

AI is coming — who is ready?

“Technology is neither good nor bad; nor is it neutral.

Historian [Melvin Kranzberg, 1986](#).

The EU needs to revise its legislation on artificial intelligence

[Initiative Urheberrecht \(Authors’ Rights Initiative\)](#) is a consortium of 42 German creative industries organisations. In a recent statement, they called on the EU to update its legislation on AI: [Authors and Performers Call for Safeguards Around Generative AI in the European AI Act \(19/4/2023\)](#).

The organisations are concerned about the impact of *generative AI* on communication on the one hand, and on artists’ rights on the other.

Generative artificial intelligence differs from other types of AI in that it actually is able to generate new outputs. It does not just compile statistics, classify and perform user-defined calculations, but it also independently creates new texts or other types of content on the basis of the data given to it.

It can use everything published on the internet — and any other digital content available to it — as raw material that it can absorb in virtually unlimited quantities.

What harm could this do? Well, if things go wrong, our world will be filled with fake texts, images, videos and chatbots generated by artificial intelligence, and we will no longer be able to tell these from real news, documentary material, speeches made by real people — or even from real people.

There is also a risk that concerns artists and other content creators that applications that use generative AI will absorb all human-generated digital or digitised material and will be able to churn out an unlimited number of adaptations and versions from the material. And since there are no direct quotes in the products made by artificial intelligence, its outputs are not governed by the current copyright legislation. Unless we significantly change the way we interpret the law.

If this is the case, the market and the earning potential may be almost exclusively in the hands of those operators who have access to the most efficient data harvesting methods and the most powerful AI applications.

It is no coincidence that man-made content is not referred to as raw material in generative AI parlance, but as *training material*. This gives the impression that AI does not actually take in or copy anything man-made but is only interested in watching what it does and uses it for training, a bit like a singer who listens to someone else singing in order to learn how to sing. This impression is misleading.

We may find ourselves in a situation in the future where people will not pay for any content, whether created by humans or artificial intelligence; we'll only pay for the use of content-generating applications instead. This could mean, for example, that large streaming services that dominate the market would not pay song makers even the pittance they do now but would start distributing robotic music produced by their own AI applications, customised for each group of listeners.

The EU completed its AI package in 2021 ([Artificial Intelligence Act / AI Act](#)). Its main purpose was to make the use of artificial intelligence safe. However, generative AI has been evolving so quickly lately that the AI Act is already outdated. The EU is still in the process of adopting the legislative package.

German organisations are now seeking to influence the form of legislation and are calling for a review and update.

The statement calls for the EU to regulate the use of generative AI throughout the production and consumption chain. The chain begins with the collection of data and ends with citizens consuming content that is created using artificial intelligence. The EU must pay particular attention to how the 'engine room' of AI applications entering the market works. Here, 'engine room' refers to the foundation models on which the learning ability of various generative applications is based.

The risks of generative AI may seem distant, but they are already real. What should be done to minimise these risks? It is not clear, and that is why the list of concrete demands compiled by the German organisations is excellent. I think that every citizen, artist and creative industry organisation should agree with those demands.

Everyone should read the statement written by the Germans. Here are my thoughts inspired by the statement and some other sources. **My suggestion for the rules for the use of generative AI would be:**

- 1. Author permission is required for the collection of content.** The bodies that use generative AI should not simply copy all digital content published by other people to feed their own apparatus, as they do now. I'm a writer and I can read and be inspired by as many books as I can get my hands on; I can use them as mental nourishment and as raw material for my own texts. This is fine as long as I do not copy verbatim or plagiarise other people's works. However, AI should not be given this right. Why? One reason is that AI is not just 'impressed' by a book that it 'reads', but copies it, word for word, into its memory. People who write books or create other content must also be able to prohibit the use of their works as raw material for generative AI. If authors do not wish to prohibit the use of their works for this purpose, they — or

organisations that represent them — must be allowed to negotiate the nature and extent of the use.

2. **Creators of content used by generative AI must be remunerated.** When authors negotiate the use of content they have created as raw material (or training material) for generative AI, they can demand financial compensation for the use of their creation.
3. **Content generated by generative AI must always be labelled in such a way that it is not accidentally mistaken for something made by a human — or for a human themselves.** AI can produce content that looks like news or research papers written by a human, opinion pieces and social media updates — or collections of poems — written by people. Audiences have the right to know whether the piece that they have read, heard or seen was created by a human or not. If an AI application has been used, consumers must be able to understand the role and extent to which the application has been used. A practical example: it is important for audiences to know whether the lyrics of a polemical song they hear were written by artificial intelligence. They should also be told if the lyrics were written by a human, but the song composed by AI.
4. **The mechanism of how generative AI works must be transparent to ensure that it is not only the owners and developers of AI who can monitor and assess aspects such as AI's ability to avoid producing lies.** This is obvious when AI diagnoses a patient's illness or assesses the likelihood of a convicted person to reoffend, because in such cases we cannot rely on 'the machine knowing best'. This same principle of transparency must, however, also apply to all other uses of AI.
5. **Companies that develop generative AI cannot shirk their responsibility for how their product is used.** At the moment, they are doing just that. Particularly bodies that develop the foundation model of generative AI are of the opinion that they cannot be held responsible if someone uses the tools they provide to deliberately generate disinformation, to steal content created by other people or to try to manipulate voters using dishonest means, for example.

6. **Developers of generative AI are, for their part, responsible for accidental errors made by the service or tool that they created, such as lies and hate speech created by AI.** It is not enough for providers of AI services to simply advise users to be cautious and to urge them to carefully check the AI outputs for possible errors or harmful content.
7. **There must be a firewall between companies that develop generative AI and companies that distribute content.** Social media companies and music streaming services, for example, should not be allowed to produce content themselves using generative AI to fill their platforms and services. Why? Because otherwise they will gain too dominant a market position and have too much power to control the consumption habits of the people who use their services and too much power to shape the way users see the world.
8. **Generative AI must not be allowed to be used to produce content that is intended to blatantly mock or vilify people or content that people create.** Why? Because, otherwise, in the future — or even now — it will be possible to commission AI to, say, spread hate speech endlessly in the form of stories, a sort of journalism, music, images, social media posts and similar.
9. **It is also important from the EU's perspective that a significant portion of the development of generative AI will be carried out in the EU and that the EU also has the required infrastructure.** For the sake of the EU's security and cultural relevance it needs to be adequately self-sufficient in this respect, too.

These were my main points. No need to read any further. Unless you are interested in my more detailed reflections on the impact of AI, especially generative AI, on humankind.

A giant leap for stupid AI

AI has recently made such advances that what seemed out-there science fiction yesterday is now available to everyone, often for free.

Online search engines are, in a way, precursors to generative AI. Search engines collect, or 'mine', data tirelessly, organising it using various classification principles so that the material required can be found as quickly as possible. Generative AI goes one step further: it is able to independently imitate and recombine all the content, the data mined, that it has collected in its memory.

Generative AI is now capable of simulating works created by humans so well that it is difficult, if not impossible, to tell the creations produced by AI from ones made by humans.

It is not that machines or clusters of bits have learned to think — at least in the sense of the word that we are familiar with — nor have they developed consciousness. On the face of it, however, some of them have learned to imitate our activities astonishingly well. One of the foundation models of generative AI is the GPT (Generative Pre-trained Transformer) language model, of which there are several versions. Applications that use GPT language models work in such a way that when they are fed with a large amount of human-generated texts, images, music or other material, they are able to produce approximately the same, yet new, material on the basis of what they have learnt. The basic idea is that when a well-fed application is given, say, a sentence or even just one word, the material it has been fed with helps it to guess what word a person is likely to type in next. The application does not really operate on the meanings of words, but rather on statistical probabilities, on paths and patterns of human activity.

The fact that they are so 'stupid' is exactly the reason why it has been possible to develop GPT-based applications so far so quickly. There is no need for them

to learn or even imitate human reasoning, problem-solving skills, value judgements or consciousness. They just need to calculate probabilities and look for correlations. They do not look for cause and effect relationships, for example, as this would require understanding. They are happy with a simple correlation: “this word often comes after that one.”

I asked an early GPT-based application: “What is a snake?” The first answer was, quite correctly, that a snake is a legless reptile. When I repeated the question, the application said that a snake is an animal that has sinned and had its legs removed by God as a consequence. Source criticism is not generative AI’s strong point.

Applications have since evolved, but are not necessarily any more reliable. They have a better answer to the question about a snake but if faced with a more complex question or assignment, they are just as unreliable as before. Or even less reliable, as AI is able to generate content such as a scientific article or news-sounding piece and supplement it with fake references that it came up with. AI does not understand the definition of a snake, for example. It merely imitates human-generated phrases in which the word ‘snake’ appears and images that have the word ‘snake’ in the caption. For this reason, the quality of the material it produces depends on both the quality of the food it is fed and the angle it is told to take.

Despite its shortcomings, generative AI is incredibly efficient in some ways: it is able to absorb an unlimited number of texts and other material in its own superficial way, and it is able to remember them all, character by character, pixel by pixel. The essence of generative AI, the foundation model, does not resemble an individual who is learning but a smoothly functioning swarm intelligence: when it learns something somewhere, it is immediately available to all users of the foundation model everywhere.

The unreliability of generative AI is not a huge problem when it is given a task such as writing bedtime stories. It can easily make up stories, and if the programme is used or supervised by humans, it is their responsibility to assess the stories' moral impact and other effects on children before children are exposed to them. Supervisors can sift out any harmful stories or modify them to suit their world view.

Generative AI is, however, most often used without any responsible supervisor. This is understandable: if I were to create works of art using generative AI, I would not want to allow the AI application I use or the works I make with it to be censored in any way in advance. But I would like to know what kind of hidden agendas or ideological tendencies the application that I use comes with.

Although generative AI is not very clever when it comes to facts and reasoning, it is already being used to write news, diagnoses and court decisions, for example. In these cases it is assisted by an individual who evaluates the meaningfulness and accuracy of the text generated and usually also other forms of AI, those that are not capable of creative activities but can perform complex calculations and are also capable of some kind of logical reasoning.

Despite all this, it is clear that the use of generative AI entails enormous risks. Generative AI is like an incredibly intelligent and charismatic person who has no grasp on reality but who can speak with great eloquence and conviction. And never shuts up.

The threat of truly intelligent AI

Will there come a time when artificial intelligence obtains consciousness and a will of its own, as if out of nowhere, automatically, as its development continues to progress? GPT-based applications may not be able to achieve this, no matter how much computing power and speed they have, as they are, at the moment,

mainly just bullshit generators that mimic human activity and thinking astonishingly well.

However, we do not know what they may develop into. Elsewhere, artificial general intelligence (AGI), 'truly intelligent artificial intelligence', is being developed, which would have actual cognitive abilities.

Psychologist and pioneer of computer science **Geoffrey Hinton**, often referred to as the "godfather of AI," recently (on 3 May 2023) said in an interview with the BBC that AI is already capable of absorbing much more information than a person, and although it is not as good at reasoning as a person, it is evolving very quickly so [we do need to worry](#):

"Right now, they're not more intelligent than us, as far as I can tell.
But I think they soon may be.

Hinton describes this possible development with the following scenario. What if an authoritarian leader like **Vladimir Putin** gave AI the ability to independently create sub-goals that served his grand goal of, say, winning the war? AI could very well end up assessing the situation in such a way that "in order to achieve the goal given to me, I need to get more power". This scenario sounds like a bad sci-fi movie to me, but if [one of the most important developers of neural networks](#) thinks that the risk is real, I should perhaps prick up my ears.

Philosophers argue among themselves about the probabilities of these speculations. So do engineers. There is no consensus on this matter. Which makes the situation so thrilling. We do not know at this point what the next generation applications will be like. [Maybe they will catch us by surprise, jump out of a bush of bits and remove us from the range of species.](#) Or put a leash around our necks. If this happens, will we even notice? How will our view of human beings change? How will our views of machines and applications change? And when will we turn into cyborgs with mechanical parts that are a bit

more amazing than artificial joints and contact lenses? I've been interested in the topic, so much so that I wrote a [chamber opera](#) about it (**Max Savikangas** composed the music, I wrote the libretto and directed the production).

Even if truly intelligent AI never obtains consciousness and starts to unleash itself from human control, it will soon revolutionise our society in various ways. Mass layoffs will be one consequence.

Advances in technology have always destroyed jobs or liberated people from the yoke of work, while new jobs have always been created elsewhere. Many economists still think that this is a kind of law of nature: when a couple of ditch digger jobs disappear, a job is created for a software engineer and another for a masseur. This assumes that even if technological advances are destroying jobs, the ever-expanding and increasingly complex technology requires an ever-increasing number of people to maintain and develop technological systems. Another cornerstone of this assumption is that the loss of jobs caused by technological advances will go hand in hand with overall growth of prosperity, and people who become wealthier will continue to have new needs and desires, and to satisfy these, there will always be a need for new types of human labour.

I don't believe a word of this. Although there is a shortage of labour in many sectors in Finland right now, which should be rectified by increasing opportunities for education and allowing more immigrants into the country, jobs will probably disappear at an accelerating rate in the long term. Why? First of all, why would people's capacity to consume always increase at the same rate as workforce productivity? Secondly, although advances in technology do not just destroy jobs but also always create new ones, new jobs will be increasingly demanding, because they will require skills that machines do not have yet. The more capable machines become, the harder it becomes for human workers to compete with them. There will be plenty of demanding jobs in the development of AI and other complex technologies for a small number of nerds. Everybody

else will be delivering pizzas to the well-to-do on electric bikes — until the bikes learn how to do the job by themselves.

This is one of the great challenges that truly intelligent AI will be throwing in the face of our civilization. Mind you, this new kind of mass unemployment and inequality cannot be remedied by the old methods alone. I believe that a basic income must be introduced that is sufficient to cover the cost of living, and some of the most advanced technologies must either be taken into joint ownership — ‘nationalised’, as it used to be called — or they must be taxed so effectively that the welfare state will not collapse but will be able to support all the ‘surplus people’ that AI and other similar technological advancements render superfluous *to the economy*.

It is quite possible that we will soon find ourselves in a situation in which a handful of the world’s richest people own — through their corporations — most of the technologies, including AI applications, that provide the digital infrastructure we need. They may also come to own, on account of AI, an increasing part of the workforce, i.e. the machines that are able to perform many tasks better than human beings. What will then happen to non-wealthy people? Do we give in to the idea that all those technologies are the private property of very few people because they have earned it through their efforts and ingenuity? Will we accept the concept that they just happen to have legitimate power over the rest of us? Or will we be calling for some kind of radical redistribution? Will we be able to create a society in which technology truly liberates us from the yoke of irksome jobs, the rat race of production and consumption, and the increasingly fierce competition on the labour market?

The outlook for the future seems to be somewhat bleak. There will be major conflicts at least. Yet I, for one, find all of this fascinating philosophically as it opens up new ways of reflecting on what thinking, intelligence, consciousness and creativity really are. Would it already be appropriate to call what AI does thinking, problem solving, intelligence, even if the applications have no

consciousness or will of their own? Applications do, after all, already replace human beings in decision-making: AI can diagnose illnesses and make decisions on investments better and quicker than humans, for example. In its current form, generative AI does not understand meanings or cause and effect relationships, but some other forms of artificial intelligence (AGI) are already better at reasoning than humans in tasks that are sufficiently narrowly defined. And how will we react if AI develops self-awareness and tells us about it? Will we try to control it or grant it human rights?

What impact will generative AI have in society now and in the near future?

I'll come back to this point. I'll forget the challenges that truly intelligent artificial intelligence poses and consider how generative AI will affect our everyday lives now and in the near future.

When pupils are given the task of writing an essay on a particular topic, they may be tempted to have AI write it for them. [ChatGPT](#) is an example of an app that can generate text on any topic, at least in English. The free version of the app is available to anyone online. If the text contains any peculiarities that humans would not write, pupils may well know how to revise it and their teacher will be none the wiser. Pupils may well be able to polish their text without understanding the content of the AI-generated essay.

Of course, AI that simulates and anticipates human creative activities and choices is already a part of our lives in an even more inconspicuous way: various algorithms constantly make choices for us, showing us content, advertisements and alternatives that the algorithm thinks are suitable for us. Well, suitable for us and especially suitable for the body that manages the algorithm.

Algorithms are often used in profiling carried out by AI, especially for targeting ads to particular recipients. On the face of it, this should be fine. If advertising is allowed, is it not actually a good thing that ads for dark chocolate, electric bicycles or atonal choral music are specifically targeted at those who might be keen, rather than pushing them on everyone?

It's not that simple. The [Facebook–Cambridge Analytica data scandal](#) is a case in point. A company called Cambridge Analytica collected Facebook users' personal data without the users' permission and helped the data to be used to influence in political contexts such as **Donald Trump's** 2016 presidential campaign in the United States. The collection of data without permission was in itself harmful as were the methods of political influence applied to manipulate people to vote for Trump. Cambridge Analytica sifted through Facebook users for potential Trump voters, who were then targeted with customised political campaign material. It was not simply a case of sending targeted campaign ads, but of also changing the digital landscape these Facebook users saw: controlling what went on their feed, which news and ads they saw, which fake sites, fake threads, fake users and fake communities surrounded them in the digital world. The campaign also involved creating fake websites, fake users and similar fake features for this very purpose. This case was a cautionary tale of how data collection, disinformation and the creation of bogus users can be an effective brainwashing method, especially when these elements are implemented at the same time.

In what ways, then, does generative AI relate to the use of algorithms for profiling and manipulating consumers and voters and for spreading lies?

There are two ways. Firstly, AI makes data collection and people profiling much more efficient.

Secondly, AI makes creating and spreading lies incredibly easy, cheap, and efficient. Instead of people having to type disinformation on their keyboard, AI

can be used to draft an unlimited amount of fake news and racist statements supposedly written by private individuals. It can create an unbelievable number of fake accounts for social media platforms that have fluent conversations with real people.

AI technology is also persuasive technology. One of the trending terms in marketing is '[targeted conversational influence](#)'. This refers to attempts to influence individual consumers through communication that A) is specifically customised to appeal to the personality, worldview and consumption habits of the individual, and B) creates an interactive, conversational relationship with the consumer, i.e. reacts to their choices in real time and learns more about the individual all the time. Artificial intelligence can do this better than any human can. Generative AI can, therefore, be an effective tool for manipulation even when it sticks to facts, as long as it selects which facts and which opinions of other people it makes visible to the individuals — especially if it seasons them with emotional comments from virtual characters created by AI itself.

This is a good time to remind ourselves of how the fake news spread via WhatsApp [has fuelled hatred between Hindus and Muslims in India](#) and led to lynchings and massacres in recent years. There are political reasons, but the phenomenon has also been caused by the rapid spread of mobile phones and [very poor media literacy](#). It is easy for us to be smug and think that such things only happen to other people in other countries, not us, who credit ourselves with a very high level of media literacy and understanding. But we are not safe, because generative AI can also alter our mediascape so that our literacy skills and understanding deteriorate. Perhaps we, too, are gullible but with regard to other things than the Indians who were tricked by rumours and fake news into killing their neighbours.

If people, no matter how well educated or media literate, cannot tell content produced by AI, be it poetry or fake news, from creative works or accurate

information provided by human beings, how can they, as citizens and consumers, protect themselves from the avalanche of lies and manipulation?

The Future of Life Institute published an open letter in March 2023 [calling for a pause on AI development](#). The letter attracted a lot of attention, because there are a lot of famous people amongst the signatories and the subject is exciting. [Some experts have been critical of the letter](#), but I have not seen anyone denying the letter's core message: a set of shared and transparent safety protocols needs to be developed to prevent people from using AI for destructive purposes — either intentionally or unintentionally:

“The letter does not claim that GPT-4 will become autonomous – which would be technically wrong – and threaten humanity. Instead, what is very dangerous – and likely – is what humans with bad intentions or who are simply unaware of the consequences of their actions could do with these tools and their descendants in the coming years.”

/

[Joshua Bengio, 5/4/2023.](#)

Generative AI is not evil in itself — no technology is good or evil in itself — but neither is it neutral. People should consider the risks of each piece of technology separately. And it is not enough that ‘enlightened citizens’ think about the risks and improve their skills in source criticism and literacy in visual and other types of content. Individuals must join other citizens and demand, before it is too late, that the state, the EU and the international community create rules for the development and use of generative and other kinds of artificial intelligence.

What impact will generative AI have on art?

A professional writer can ask an AI application to produce a hundred poems on a given topic, using styles and references chosen by the author. The writer can use these machine-generated products either as they are or — probably more likely at least for the time being — as a basis for their writing, as first drafts for something that they can reform using their brain and bodily experiences.

Paula Havaste, Vice-chair of the Union of Finnish Writers, recently wrote in a letter to the Union’s members about artificial intelligence and art, stating, among other things:

“But how can we ultimately tell a work made by artificial intelligence from a work written by an author?”

The answer is quality. A human writer can combine different parts in surprising and creative ways. This creates content that is much more diverse, challenging and profound than artificial intelligence could generate.

[...]

And AI does not do anything on its own. It requires a specific, cleverly formulated question, wish or command, which must have a clever approach to the topic. That’s why I’m not terribly worried. AI is a great tool, a bit like a knife. A knife can be useful and handy, but it can also be dangerous in the wrong hands.”

A letter to the members of the Union of Finnish Writers, 2 May 2023.
“Greetings from the Vice-chair: A knife and AI.”

AI can also generate music and images. Earlier this year, German artist **Boris Eldagsen** [won a prize with an image that looked like a photograph but was generated by AI](#). Eldagsen refused to accept the award, saying that his work was intended to provoke debate about the nature, definition and future of photography.

I think that most artists are curious, rather than worried, about the advent of AI. It's thought to be a bit like a camera, a synthesizer or the numerous programmes that mimic actual instruments and are now used to create most electronic music. They are just brainless, passive machines. They do absolutely nothing without the artist's desire, insight and choices. Users of a good piece of music making software will never have time to familiarise themselves with all the samples, loops and other effects that come with the application, but the programme still does not make music by itself. However, generative AI is fundamentally different from all previous tools for artistic creation. It really does make music if it is told to do so. It does not have a will of its own — at least not yet — but if you give it, say, five loaves and two fish — oops! five words and two chords — it can use those to generate a three-minute audio file, which will be impossible to distinguish from a three-minute recording by one person singing and a couple of others playing real instruments.

On 4 April 2023, a person using handle **Ghostwriter977** released a song called [Heart on My Sleeve](#) on several streaming platforms. The song is a [fictional duet](#) by rapper **Drake** and singer **The Weeknd**. Drake and The Weeknd are popular Canadian musicians, but they had nothing to do with the song. It was Ghostwriter977 who made AI invent the lyrics and simulate what Drake and The Weeknd's duet might sound like. The song became a hit for a while, but Drake and The Weeknd's record label, Universal Music Group, demanded that the streaming platforms take the song down, citing copyright as the reason. They did, but the song was immediately back in distribution, without permission. It is still not clear how copyright laws should be applied to products generated by artificial intelligence. "Heart on My Sleeve" was well suited to stir up this debate

as it is a weird borderline case: it refers to two existing people who make music, imitates their sound and sort of exploits their work but does not directly borrow any rhymes, melodies or beats. The song is, of course, popular because it refers to two popular people who make music, not because of the song itself. However, this case is a concrete example of how clever AI can already be at artistic tasks.

It would be sad if AI, especially generative AI, were to be used only to create slightly modified versions of human-generated hit products; AI would simply be a way to circumvent copyright laws and make artists unemployed. Also, it would not produce anything really new but would try to milk old, familiar and popular works endlessly for the profits.

Fortunately, many artists try to achieve something [completely different](#) when they use AI. They use or work with AI to make works of art that could not be done without AI and that somehow sound or look truly new.

Composer, musician and artist [Holly Herndon](#) is a great example. She has created works of art using both traditional methods as well as AI-based applications. She has done baffling things with voice editing and generative and real-time audio production. The easiest way to grasp the idea of how Herndon uses AI is [Holly+](#). It is simply a programme and service that allows you to make the virtual Holly sing anything, either from a recording or in response to live sound. Holly+ is a compact version of the larger AI software Spawn, which Herndon developed with software developer **Jules LaPlace**. Spawn learns people's singing voices and techniques and is able to create new music based on them, either under the strict guidance of the user or rather independently. I find the outputs interesting and sometimes even great, and they never sound like imitations of existing music. I can highly recommend Herndon's album [Proto](#) (2019), for example. An example of a simple use of Holly+ is Herndon's version of Dolly Parton's [Jolene](#) (2022).

Together with **Mathew Dryhurst**, Herndon has also developed an online service [Have I been Trained?](#), which allows artists to try to find out if their works have been used to train generative AI. The service does not cover all works but at least it shows if an artist's images or texts are included in the generative AI dataset [Laion-5b](#), which is the largest of its kind in the world. I put the site to the test. It showed that dozens of my paintings and photographic works have been used to train generative AI. I also found numerous images, digital paintings, which were clearly generated using my works as one of the raw materials.

Artists can also opt out from the Laion dataset on the Have I Been Trained? website. It's great that one artist and a couple of developers have managed to create such a control tool, but it is unfortunate that governments and organisations have not created a similar, more comprehensive system and made its use mandatory. Have I Been Trained? is only a very small, single step forward. Laion-5b is just one of the datasets used to train generative AI, and while Laion is a non-profit organisation that aims ["to make large-scale machine learning models, datasets and related code available to the general public"](#), it alone is not a solution to AI-related copyright issues and for human artists to receive financial compensation for their works.

Meanwhile, something completely different is happening elsewhere, also based on AI, but much less artistic. For example, AI-based services have emerged that can be used to compose background music for films, ads, podcasts and similar media. The service provider makes money while the service user obtains background music for less money than if they used music made by human beings. And the works by AI are 'unique', tailor-made especially for a particular assignment. I think the outputs are terrible but no worse than the average elevator music, which is usually either music made by musicians under their own names or anonymous ["stock music"](#).

What should we think about all this? It would be ridiculous to say that lift music generated by AI is automatically wrong, somehow; an affront to humanity or to

musicians. AI taking the job of a jobbing bassist is similar to tractors, combine harvesters and other machines replacing most farmhands.

Some people see no need for any restrictions. The current copyright laws prohibit direct copying and plagiarism, but people are allowed to take inspiration from all kinds of influences, which they can then make their own. I have been inspired by a number of my favourite painters, favourite poets, favourite directors, and I have created works as a painter, writer and director from which an expert could immediately deduce that “Teemu clearly likes **Bertolt Brecht’s** plays and poems, **Otto Dix’s** paintings as well as films by **Agnes Varda** and **Pier Paolo Pasolini**, and now he’s doing something similar, mixing influences from various sources with an odd Lapua twist”. Why shouldn’t AI do the same? Or why shouldn’t an individual or a corporation generate artistic products using AI and the stylistic choices mentioned above and put them on the market? Then just wait to see if the demand is there?

It’s not that simple, however. The AI applications we are talking about here need a huge amount of cultural training material created by humans, otherwise they would not be able to guess which word follows another word or what a three-minute rap song might be like; they would have nothing on which to base their simulations and versions. When AI absorbs all the music made by humans and then pops out an endless number of new songs, the situation is quite different from a tractor replacing a hundred men with shovels. A tractor does not make use of the tricks of the shovelling trade, nor does it use and anonymously take credit for the intellectual capital and creative work of the workers.

I do not think it is a good idea to give AI, or an individual or corporation using AI, unlimited access to all human-generated content that it wants to use as raw material, or ‘training material’, to produce its own content. Even if AI does not directly plagiarise or steal storylines written by human beings, faces painted by human-beings, bass riffs composed by human beings, verses drafted by human beings, even if it strictly avoids

appropriating any material prohibited under current copyright laws, it still creates something new from human-generated cultural material. People must be allowed to decide how much of their creative output can be used as food for AI and the remuneration they receive for such use.

I can hear someone asking: “What’s the harm, really, if AI is allowed to freely use as its raw material any content created and published by human beings?” The harm would be the unfairness of it all: you write a poem and as soon as it is published, AI generates a thousand poems based on your poem and poems by a couple of other writers, yet the poems written by AI are so different from the human-written originals that the application cannot be accused of plagiarism. The AI owner would have the control and earning potential. And while ChatGPT, for example, is a free online service now, the industry’s long-term goal is, of course, to make people pay for it — and the more advanced the AI application, the more people will have to pay for it.

You could not even make a spoon without help. Even if you did not know what a spoon was, you might come up with the concept of spoon, but if you did not have access to certain tools developed by other people, such as a knife, it would be unlikely that you could carve a wooden stick into a utensil with a bowl at one end that would fit into your mouth. The bowl is probably the most difficult part to carve, no matter how carefully you try to scrape it with a sharp stone. In this fundamental sense, all the achievements that humankind has accomplished are shared achievements, from the spoon and taxation to the string quartets of **Dmitri Shostakovich**. In most cases, we do not need to pay for innovations made by others, and copyright laws only protect outputs, not ideas, and I’m happy with that. We can get inspired by anything and everything, and we are allowed to use all the information and other material we have acquired in our own creative activities free of charge. Generative AI — and, of course, AI in general — is fundamentally changing this situation, requiring new solutions and a new kind of fairness.

If generative AI were equally easily available to all people for free, an anarchist might think that there is no need for any controls or restrictions related to the rights of human artists. If everyone could use AI to generate whatever content they want, people would be completely equal in this regard. There would be those moaning and asking how human artists are supposed to support themselves if everything they do in the future will be free for all to exploit? The anarchist would say that perhaps motivations other than money should be behind art-making, anyway, and, besides, AI will probably soon learn to use works it has generated itself and to create works based on materials that are not human-generated, so what's the problem here?

I disagree with the anarchist of my imagination. One of the reasons for this is that it would be a mistake to think that AI and human beings absorb knowledge and influences in the same way and equally efficiently. And the anarchist's dream of an AI that would be equally accessible to everyone, like breathable air, will never come true. (Granted, not everyone has access to equally breathable air, but that's another story.) It is more likely, as I said above, that a very limited number of bodies will own the most powerful versions of generative AI, and it will be only in the interest of those few to obtain the raw material required in their production processes free of charge. Journalist, author and professor **Naomi Klein** [described this problem in a recent article](#), saying that we are perhaps witnessing the largest theft of all time, in which a number of the world's wealthiest companies (Microsoft, Apple, Google, Meta, Amazon, etc.) are trying to seize all of the human knowledge ever produced that exists in digital form to turn it into proprietary products and services. Continuing Klein's line of thought, one might think that those who are spreading the joyful news of generative AI are surreptitiously trying to turn us from creators into consumers who are happy to pay for content they created themselves in the first place.

Not everyone recognises this danger. Paula Havaste, for example, drew attention to the point that an application only generates something when it is given an assignment, and she said that she believed that the quality of human-

made works would continue to stand out from AI-generated outputs, and for that reason she is not overly concerned. I think she is mistaken.

On the face of it, you might think that AI is only capable of generating something interesting or sellable when a human being gives it an insightful question or task. A clever briefing certainly helps, but it is not necessary since AI is able — if not right now, then probably very soon — to generate a million outputs in a few seconds even on the basis of very rudimentary instructions: “Write a novel about refugees in a style that is a combination of [F. E. Sillanpää](#) and [Kathy Acker](#).” I use Sillanpää and Acker in my imaginary assignment as I’d love to know what AI would come up with if it merged the voices and worldviews of a Finnish Nobel Prize-winning author, born in 1888, who specialised in rural life and nature, and a postmodern New York City punk-rocker lesbian, born in 1947.

A million — or even a thousand — pieces generated under those instructions would present a challenge: who will go through them, who will select, edit and polish the one that will be published? A human being or AI, and what would the job title be: author or editor? It is quite likely that the job descriptions of editors, artists, curators and similar professions will overlap. When AI produces a pile of half-finished prototypes or completed works, and a person with specialist artistic expertise wades through the stack and selects and, if necessary, edits those that are launched on the world markets with a great deal of fanfare, it is hard to say when and to what extent it is reasonable to describe the job as art-making and when to describe it as editing or curating.

The United States Copyright Office has recently announced that works made by AI will also be protected by copyright, as long as a human author has been involved in making the work ([The Human Authorship Requirement](#)). However, since AI applications are used by humans, it may be open to interpretation as to whether a particular work can be considered as having been *entirely* made by

AI and left without any copyright protection. It is also the Copyright Office's view that copyright law *does not* protect those parts of a work that AI has created independently. But how is it possible to determine AI's independent role in the process and the output?

What does this mean in practical terms? One dystopian possibility is for the largest corporations to acquire the most powerful AI applications that are suitable for the creation of artistic content and to fill the market with their outputs, thus supplanting the people in this industry altogether as a consequence — or rather marginalising the people in the fields that were previously considered to be activities particularly characteristic of humans. Streaming services are now the most popular channels for listening to music. These companies are extremely profitable but pay music-makers infuriatingly [small fees](#). A million streams generate an income of about EUR [500](#) for a musician, provided that they are the sole author of the song. In principle, it would be even more profitable for Spotify and other such services to own an AI application that would generate the content for distribution. The platform would not be just a platform but also a primary sector producer, a processor and the holder of all rights. And if music generated by AI were at risk of not being protected by copyright as outlined by the U.S. Copyright Office, the platform could simply involve enough human input to 'reach the threshold of originality' — but the rights would still remain with the corporation.

I and musicians think that this would be a terrible outcome, but not everyone agrees. Artists should realise that there are people who will embrace the surge of AI-generated works even if it displaces art made by humans. Robotic artists will have fans just like human artists.

The 'Yellow Library', published by Tammi, is the largest series of translated fiction published in Finland. Its motto is: "The best literature from across the world since 1954." [532 works](#) have been published in the series, including by thirty Nobel Prize-winning authors. I would be very interested to see what kind

of novel a clever AI application would come up with if it was trained with all those books and given the task of writing a couple of new ones, suitably limited in terms of the topic, style, era, worldview and sex. Although it is obvious that generative AI does not really understand any philosophies, it might still be able to sift through the mass of data to find some external characteristics of a wide range of views and then get down to work as instructed.

So let's imagine that the publishing house Tammi, which is owned by WSOY, which is owned by Bonnier Books — owned, sometime in the future, by an even larger organisation that publishes and distributes content — started commissioning AI to write new works for the Yellow Library series in the near or distant future. So what? I would be very curious to read at least one such novel. But maybe no more.

Why? Because art is important to me mainly as an especially versatile and flexible vehicle for philosophical reflection. By making and consuming art, we muse on questions such as: "What is the world really like? What could it be like? What would I like it to be like in the future? What am I like and what do I want to be like? How should I live? What kind of society would be a better society? How is it created? What is justice? Why do I want to continue living? What is a good life?"

A work created by generative AI would not reflect on these types of issues, it would not try to ask good questions nor would it offer even tentative answers. It is just a bunch of stimuli that people either find fascinating or not.

However, this lack of ideological content and aspiration does not bother all consumers of art. Many people think that art is just pure entertainment, something that you can look at for a little longer than a lava lamp, something that is not expected to carry any wisdom or other meaningful content. They do not automatically see any difference between art made by AI and art made by humans, no matter how much makers of critical art whimper in disagreement. If

a work of art is simply seen as just a bunch of entertaining stimuli — and if that's all it is supposed to be — AI might be just as suitable a maker as a human being.

We could also see a work of art created by AI as a kind of empty signifier. I love many subtypes of postmodernism and works of art they have produced but I'm not a fan of the kind that celebrates the disappearance of meaning and has the idea that a work of art, at its best, should specifically be an attractive and delicious yet empty signifier, free from all intentions and ideological baggage. If you believe it possible that there can be art that is ideologically and politically neutral, then AI is, of course, well suited for generating such content. Such content, however, does not help us to solve the most difficult questions, make good choices and create a good life — or even alternative definitions for a good life.

Despite all this, it is, obviously, possible to do just those things – reflect on the ways of the world and the meaning of life while enjoying a work generated by AI. But if the work lacks a human artist's touch, if the work is not an attempt by a subject who writes, paints or plays music to address the most difficult existential questions, how could I, as a recipient and a consumer of the work, really be bothered to chew on these questions?

Yet I do not see works of art as statements made by their creator or messages sent by their maker, but rather as stand-alone entities. If **Marcel Proust** was brought here in a time machine, he would not be the supreme authority to explain what the 'real' content of *In Search of Lost Time* is. Each work has its own implied author that the actual author/composer/painter/etc. cannot properly control or even understand. The author of a work owns the copyright, but from an artistic point of view, their relationship with the work is more like a parent's relationship with an adult child than a trumpeter's relationship with a trumpet. Great works are wiser than their creators, or at least they contain the best elements of their creators' thinking.

Looking at it this way, one might wonder that if works of art are not just statements made by their creators, simply their experiences and views expressed in an artistic form, if a work mainly lives a life of its own and is not only the brainchild of an auteur, a human individual who held the pen or similar instrument, but also a product of several other influencing factors, to some extent a product of its era and location and context... ... then would it not be possible, in the future, for AI to be able to sift through a massive amount of material and mix it into truly fascinating summaries and visions of the spirit of a particular era or the aspirations of an artistic or ideological movement? In a way, yes... ...but maybe not. For a work to be able to explain how we should live and what a good life is, the individual who wrote (or painted, etc.) it must have a vested interest in it, everything at stake.

A work of art emerges from its creator's sincere attempt to make life more meaningful, fairer. Or something. AI does not — at least not yet — have such concerns and aspirations. That is why it is difficult for me to take it seriously as an artist. I'll change my mind as soon as AI obtains awareness and a will of its own. I'm not sure if it would also need a body to be able to share its deepest concerns, dreams and thoughts with people. Maybe.

That's all very interesting, but fortunately, it won't have any impact on my line of work

Visual artists and artists in many other fields may be reading this text, thinking, "That's all very interesting, but fortunately, it won't have any impact on my line of work, at least not in my lifetime."

A person who paints oil paintings that are exhibited in galleries and museums and are sometimes sold may find it difficult to believe that AI could be doing the same any time soon, if ever.

The majority of visual artists create unique physical objects by hand, which are valued for the artists' personal touch, for their personal efforts and reflections. If someone were to let AI generate a hundred sketches for paintings and then printed them out on canvases using museum-grade paint, it would still be considered human-made art. The point would be what kind of instructions the human artist had given to the AI application and which works from among the sketches the artist selects for printing. It would be a limited edition — no more than five to be printed — and the human artist would still be the one to achieve fame and fortune. But it would be foolish to think that the development stops there.

Many people also hope that the advent of AI and digitalisation that has been spreading for a long time will cause a natural backlash as people will hanker for and be more appreciative of handmade, 'analogue' and unique products. I'm sure that one of the main reasons why people go to concerts and dance, theatre and opera performances is the fact that they offer experiences that are completely different from anything they experience using digital devices.

Many artists probably also think that the way generative AI threatens communications, democracy and peace in society is a political challenge, whereby individual artists, voters and consumers can do nothing else but be more vigilant and develop their media literacy skills.

I, for one, hope that each one of us would take the risks of AI much more seriously. AI is already here. No one was really prepared for it, but if we do not prepare ourselves for the next waves of AI, things could go very wrong for us, both as artists and as human beings.

There is also a more mundane and cultural-political aspect to the matter. Generative AI absorbs content from all art forms fairly equally, but its immediate downsides have been unevenly distributed until now. AI is already a notorious thief and a major competitor for music makers and screenwriters. In the field of applied arts, designers and illustrators are already witnessing how employers can use AI applications to get rid of most human designers. In the media, journalists are becoming a more endangered profession by the day, not to mention the other professions.

The statement by the German organisations has a section on fine arts, which repeats the demands made in the general section. Developers of generative AI must pay for the use of copyrighted works collected by applications. Creators must be able to opt out of crawling, and application owners must keep a public register of the materials they have collected to train their AI. There are, however, also demands that are not mentioned in any other part of the statement. One of them is:

“We demand from developers of generative AI: the obligation to contribute to the artist social security system as it applies for all exploiters by law.”

So they insist that developers of generative AI have an obligation to contribute to financing social security for artists since they benefit from artists' work. I think this is a fascinating and radical idea.

I assume that it is based on both pragmatic and ideological considerations. From a practical point of view, it will probably be difficult to develop a copyright system that would operate on the idea that whenever generative AI uses works, such as a poem I have written or a painting I have painted, as raw material, I would receive a fee that would be proportional to my work's contribution to a particular AI-generated output. This might be impossible in most cases, just like it is impossible for me to estimate the impact of **Eeva-Liisa Manner's** poems on

my poems. Therefore, the remuneration system will probably be founded, at least in part, on a collective principle, on some kind of copyright-based 'copyright tax' levied on developers and users of AI and paid to artists by collecting societies, either as remuneration determined by the use (if it is possible to determine the extent) or as dividend-like payments, grants or by some other method.

The development of a social security system for artists goes well with a remuneration system for the use of copyrighted works, for it is logical and fair that when AI needs and uses works by human artists as its raw material, its owners and users must also bear the responsibility for the realisation of human artists' rights and basic income, i.e. social security. This is not something that can be left for governments to handle. It is great that in Germany it was the association of visual artists that came up with the idea of demanding these conditions on behalf of all artists.

The art sectors should, in the name of their own interests and for the common good, form a front and use the collective force to lobby for AI legislation and other rules not only in Germany, but in all other countries as well.

If visual artists do not act now, it will backfire later. A film director may think, "Oh well, you never know, AI might offer me better scripts than human writers, and I can't really see a robot sitting in the director's chair since this job requires a human touch." This attitude will backfire later. Unfortunately, the history of art policies knows at least as many examples of disloyalty between the art forms as those of compassion and understanding the common good.

Musicians, filmmakers and performing arts professionals have found it difficult to understand the financial plight of visual artists — difficult to understand and even more difficult to get up and do something to help.

Professionals of other art forms tend to say to visual artists: “Your system is odd as you don’t earn any money for displaying your works at exhibitions or online, or you receive a ridiculously small fee, but... .. you’ve agreed to this and maybe you should just learn to sell your works — the money is great when someone actually decides to buy them. So, best of luck.”

Attitudes have changed slightly in recent years as terms of employment and other contractual terms have deteriorated in the music, film and TV industries and in institutional theatres, and the number of permanent and long-term jobs has decreased. As music makers’ earnings from recordings have plummeted due to factors such as streaming services, many of them have faced poverty or lack of income opportunities; a situation that has long been the reality for most visual artists.

Yet artists in different sectors often still find it difficult to take each other’s financial problems and copyright issues seriously.

Many visual artists play the victim and wonder how anyone could be against paying exhibition fees; why people cannot understand that visual artists often work without receiving any financial compensation for their work, even if thousands of people visit their exhibitions. And yet most of these artists, who lament how miserable their lives are, listen to music mainly or exclusively through streaming services that do not generate any income for the majority of music makers.

It is often easy for people to recognise the forms of injustice they suffer themselves and much more difficult to recognise the injustices that other people suffer and their own role in the creation of those injustices.

I hope that the situation improves and that there will be greater solidarity among artists. As the damage caused by the unfair spread of digitalisation affects more

people and an increasing number of artists in all fields are freelancers or self-employed, it is easier to see the common good and the need to unite forces.

Generative AI affects all of us now and will have a far more severe effect on us in the future. It has the potential for good and bad, and we should work together to build structures that can take advantage of the positive aspects of AI, in arts as well as in other areas of life. We also need to build structures that prevent AI from destroying communications, democratic society, opportunities for making art and art's place in the human world.

[TEEMU MÄKI](#)

Helsinki, 9 May 2023

The author is a writer, visual artist, theatre and film director and researcher. He is also Chair of the Board of [the Artists' Association of Finland](#) and President of [the IAA Europe](#).

References

Joshua Bengio: "Slowing down development of AI systems passing the Turing test". 5/4/2023. <https://yoshuabengio.org/2023/04/05/slowing-down-development-of-ai-systems-passing-the-turing-test/>

ChatGPT. <https://chat.openai.com/>

Rohit Chopra: "In India, WhatsApp is a weapon of antisocial hatred". The Conversation. 23/4/2019. <https://theconversation.com/in-india-whatsapp-is-a-weapon-of-antisocial-hatred-115673>

Facebook–Cambridge Analytica data scandal.

https://en.wikipedia.org/wiki/Facebook–Cambridge_Analytica_data_scandal

Future of Life Institute: “Pause Giant AI Experiments: An Open Letter”.

22/3/2023. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

Jamie Grierson: “Photographer admits prize-winning image was AI-generated”.

The Guardian. 17/4/2023.

<https://www.theguardian.com/technology/2023/apr/17/photographer-admits-prize-winning-image-was-ai-generated>

Have I Been Trained? <https://haveibeen trained.com>

Heart on My Sleeve (ghostwriter977 song).

[https://en.wikipedia.org/wiki/Heart_on_My_Sleeve_\(ghostwriter977_song\)](https://en.wikipedia.org/wiki/Heart_on_My_Sleeve_(ghostwriter977_song)) &

<https://www.youtube.com/watch?v=iYU9h7FEw0>

Douglas Heaven: “Deep learning pioneer Geoffrey Hinton quits Google”. MIT Technology Review. 1/5/2023.

<https://web.archive.org/web/20230501125621/https://www.technologyreview.com/2023/05/01/1072478/deep-learning-pioneer-geoffrey-hinton-quits-google/>

Holly Herndon. <http://www.hollyherndon.com/>

Holly Herndon: “AI and Music – Holly Herndon presents Holly+ feat. Maria Arnal, Tarta Relena and Matthew Dryhurst”. Sonar 2023 Festival.

<https://www.youtube.com/watch?v=Wk6T2WmhuJw>

Holly Herndon: *Jolene*. 2022. <https://songwhip.com/holly-herndon/jolene>

Holly Herndon: *PROTO*. 2019. <https://hollyherndon.bandcamp.com/album/proto>

Information Is Beautiful: “Money too tight to mention — major music streaming services compared”. 3/3/2028.

<https://informationisbeautiful.net/visualizations/spotify-apple-music-tidal-music-streaming-services-royalty-rates-compared/>

Initiative Urheberrecht. <https://urheber.info/about-us>

Initiative Urheberrecht: “Authors and Performers Call for Safeguards Around Generative AI in the European AI Act”. 19/4/2023.

https://urheber.info/media/pages/diskurs/ruf-nach-schutz-vor-generativer-ki/e9ae79fd37-1681902889/final-version_authors-and-performers-call-for-safeguards-around-generative-ai_19.4.2023_12-50.pdf

Definition of stock/production music.

https://en.wikipedia.org/wiki/Production_music

Naomi Klein: “AI machines aren’t ‘hallucinating’. But their makers are”. The Guardian. 8/5/2023.

<https://www.theguardian.com/commentisfree/2023/may/08/ai-machines-hallucinating-naomi-klein>

Zoe Kleinman & Chris Vallance: “AI ‘godfather’ Geoffrey Hinton warns of dangers as he quits Google”. BBC News. 3/5/2023.

<https://www.bbc.com/news/world-us-canada-65452940>

Laion. <https://laion.ai/>

Lucas Mearian: “Q&A: Google’s Geoffrey Hinton — humanity just a ‘passing phase’ in the evolution of intelligence”. Computerworld, 4/4/2023.

<https://www.computerworld.com/article/3695568/qa-googles-geoffrey-hinton-humanity-just-a-passing-phase-in-the-evolution-of-intelligence.html>

Louis Rosenberg: “The Metaverse and Conversational AI as a Threat Vector for Targeted Influence”. IEEE 13th Annual Computing and Communication Workshop and Conference. 8/3/2023.

https://www.researchgate.net/publication/368492998_The_Metaverse_and_Conversational_AI_as_a_Threat_Vector_for_Targeted_Influence

Jim Samuel: “A Quick-Draft Response to the March 2023 “Pause Giant AI Experiments: An Open Letter” by Yoshua Bengio, signed by Stuart Russell, Elon Musk, Steve Wozniak, Yuval Noah Harari and others...”

https://www.researchgate.net/publication/369803272_A_Quick-Draft_Response_to_the_March_2023_Pause_Giant_AI_Experiments_An_Open_Letter_by_Yoshua_Bengio_signed_by_Stuart_Russell_Elon_Musk_Steve_Wozniak_Yuval_Noah_Harari_and_others

Akanksha Saxena: “India fake news problem fueled by digital illiteracy”.

Deutsche Welle. 2/3/2021. <https://www.dw.com/en/india-fake-news-problem-fueled-by-digital-illiteracy/a-56746776>

Philip Sherburne: “Will AI Lead to New Creative Frontiers, or Take the Pleasure Out of Music?” Pitchfork. 24/5/2022. <https://pitchfork.com/features/article/ai-music-experimentation-or-automation/>

Tammen Keltainen kirjasto (Yellow Library).

<http://keltainenkirjasto.fi/kirjat/keltaisen-kirjaston-kirjaluettelo/>

Teemu Mäki & Max Savikangas: *Ihmisen jälkeen / Posthuman*. 2023.

<https://teemumaki.com/theater-posthuman.html>

Markku Tuhkanen: "Musiikintekijät pelkäävät, että Suomen tulkinta EU:n uudesta tekijänoikeusdirektiivistä suosii digijättejä". 2/3/2021.

<https://voima.fi/artikkeli/2021/musiikintekijat-pelkaavat-etta-suomen-tulkinta-eun-uudesta-tekijanoikeusdirektiivista-suosii-digijatteja/>

United States Copyright Office: "Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence". 16/3/2023.

https://www.copyright.gov/ai/ai_policy_guidance.pdf