

Background & explainers

- [AI woordenlijst](#)
- [Hoe werkt generatieve AI?](#)
- [101: Ecosystem of AI](#)
- [Boeken](#)

AI woordenlijst

Een verzameling AI-jargon om je op weg te helpen!

Algoritme

Een set instructies of regels op basis waarvan een computer (of mens) een berekening kan maken of een probleem kan oplossen, vergelijkbaar met een recept. Zie ook [deze video](#), voor een humoristische interpretatie (en consequenties) hiervan.

Bias

Generatieve AI werkt vanuit een 'model' dat de computer zichzelf geleerd heeft op basis van een bepaalde dataset. Dit proces heet trainen, en de dataset kan een verzameling teksten zijn (boeken, blogposts, e-mails), een verzameling afbeeldingen (profielfoto's, kunstwerken), een verzameling geluiden (muziek, soundeffects). Een verzameling data is nooit compleet, er ontbreken altijd dingen die niet gedocumenteerd zijn, of die om verschillende redenen expres verwijderd zijn. Dit kan er voor zorgen dat bepaalde onderwerpen of (bevolkings)groepen niet of verkeerd gerepresenteerd worden in de data en daarmee in het model. Dit kan grote gevolgen hebben, [lees hier](#) een technische achtergrond.

Machine Learning

Termen als AI, Machine Learning, Deep Learning en generatieve AI worden vaak en veelvuldig door elkaar gebruikt. Wat is wat? In een sterke versimpeling gaan al deze termen over dezelfde processen, maar is elke term een sub-onderdeel van de andere. AI (Kunstmatige intelligentie) is de term voor alle (computer)systemen die beslissingen kunnen nemen op basis van vooraf gedefinieerde regels. Daarbinnen valt Machine learning, waarin deze systemen patronen leren herkennen, zoals spraak- en beeldherkenning. Onderdeel daarvan zijn Deep Learning en/of neurale netwerken, waarbij de patroonherkenning complexer wordt en over meerdere lagen loopt. Generatieve is de meest specifieke set binnen deze termen, waarin de netwerken ingezet worden om *nieuwe* data te genereren gebaseerd op de geleerde patronen. Zie ook deze afbeelding:

The Original Benny G. CC BY-SA 4.0 <<https://creativecommons.org/licenses/by-sa/4.0/>>, via Wikimec
Afbeelding [Wikimedia Commons](#)

Prompt

De opdracht die je invoert om een (generatieve) AI aan het werk te zetten. Bijv: "schrijf een essay over de opkomst van AI", "teken een paard op een astronaut", of "schrijf een liedje over het verdriet van de zon".

Training

Voordat een AI kan interpreteren of genereren moet deze getraind worden. Dit gebeurt op basis van een trainingset (boeken en webpagina's voor tekst, foto en film voor beeld, geluid en muziek voor audio) en vindt plaats *voordat* de eindgebruiker ermee in aanraking komt. De training is meestal het meest energie-intensieve deel van AI.

Het model waar jij als gebruiker mee werkt is al getraind en kan in die zin niet bij leren. Huidige systemen zoals Chat-GPT kunnen extra informatie aan hun proces toevoegen tijdens het gesprek, maar dit wordt *niet* toegevoegd aan de dataset. Meestal wordt dit opgelost door deze informatie op de achtergrond (en dus voor de gebruiker onzichtbaar) toe te voegen aan de ingevoerde prompt.

Hoe werkt generatieve AI?

Een zeer uitgebreid overzicht AI tools en achtergronden vind je [hier](#).

Generatieve AI

Generatieve AI is een vorm van kunstmatige intelligentie die nieuwe gegevens kan maken op basis van een model dat getraind is op een (grote) dataset. Die nieuwe gegevens kunnen vele vormen aannemen, zoals tekst, beeld en geluid. Hieronder verzamelen we een aantal explainers die in vogelvlucht uitleggen hoe de meest voorkomende modellen werken. Er is nog veel meer over te vinden, maar deze video's zullen je in elk geval een basisbegrip en het vocabulaire kunnen geven om je op weg te helpen.

Tekst (ChatGPT, Copilot, Claude)

Huidige taalmodellen zoals ChatGPT en Copilot lijken op chatbots. Je kunt ze van alles vragen, en ze geven hun antwoorden met zulke stelligheid dat je eraan twijfelen soms niet eens meer in je op komt. Waar komen deze teksten vandaan? En klopt het wel wat ze zeggen?

<https://www.youtube.com/embed/aQguO9IeQWE>

Afbeeldingen (Midjourney, Dall-E, StableDiffusion)

Je hebt vast wel een keer een zin in een browserscherm ingetypt en verwonderd (of niet-onder-de-indruk) gekeken naar het plaatje dat tevoorschijn kwam. De afbeelding die gemaakt wordt door generatieve AI is geen bestaande afbeelding, geen collage, en ook geen willekeur. Hoe komen deze plaatjes tot stand? In deze video krijg je een korte geschiedenis van de beeldgeneratie. [Vanaf 8:46](#) gaat het over de huidige tools als DALL-E, Midjourney en Stable Diffusion (de zogenaamde Diffusion models)

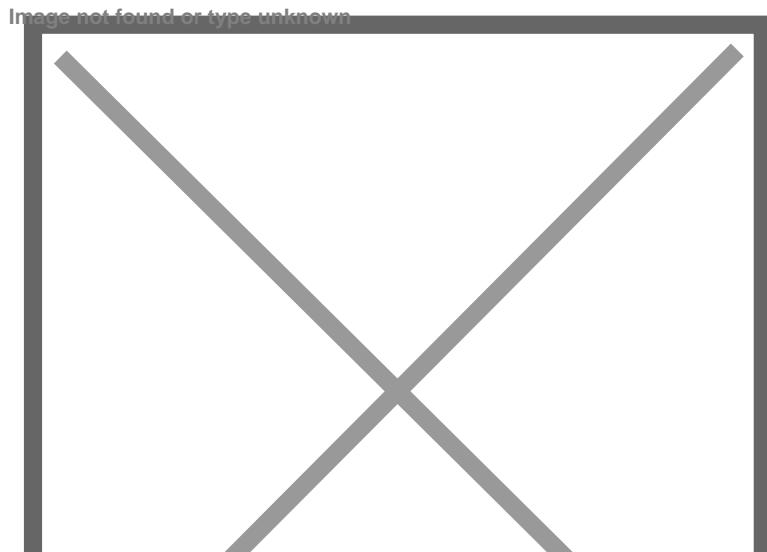
https://www.youtube.com/embed/Rke0V_VkF3c

101: Ecosystem of AI



The massive ecosystem of AI relies on many kinds of extraction: from harvesting the data from our daily activities and expressions, to depleting natural resources and to exploiting labor around the globe so that this vast planetary network can be built and maintained.

This guide gives you insights, numbers and examples of art projects to help you as a maker navigate this field.



Cartography of Generative AI shows what set of extractions, agencies and resources allow us to converse online with a text-generating tool or to obtain images in a matter of seconds, by Estampa
“Cutting-edge technology doesn’t have to harm the planet”

Experimenting with AI

As a maker it can be daunting to experiment in this field, some tips (*this is a growing list*):

1. Be picky

Using large generative models to create outputs is far more energy intensive than using smaller AI models tailored for specific tasks. For example, using a generative model to classify movie reviews according to whether they are positive or negative consumes around 30 times more energy than using a fine-tuned model created specifically for that task.

The reason generative AI models use much more energy is that they are trying to do many things at once, such as generate, classify, and summarize text, instead of just one task, such as classification.

Be choosier about when they use generative AI and opt for more specialized, less carbon-intensive models where possible.

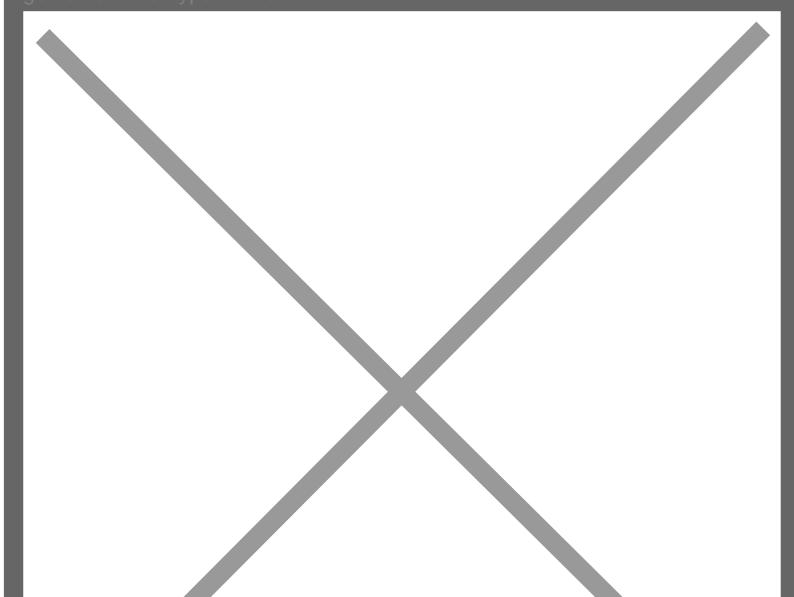
2. Use tools to keep track

Code Carbon makes these calculations by looking at the energy the computer consumes while running the model: <https://codecarbon.io/>

3. Run your models locally

Running AI models locally helps you to have increased control over energy usage and resource allocation. By managing models on personal or dedicated hardware, you can optimize efficiency and reduce energy consumption compared to relying on large, centralized data centers. This localized approach also minimizes the environmental impact, as it circumvents the need for extensive data transmission and the associated carbon emissions. Additionally, local operation enhances privacy and security, allowing sensitive data to remain on-site rather than being transmitted over potentially insecure networks. [Read more in this book about how to do that.](#)

Image not found or type unknown



Inspiration project: Solar Server is a solar-powered web server set up on the apartment balcony of Kara Stone to host low-carbon videogames. <https://www.solarserver.games/>

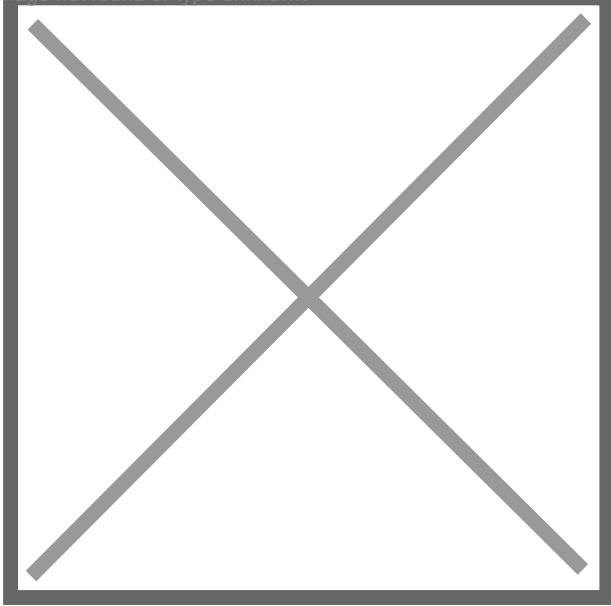
At what costs...

Artificial intelligence may invoke ideas of algorithms, data and cloud architectures, but none of that can function without the minerals and resources that build computing's core components. The mining that makes

AI is both literal and metaphorical. The new extractivism of data mining also encompasses and propels the old extractivism of traditional mining. The full stack supply chain of AI reaches into capital, labor, and Earth's resources - and from each demands enormous amounts.

Each time you use AI to generate an image, write an email, or ask a chatbot a question, it comes at a cost to the planet. The processing demands of training AI models are still an emerging area of investigation. The exact amount of energy consumption produced is unknown; that information is kept as highly guarded corporate secrets. A reason is that we don't have [standardized ways of measuring](#) the emissions AI is responsible for. But most of their carbon footprint comes from their actual use. What we know:

Image not found or type unknown



[The AI Index](#) tracks the generative AI boom, model costs, and responsible AI use. 15 Graphs That Explain the State of AI in 2024:

Usage

- Making an image with generative AI uses as much energy as [charging your phone](#). Creating text 1,000 times only uses as much energy as 16% of a full smartphone charge.
- Generating 1,000 images with a powerful AI model, such as Stable Diffusion XL, is responsible for roughly as much carbon dioxide as driving the equivalent of 4.1 miles in an average gasoline-powered car.
- A search driven by generative AI uses four to five times the energy of a conventional web search. Google [estimated](#) that an average online search used 0.3 watt-hours of electricity, [equivalent](#) to driving 0.0003 miles in a car. Today, that number is higher, because Google has integrated generative AI models into its search.
- It took over 590 million uses OF Hugging Face's multilingual AI model [BLOOM](#) to reach the carbon cost of training its biggest model. For very popular models, such as ChatGPT, it could take just a couple of weeks for such a model's usage emissions to exceed its training emissions.
- According to [some estimates](#), popular models such as ChatGPT have up to 10 million users a day, many of whom prompt the model more than once.

Build/Run/Host

- Running only a single non commercial natural language processing model produces more than 660.000 pounds of carbon dioxide emissions, the equivalent of 5 gas powered cars over their total lifetime (incl manufacturing), 125 round trips New York - Beijing. This is a minimum baseline and nothing like the commercial scale Apple and Amazon are scraping internet-wide datasets and feeding their own NLP systems.*(AI researcher Emma Strubell and her team tried to understand the carbon footprint of natural language processing in 2019)*
- ChatGPT, the chatbot created by OpenAI, is already consuming the energy of 33,000 homes. OpenAI estimates that since 2012 the amount of compute used to train a single AI model has increased by a factor of ten every year.
- Data centers are among the world's largest consumers of electricity. China's data center industry draws 73 percent of its power from coal, emitting about 99 million tons of CO₂ in 2018, and the electricity is expected to increase two-thirds by 2023.
- Roughly two weeks of training for GPT-3 consumed about **700,000 liters of freshwater**. The global AI demand is projected by 2027 to account for 4.2-6.6 billion cubic meters of water withdrawal, which is more than the total annual water withdrawal of Denmark or half of the United Kingdom.
- Water consumption in the company's data centres has increased by more than 60% in the last four years, an increase that parallels the rise of generative AI.

In-depth breakdown

The dirty work is far removed from the companies and city dwellers who profit most. Like the mining sector and data centers that are far removed from major population hubs. This contributes to our sense of the cloud being out of sight and abstracted away, when in fact it is material affecting the environment and climate in ways that are far from being fully recognized and accounted for.

Compute Maximalism

In the AI field it's standard to maximize computational cycles to improve performance, in accordance with a belief that bigger is better. The computational technique of brute force testing in AI training runs or systematically gathering more data and using more computational cycles until a better result is achieved, has driven a steep increase in energy consumption.

Due to developers repeatedly finding ways to use more chips in parallel, and being willing to pay the price of doing so. The tendency toward compute maximalism has profound ecological impacts.

[The \[uncertain\] Four Seasons](#)

The [Uncertain] Four Seasons is a global project that recomposed Vivaldi's 'The Four Seasons' using climate data for every orchestra in the world. <https://the-uncertain-four-seasons.info/project>

Consequences

Within years, large AI systems are likely to need as much energy as entire nations.

Some corporations are responding to growing alarm about the energy consumption of large scale computation, with Apple and Google claiming to be carbon neutral (meaning they offset their carbon emissions by purchasing credits) and Microsoft promising to become carbon negative by 2030.

In 2023, Montevideo residents suffering from water shortages staged a series of protests against plans to build a Google data centre. In the face of the controversy over high consumption, the PR teams of Microsoft, Meta, Amazon and Google have committed to being water positive by 2030, a commitment based on investments in closed-loop systems on the one hand, but also on the recovery of water from elsewhere to compensate for the inevitable consumption and evaporation that occurs in cooling systems.

Image not found or type unknown



Deep Down Tidal is a video essay by Tabita Rezaire weaving together cosmological, spiritual, political and technological narratives about water and its role in communication, then and now.

More about water

Reflecting on media and technology and geological processes enables us to consider the radical depletion of nonrenewable resources required to drive the technologies of the present moment. Each object in the extended network of an AI system, from network routers to batteries to data centers is built using elements that require billions of years to form inside the earth.

Water tells a story of computation's true cost. The geopolitics of water are deeply combined with the mechanisms and politics of data centers, computation, and power - in every sense.

The digital industry cannot function without generating heat. Digital content processing raises the temperature of the rooms that house server racks in data centres. To control the thermodynamic threat, data centres rely on air conditioning equipment that consumes more than 40% of the center's electricity (Weng et al., 2021). But this is not enough: as the additional power consumption required to adapt to AI generates more heat, data centers also need alternative cooling methods, such as liquid cooling systems. Servers are connected to pipes carrying cold water, which is pumped from large neighboring stations and fed back to water towers, which use large fans to dissipate the heat and suck in freshwater.

The construction of new data centers puts pressure on local water resources and adds to the problems of water scarcity caused by climate change. Droughts affect groundwater levels in particularly water-stressed areas, and conflicts between local communities and the interests of the platforms are beginning to emerge.

Curious about more? Reading tips

[Atlas of AI - Kate Crawford](#)

[A Geology of Media - Jussi Parikka](#)

[Hyper objects - Timothy Morton](#)

Sources:

[Atlas of AI - Kate Crawford](#)

[Podcast: Kunstmatig #28 - Tussen zeespiegel en smartphone](#)

[Datacenters and water: KUER](#)

[Technology Review: Making an Image with Gen AI uses as much energy as charging your phone](#)

<https://www.washingtonpost.com/technology/2023/06/05/chatgpt-hidden-cost-gpu-compute/>

[https://arxiv.org/pdf/2206.05229](https://arxiv.org/pdf/2206.05229.pdf)

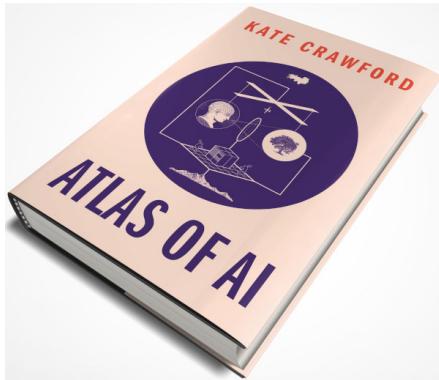
[Technology Review: getting a better idea of gen AIs footprint](#)

<https://spectrum.ieee.org/ai-index-2024>

[Standford EDU: Measuring trends in AI](#)

Boeken

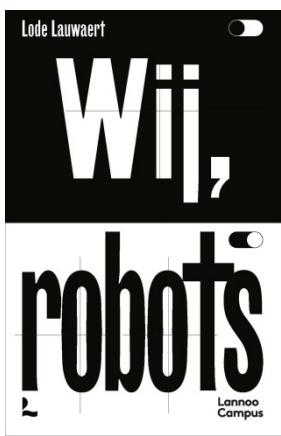
Atlas of AI - Kate Crawford



A history of the data on which machine-learning systems are trained through a number of dehumanizing extractive practices, of which most of us are unaware.

*In *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* Crawford reveals how the global networks underpinning AI technology are damaging the environment, entrenching inequality, and fueling a shift toward undemocratic governance. She takes us on a journey through the mining sites, factories, and vast data collections needed to make AI "work" — powerfully revealing where they are failing us and what should be done.*

Wij robots - Lode Lauwaert

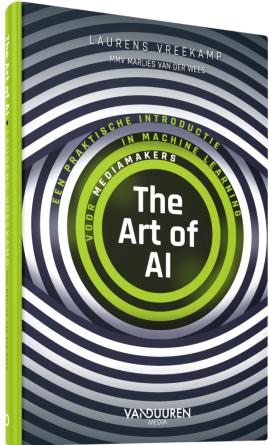


A philosophical view on technology and AI.

*When discussing technology and AI, the image often arises of a future with malevolent, hyper-intelligent systems dominating humanity. *Wij, robots*, however, focuses on what is happening right before our eyes: (self-driving) cars, smartphones, apps, steam engines, nuclear power plants, computers, and all the other machines that we surround ourselves with.*

This book raises fundamental questions about the impact of new and old technology. Is technology neutral? Are we sufficiently aware that Big Tech knows our sexual orientation, philosophical preferences, and emotions? And since the digital revolution, is the world governed by engineers, or are their inventions merely a product of society?

Art of AI - Laurens Vreekamp

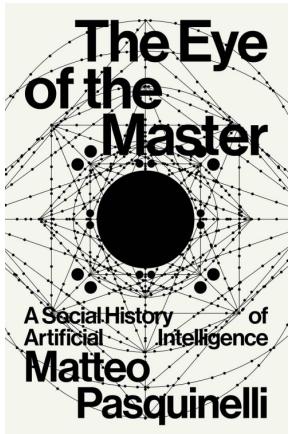


A practical look at how the creative process and AI go hand in hand, outlining the opportunities and dangers of artificial intelligence.

Automatically searching images in your own archive. Editing podcasts as if you were modifying a text in Word. Generating illustrations with text, coming up with alternative headlines for your article, or creating automatic transcriptions of dozens of audio recordings simultaneously; the possibilities of AI are endless.

In *The Art of AI*, authors Laurens Vreekamp and Marlies van der Wees take a practical look at how the creative process and AI go hand in hand, outlining the opportunities and dangers of artificial intelligence. This book is accessible even if you have no programming knowledge, mathematical aptitude, or a preference for statistics.

The Eye of the Master - Matteo Pasquinelli



A social history of AI that finally reveals its roots in the spatial computation of industrial factories and the surveillance of collective behaviour.

What is AI? A dominant view describes it as the quest "to solve intelligence" - a solution supposedly to be found in the secret logic of the mind or in the deep physiology of the brain, such as in its complex neural networks. The Eye of the Master argues, to the contrary, that the inner code of AI is shaped not by the imitation of biological intelligence, but the intelligence of labour and social relations, as it is found in Babbage's "calculating engines" of the industrial age as well as in the recent algorithms for image recognition and surveillance.