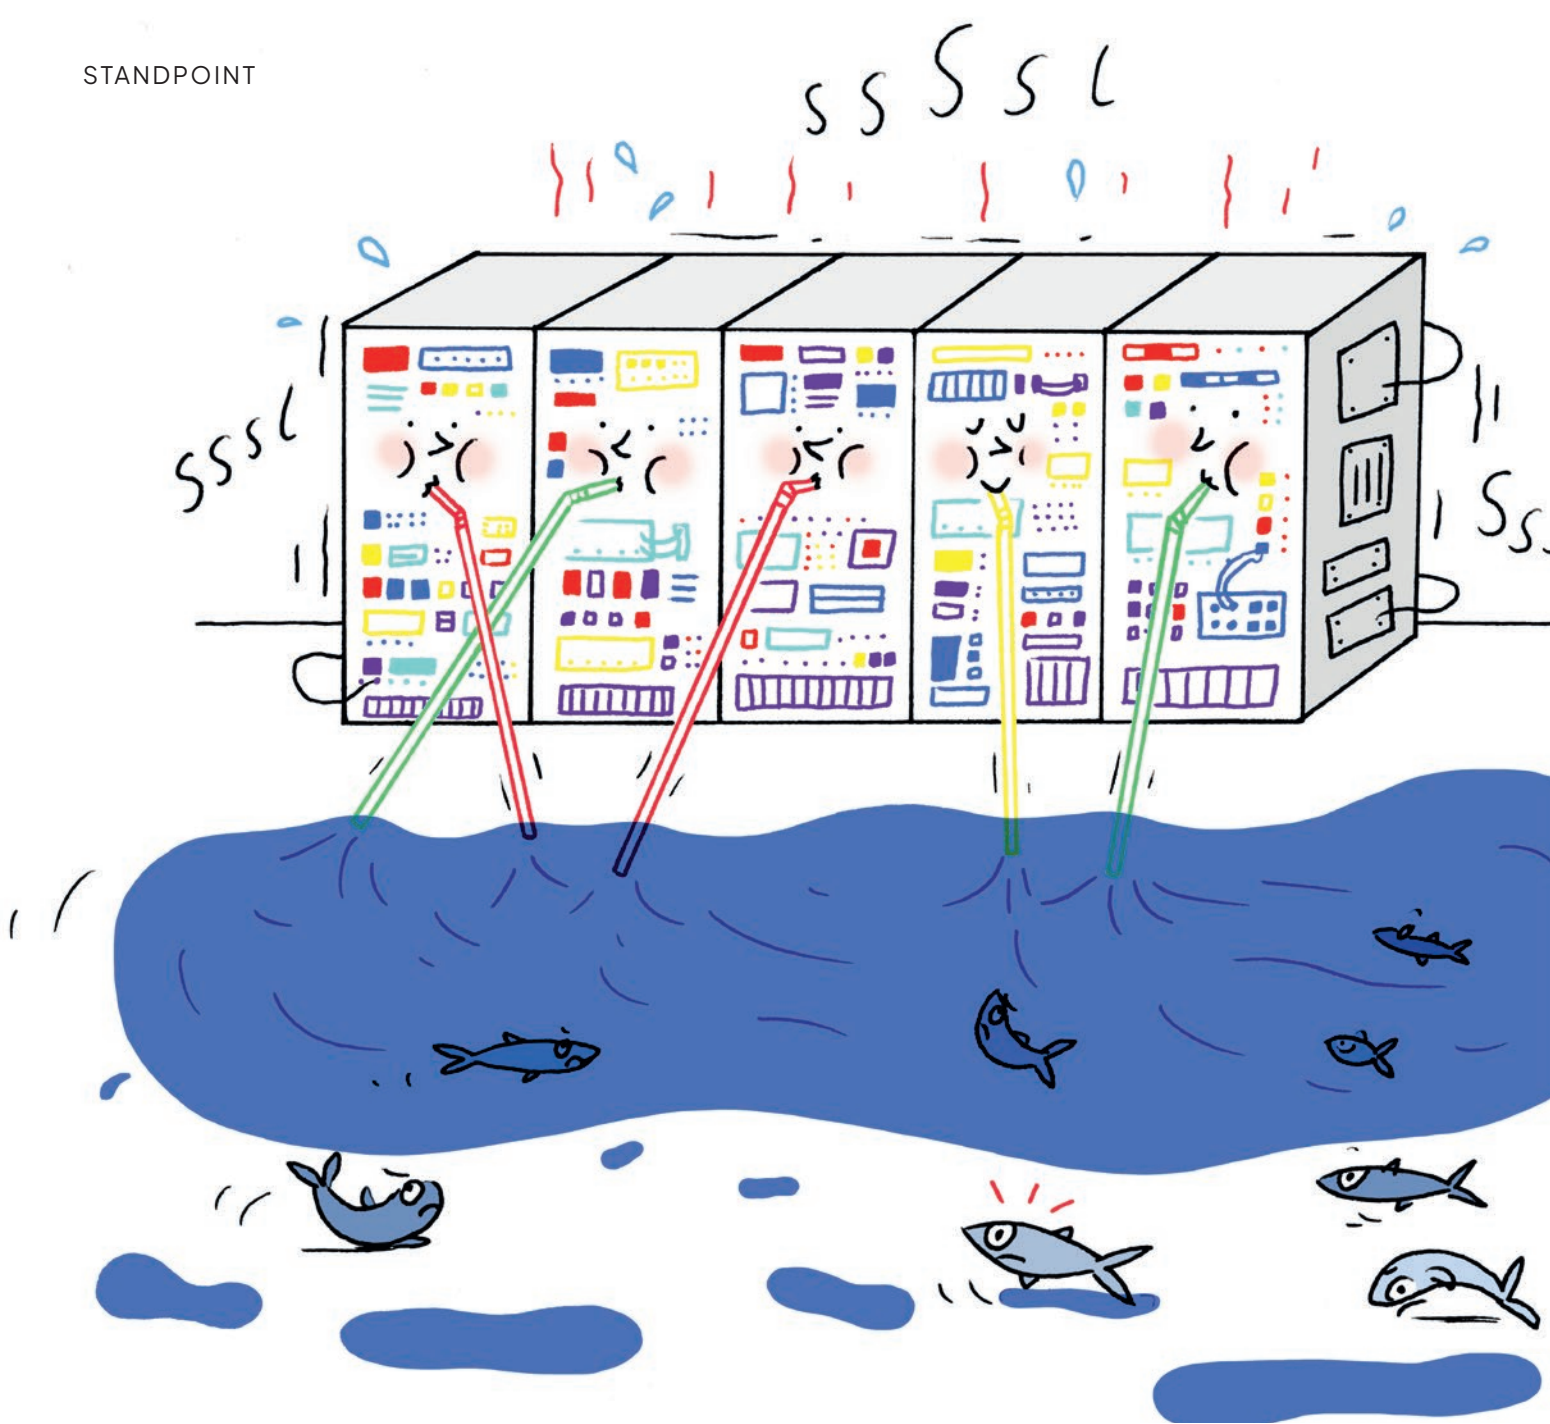
We can’t expect the advertising industry to voluntarily move away from what has proven to be a cash cow. We urgently need policy makers to outlaw the targeting of ads based on personal data, similar to the provision of the European Union’s Digital Services

Act provision banning the targeting of children. Only then can the climate impact of AdTech be significantly reduced.



**DUNCAN
MCCANN**

... is Head of Accountability at 5Rights, a non-governmental organisation working on the rights of children and young people in a digital world. Duncan joined 5Rights from the New Economics Foundation, where he led the Digital Economy Programme. Duncan has a legal background as well as a decade of experience working in the tech industry for Cisco Systems.



Water Is the New CO₂

With advances in Deep Learning, which uses artificial neural networks to analyze large data sets, Artificial Intelligence has become a real game changer. Advances in technological capabilities are leading to scientific breakthroughs and accelerating business growth, and they appear to be offering solutions to global challenges in important areas such as the climate crisis.

But the success of AI relies heavily on computationally intensive calculations to learn useful patterns from data during training and to check whether the predictions based on them are accurate during inference, the application phase of AI systems. As such, AI models, especially large generative models like GPT-3 and LaMDA, are typically trained on large clusters of servers, each with multiple graphic processing units (GPUs), creating an enormous appetite for energy. To curb the tremendous energy demand of AI systems, it is time we address their environmental footprints in their entirety.

While a low carbon footprint has now entered the public consciousness as an indicator of sustainability, the water footprint of AI systems – the fresh water consumed for generating

electricity and cooling servers – is still given too little attention. Even putting aside the significant water toll of chip manufacturing, training a large language model like GPT-3 and LaMDA can easily evaporate millions of liters of fresh water for cooling the power plants and AI servers. This is all the more concerning as water becomes increasingly scarce due to rapid population growth and/or outdated water infrastructures, especially in drought-prone areas. Water scarcity has become one of our greatest global challenges. Despite all the efficiency gains being made in the field of Artificial Intelligence in terms of resource consumption, the exponential growth in demand is resulting in an ever-increasing water footprint. For example, Google's direct water consumption increased by 20 percent between 2021 and 2022, and even doubled in certain drought-hit areas. Microsoft saw a 34-percent increase in its direct water consumption over the same period.

The Scope-2 indirect carbon footprint caused by electricity usage for training is routinely recorded in the model cards of published AI models. Yet, not even the direct water consumption during AI model training is included in the model card, not to mention the indirect water consumption tied to its electricity usage. To some extent, withholding information about AI's water footprint is comparable to not including calorie content in the nutrition facts label of a food product. Such a lack of transparency is more than just an impediment to innovations that could improve water sustainability. It's also difficult to reconcile with recent statements made by major technology companies in regard to water. Google, for example, announced its intention to become water neutral by 2030.

Developers of AI models need to take urgent action to curb growing water consumption. A first and crucial step would be to increase transparency and publicly disclose how much water is used to train and infer AI models, both directly for cooling AI

servers and indirectly for generating the electricity to power them. An AI model's water footprint should be recorded in its model card. Only with the availability of this information will it be possible to holistically benchmark AI's environmental footprint. This measure would complement current efforts to make water

Thirsty AI

ChatGPT needs about 500ml water for a simple conversation of 20-50 questions and answers. Since the chatbot has more than 100 million active users, each of whom engages in multiple conversations, ChatGPT's water consumption is staggering. And it's not only the application's operational mode: Training GPT-3 in Microsoft's state-of-the-art U.S. data centers would directly consume 700,000 liters of clean freshwater (enough for producing 370 BMW cars or 320 Tesla electric vehicles) and water consumption would be three times that if training were performed in Microsoft's Asian data centers.

These estimates are taken from the yet to be peer-reviewed study *Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models*. The study's authors provide a methodology to evaluate the water footprint of AI models, using public data sources. They also explain how developers could reduce the water footprint of their AI models and increase water efficiency by scheduling AI model training and inference in different places and at different times.

["Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models," by Pengfei Li, Jianyi Yang, Mohammad A. Islam, Shaolei Ren](#)

supplies more sustainable, such as integrating more water-efficient techniques for cooling servers into AI systems. In addition, if the AI water footprint became more transparent, developers could exploit the spatial and temporal flexibilities of AI and train and deploy AI models in places where their footprint will be smaller. It also enables flexible trade-offs: If the AI model is deployed in a water-stressed area, it would probably make more sense to use a compact model with a smaller water footprint than a full, more resource-intensive model. Knowing AI's water footprint data could also mitigate the environmental inequity that is accelerated by AI systems. We could move AI workloads around to equitably balance AI's water footprint across different regions rather than letting a few disadvantaged and drought-stricken areas disproportionately bear the negative impact.

We can no longer allow the water footprint of AI systems to remain under the radar. It must be prioritized as part of the global fight against water scarcity. The first step is simple: We need to measure AI's water footprint and make that information public.



SHAOLEI
REN

... is an Associate Professor of Electrical and Computer Engineering at the University of California, Riverside. His research broadly focuses on AI and Sustainability, with the goal of building an environmentally sustainable and equitable future.

The SustAI In Project Is Ending, But Work on Sustainable AI Is Just Beginning



The third issue of SustAI In magazine marks the end of our project “SustAI In: The Sustainability Index for Artificial Intelligence.” We have posted the project’s main conclusions at <https://sustain.algorithmwatch.org/>, where you will also find:

- ▶ the three issues of SustAI In magazine
- ▶ our Self-Assessment Tool for organizations for determining how sustainable their AI systems are
- ▶ our guidelines for the sustainable development of AI systems

The project has laid the foundations for launching a broader discussion on AI sustainability. Our indicators show that the sustainability of AI systems isn’t defined exclusively by their energy consumption. Other factors must also be taken into account along the AI value chain: issues like data protection, unfair global conditions, the working environment for those involved in the systems’ production and the AI industry’s high barriers for market entry.

It is now up to decision-makers in politics, society and industry. We have presented suggestions on the ways AI systems can be made more sustainable – and AI practitioners and regulators must now take things from here.

Many questions remain unanswered: How do we deal with contradictions between different sustainability criteria? In relation to a given AI system’s high-level performance, up to what point is enormous energy consumption still justified? To what extent does an AI system’s performance justify its enormous energy consumption? What countermeasures can be taken against market concentration in the AI industry? To answer those questions, more research and a greater willingness to make AI more sustainable are urgently needed.

Imprint

Sustain Magazine #3

October 2023

Available online at

<https://bit.ly/SustainEN>

Publishers

AW AlgorithmWatch gGmbH

Responsibility (according to § 5 TMG and § 55 RStV): Matthias Spielkamp

Linienstr. 13

10178 Berlin

Germany

Content Lead

Anne Mollen & Kilian Vieth-Ditlmann

Editorial Lead

Waldemar Kesler

Content Support

Vivian Frick (Institut für ökologische Wirtschaftsforschung)

Gesa Marken (Institut für ökologische Wirtschaftsforschung)

Andreas Meyer (DAI Laboratory, TU Berlin)

Friederike Rohde (Institut für ökologische Wirtschaftsforschung)

Frieder Schmelzle (Institut für ökologische Wirtschaftsforschung)

Josephin Wagner (Institut für ökologische Wirtschaftsforschung)

Illustrations

Juliette Baily, Kevin Lucbert, Beate Autering

Media

Marc Thümmeler

Design / Layout

Beate Autering

Translation and Copy Editing

Dacha Media GbR: Daryl Lindsey, Charles Hawley

Output from the project

"SustAI: The Sustainability Index for Artificial Intelligence"



Project partners:



This publication is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License

<https://creativecommons.org/licenses/by-nd/4.0/>

Supported by:



based on a decision of
the German Bundestag



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/>



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/>



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/>

