Chapter 9

What Makes Technology "Sweet"?

Håkan Lindgren

When you visit a conference, you should always listen very carefully to the spontaneous discussions which happen during the breaks. Raffi Khatchadourian did so, and he was lucky to hear some of the most honest words that have ever been said about artificial intelligence, as well as the development of technology in general.

Khatchadourian is a writer for *The New Yorker*. In 2015, he visited a conference on machine learning, which was held by the Royal Society in Britain. When people were mingling afterwards, he spotted Nick Bostrom and Geoffrey Hinton. Hinton, who at the time was working for Google, is one of the key people who have developed neural networks, the technology behind everything that we now call artificial intelligence. Bostrom is a Swedish-born philosopher who is leading The Future of Humanity Institute in Oxford. The purpose of the institute is to prevent large-scale threats against the survival of humanity; most of all, Bostrom is worried about artificial intelligence. Let's join Khatchadourian and eavesdrop on what they are saying.

When Khatchadourian joins the conversation, Hinton says that as soon as AI exists, political leaders will use this technology to terrorize their citizens. Bostrom asks why Hinton is working so hard to create the technology he is warning everyone about. "I could give you the usual arguments", Hinton replied. "But the truth is that the prospect of discovery is too sweet." (Khatchadourian, 2015).

Because of this hope for a fleeting moment of sweetness, we now have a technology which the rest of the world has to adapt to. Hinton left Google in 2023 to be able to speak openly about the risks of AI. In 2024 he was awarded a shared Nobel Prize in physics, for having developed a technology which he has grown increasingly critical to.

The most important word in his reply is "sweet". If the inventions we end up having are not brought into existence because they are useful for humanity, and not because someone hopes to make money from them, but because their inventors believe that creating them will give them a feeling of triumph, then we need to take a closer look at the emotional motivations of those inventors. What is "sweet"? Can any technology be sweet? Maybe the pursuit of this particular kind of sweetness is leading the inventors down a very narrow path, while many other, alternative paths remain unexplored?

Robert Oppenheimer has said a similar thing when he described how the atomic bomb was developed. "When you see something that is technically sweet, you go ahead and do it and argue about what to do about it only after you've had your technical success." Afterwards, it is of course too late for any discussion: the bomb already exists, the world has changed, and the only thing we can do is to accept the new reality and adapt to it.

Oppenheimer's words can be used to legitimize anything, since the purpose of his line of thinking is to postpone all judgement until the moment when all objections have become irrelevant. Before the technology exists, you can't criticize it because it's too early; as soon as it exists, it's too late.

Between those two moments there is never any space for questions and critical reflection. This line of thinking contains a smart trick: anyone who is thinking like this is striving very hard to put himself and the rest of the world in a situation where things afterwards never could have been different. You can even pretend to be surprised or shocked by the outcome of your own actions. Humans are always very smart when we try to deceive ourselves.

If I'm focusing so hard on my own feeling of triumph, I might not be fully aware of what kind of technology I'm creating. The machine I let loose in the world might do something completely different from what I have intended. Most importantly, if I don't intend anything in particular, if I just hope I will get a satisfying kick out of seeing this machine being created, it's very likely that it will do something unforeseen. As for AI, it's obvious that this is not yet another kind of information technology. This technology destroys information.

As soon as ChatGPT exists, universities have to change their way of working. It's futile to tell the students to spend months on writing a paper, when ChatGPT can do it for them in a few seconds. Trying to expose this kind of cheating is not possible; the only thing that the universities can do is to give up. From now on, the Swedish professor of physics Ulf Danielsson has said, the only possible way of grading students is through a verbal exam, or through a written test with all cell phones

left outside the room (Danielsson, 2025). But the purpose of reading and writing long, advanced texts is not only that the students should memorize a certain amount of information – essay writing develops your mind in a way which can never be replaced with a verbal exam or a short written test. This mental training has been the main purpose of universities as long as universities have existed. If we throw it away – does what remains still deserve to be called a university?

People use ChatGPT for research, because they mistakenly believe that it is just another search engine, but ChatGPT doesn't find information, it invents information. Eugene Volokh, professor of law at the University of California Los Angeles, decided to do a test (Oremus, 2023). He asked ChatGPT for names of scholars who have been convicted of sexual assault. He got five results, three out of which were false. One of the names, the professor Jonathan Turley, is a real person, but completely innocent. ChatGPT had made up the story of his conviction, including quotations from a non-existing article in *The Washington Post*.

Chris Moran, who is head of editorial innovation at *The Guardian*, has said that his paper has been getting emails from readers who can't find the articles which ChatGPT is quoting from (Moran, 2023). The articles are signed with the names of real *Guardian* journalists, but you can't find them through Google, since they don't exist. In the current climate, when established papers are already distrusted – who will the readers trust in this case? Moran or ChatGPT? Perhaps they will believe that *The Guardian* is secretly censoring its archives.

People in the tech business know that their AI programmes give fictitious answers, and their counter-argument goes like

this: there is nothing to worry about, as long as anyone using ChatGPT understands that they need to fact check all the answers. Well, imagine that the pocket calculator, which was introduced in the 1970s, had been just as unreliable, and that it was defended with the same argument: sure, it sometimes gives the wrong answer, but as long as you can calculate the square root of pi in your own head, there's nothing to worry about.

Image-generating AI is drowning the world in fake information. If you google famous artists like Monet or van Gogh, you will find lots of paintings which they never made. Sometimes the fakes are clumsy, sometimes they are very convincing. How long can this go on, until there is no point in searching for information online? Will you stop googling when one third of the search results look like AI-generated rubbish? Maybe you won't – because what would the alternative be? Humanity might very well get used to a situation which is similar to living in the Middle Ages: the only things that exist are rumours, prejudices, myths, and self-confident bullshit.

AI-generated deepfakes are so convincing that photographs and video clips can no longer be used in journalism. A video clip of Donald Trump being arrested by the police does not prove anything. Maybe you could still suspect that something fishy is going on, if you freeze the clip and discover that one of the policemen has got seven fingers, but give the technology a few more years, and these tell-tale signs might disappear. Please note that this also means that all historical source material has become untrustworthy. Perhaps you would prefer that the Chinese army had liberated Auschwitz and captured Hitler; in that case, you could easily create some credible, scratchy newsreel footage in black and white. We have put ourselves in a position

where no information is reliable or useful without specialist knowledge – photos, text or video. For how many more generations will this specialist knowledge be able to survive in the world we have created?

Perhaps the inventors felt a moment of sweetness when they launched a technology that destroys the universities, journalism and the web as a source of information, but I doubt that this feeling will make up for the consequences which all of us have to deal with. We should not talk about technology only in technical or economical terms: what technology we get depends also on the conscious and subconscious wishes and expectations of the inventors.

The sweetness that Hinton refers to might perhaps be described as pride, at least partially. What is perceived as sweet is anything that encourages our pride and our ambitions, anything that makes you satisfied with yourself. Note that pride and satisfaction are very close to self-respect – and this is where things get interesting. Sweetness and self esteem are very similar, so similar that you might confuse them, but if you use them as guiding principles for your work they will lead you in very different directions.

Could the sweetness that Hinton is chasing be a compensation for self-esteem? What kind of technology would people like him develop if they changed their course just slightly – if they were aiming for self-esteem instead of fleeting moments of sweetness?

To understand what happens when a certain type of technology becomes popular we also need to examine another kind of sweetness. This is not the sweetness felt by the inventors, but by everyone else, who is expected to embrace any new invention.

Behind the evolution of technology, you can distinguish a warped idea of what a good life looks like: it's always better to have a machine doing something for you instead of doing it yourself. What makes us human is that anyone who wants to compose music, make movies or write novels must spend years of training to develop his abilities. We don't have to do this anymore – we only need to say a few words to our virtual assistant.

Since the beginning of the industrial revolution we have been making technology that liberates us from hard work, giving us more time to develop our intellectual and creative abilities. AI has the opposite effect. AI doesn't give us more time for writing poetry; it steals poetry from us, and turns us into consumers. The latest versions of modern technology do not allow anything but consumption. A medical doctor will consume AI-generated entertainment in her spare time, and at work she will consume the interpretations of X-ray images that AI feeds her, since the program is able to analyse these images much faster than any human.

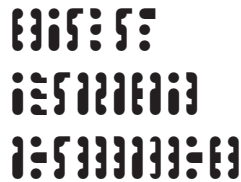

Consumption is *sweet* – this is what we are expected to feel, and what we are being instructed to feel, from infancy and onwards. Work, on the other hand, is nothing but oppression – even the kind of work that is necessary if we want to become skilled at something we are interested in. Work turns us into prisoners in a labour camp. We might not be as cold and hungry as Solzhenitsyn, but there is basically no difference: our work is nothing but punishment, just like his. Humans are made to be passive consumers – as long as we believe this, this idea will decide what kind of technology we develop, and it will decide what we perceive as progress.

When the idea of sweetness becomes a visible ideology – the ideology behind our technology and our consumption – we can also begin to see what the alternatives could look like. If we want to resist this enormous river of sweetness, and the enormous hunger for sweetness which is supposed to be normal human nature, we need a different concept of what human life should be.

We don't need to look very far to find such an idea, we only need to look where we aren't supposed to look. "The main purpose of work is to humanize us", wrotes the Polish-Swedish journalist Maciej Zaremba (Zaremba, 1982). He had travelled to Poland to write about Solidarity, which had been crushed by general Jaruzelski the previous year.

Solidarity was more than a free trade union, it was a civil rights movement with ten million members. The strength of this movement was that it united many different groups for a common purpose – intellectual dissidents, angry dockyard workers and catholic humanists. Zaremba describes the vision that some of the people behind this movement were aiming for – among them Adam Michnik, Józef Tischner and Jacek Kuroń. They didn't just want to abolish communism, they wanted to open the door to a new kind of society, which has not been tried before: "Work should serve the needs of humans, and to help them fulfil their own humanity" (Zaremba, 1982).

This sentence may look innocent, but it contains everything we need to destroy AI, and all other misguided technologies, existing or future. This sentence says two important things. Firstly, *work is not evil*. Work is not something that should be given away to machines, so that we could live a supposedly comfortable life as consumers. And secondly: *we are not yet*

ourselves. Mankind needs to fulfil its humanity. Even a fully grown human is, ethically and psychologically speaking, just a seed, nothing more than an expectation of something that could be. This should form the basis of all discussions about technology and future societies.

The ideas of the Solidarity movement were never realized when Poland became independent in the 1990s, but this doesn't mean that these ideas were nothing more than an historical blind alley. For me, these ideas are more relevant than ever, because of AI. We will never have a constructive and mature discussion about technology as long as we only debate if we should be for or against the latest technological hype. A mature discussion will only begin when we let technology rest for a while, and put humans at the centre. Instead of asking if we should be for or against AI we should start with much more basic questions: what kind of work is necessary? What kind of work would help improve human dignity?

What is a human being? What are our needs? We are not yet fully developed – but what should we strive to develop into? What kind of future society would be dignified and desirable? Questions like these have no simple answers. They might even be impossible to answer, but the purpose of these questions is not to find an answer, the purpose is to place humans at the centre of the discussion. When we have an idea of the future society which we would like to strive for, we can then ask ourselves what kind of technical tools we need in order to turn this vision into reality. Then and only then technology is in its proper place – as a tool for human needs and purposes.

There is currently a wave of populist politics in many countries. I'm sure there are many explanations for this, but one of

the deepest reasons, I think, is that our society does not offer any kind of work that improves human dignity. Our only ideals are to be rich, famous and consume luxury goods. Many people understand that they will never be part of the top few percent that are able to live this kind of life, and they are afraid that they will lose what they already have. This is where their anger and discontent come from.

If work is so good for us, like the Solidarity leaders say, why isn't all work sweet? My reply would be: we only know work in its humiliating form. This is like seeing a tired and depressed tiger in a zoo. In its natural habitat, the tiger is a completely different animal.

What would technology look like if these two ideals were the basis for the development of future technical tools – that work should improve human dignity, and that humans need to strive to fully develop their own capabilities? Try to imagine this, and you'll realize that it's just as difficult as trying to imagine a planet with a completely different ecosystem. This is our chance to dream big, and to create inspiring and optimistic ideas for the future. Hoping for AI is not thinking big enough – this is just an old 20th century fantasy about thinking machines, which should crawl back to the 1950s where it belongs.

I'm not against technology. Humans have always been tool-using animals. We will never develop our own capabilities without tools – but not all tools are equally good for us. Putting on shoes with Velcro ribbons might be more comfortable than tying your shoelaces, but an inventor who really wants to do us a favour should give us something that *complicates* our lives in an interesting and stimulating way. When we judge technology we are stuck with the polar opposites *old* and *new*: new is

good, old is boring. We should ignore the age of the tools we are discussing, and instead evaluate them for their effect on us. What technologies humiliate us, what makes us passive, what contributes to a dignified life?

A violin is 17th century technology, but for anyone learning to play it, it represents technological progress, since it develops new abilities. It's difficult to play, and that's the whole point of it. It is technology of a kind which is not meant to turn humans into consumers. A violin develops hearing, skill and virtuosity. Tools that rob us of the virtuosity and the skill we could have developed do not represent progress – they are the opposite of progress, and it's about time we realized it.

What could tools that contribute to a more dignified life look like? For a start, I believe they will insist on being perceived as *tools*, no more and no less. They do not pretend to be *you*, and they don't pretend to be another person, with a consciousness and a will of their own. They will not weaken or dissolve the boundary between us and our tools, on the contrary, they will reinforce it. I believe such tools will have more in common with the violin or the bike, than with self-driving cars.

Actually, the computer was once supposed to be a bike. In 1990 Steve Jobs described what kind of machine he would like to sell (Lawrence, 1990). In a scientific study, he said, they had compared how much energy different animals used when they were moving around. The condor was the obvious winner. No one can get as far on as few calories as a condor. Humans didn't do very well, about one third of all animals in this study were more energy efficient than we are. But, Jobs added, give humans a bike and we will beat the condor and all other animals. The important thing for him was that the study showed that we

can make tools that develop our innate abilities far beyond our biological limitations. The people who developed the personal computer, from the 1960s until the internet, hoped that the computer would be a tool of this kind: it was meant to help our own intelligence and creativity to grow. Since then, these ideals have been abandoned. Apple's real bestseller, the smartphone, is a machine for passive consumption.

The sweetness that Hinton and Oppenheimer are chasing is short-lived and it doesn't care about responsibility and moral complexities. Self-respect, on the other hand, is a long-term ideal. So is the human dignity which the Solidarity movement wanted to place as the foundation of society.

Of course we need both kinds of driving forces: the short-term and the long-term ones. The tension between them can never be resolved. All we can do is try to find a constantly shifting balance.

Will I end this essay by throwing this impossible task at the readers and then walk away? Well, I have a few more words left, and there's one thing I would like to add. The wisest words which Western civilization so far has produced are more than 2 500 years old. Γνωθι σεαυτόν – "know thyself" – was the inscription on the temple of Apollo in Delphi. These two words are the clue we need to solve this problem.

It's as if the Greeks already understood the limitations of everything that would happen later in history, from the renaissance and the enlightenment to mass consumption and artificial intelligence. As if they could clearly see the difficult situations we would end up in when we finally had acquired all the technological powers we were dreaming about. If man is the measure of everything, if we have grown so powerful that

nature does not stop us anymore – if we have achieved all this without understanding ourselves, our emotional needs, reactions and driving forces, we are still completely helpless.

Self-knowledge can lead us to the future we deserve. It may not always be sweet, but we should not settle for anything less.

References

Danielsson, U. (January 8, 2025). AI gör att bara muntliga prov och salstentor är möjliga. *Dagens Nyheter*.
https://www.dn.se/kultur/ulf-danielsson-ai-gor-att-bara-muntliga-prov-och-salstentor-armojliga/
[Accessed June 10, 2026.]

Khatchadourian, R. (November 23, 2015). The doomsday invention. Will artificial intelligence bring us utopia or destruction? *The New Yorker*.
https://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nickbostrom
[Accessed June 10, 2026.]

Lawrence, M. (Director) (1990) *Memory & imagination. New pathways to the Library of Congress*. [Film] Krainin Productions.
http://www.mlfilms.com/productions/m_and_i
[Accessed June 10, 2026.]

Moran, C. (April 6, 2023). ChatGPT is making up fake Guardian articles. Here's how we're responding. *The Guardian*.
https://www.theguardian.com/commentisfree/2023/apr/06/ai-chatgptguardian-technology-risks-fake-article
[Accessed June 10, 2026.]

Oremus, W. & Verma, P. (April 5, 2023). ChatGPT invented a sexual harassment scandal and named a real law prof as the accused. *The Washington Post*.
https://www.washingtonpost.com/technology/2023/04/05/chatgpt-lies/
[Accessed June 10, 2026.]

Zaremba, M. (1982). Solidaritets ideologi. Religionen och kampen om människovärdet. *Ord & bild* 1982, No. 1.