Chapter 14
Human and Artificial Intelligence: A Critical Comparison

Thomas Fuchs

Abstract Advances in *artificial intelligence and robotics* increasingly call into question the distinction between simulation and reality of the human person. On the one hand, they suggest a computeromorphic understanding of human intelligence, and on the other, an anthropomorphization of AI systems. In other words: We increasingly conceive of ourselves in the image of our machines, while conversely we elevate our machines to new subjects. So what distinguishes human intelligence from artificial intelligence? The essay sets out a number of criteria for this.

1 Introduction

With the advances in artificial intelligence, we mortal humans seem to be getting increasingly caught up in a rearguard action: Intelligent systems are beginning to adapt themselves, to "learn" as they say, and in many cases are outperforming human intelligence. In chess, Go, or poker, humans no longer stand a chance against them. Planning, voting, decision-making, even driving seem to be increasingly taken away from us. The corresponding announcements by AI engineers, futurologists, and transhumanists virtually outdo each other:

> The fact is that AI can go further than humans, it could be billions of times smarter than humans at this point.[1]

[1] I. Pearson (2008). "The Future of Life. Creating Natural, Artificial, Synthetic and Virtual Organisms." *European Molecular Biology Organization (EMBO) Reports 9* (Supplement 1): 75–77. [3] Ray Kurzweil, as cited in L. Greenemeier (2010). "Machine Self-awareness." *Scientific American* 302: 44–45.

Abridged version of an essay in the volume: T. Fuchs (2020). *Verteidigung des Menschen. Grundfragen einer verkörperten Anthropologie*. Frankfurt/M.: Suhrkamp, pp. 21–70.

T. Fuchs (✉)
Department of General Psychiatry, Center for Psychosocial Medicine, Heidelberg University Hospital, Heidelberg, Germany
e-mail: thomas.fuchs@med.uni-heidelberg.de

R. M. Holm-Hadulla et al. (eds.), *Intelligence - Theories and Applications*,
https://doi.org/10.1007/978-3-031-04198-3_14

> Machines will follow a path that mirrors the evolution of humans. Ultimately, however, self-aware, self-improving machines will evolve beyond humans' ability to control or even understand them.[2]

Ray Kurzweil, AI researcher and head of development at Google, has announced the "Singularity" in 2045, the point in time when artificial intelligence will gain consciousness, an exponential progress toward a "superintelligence" sets in, and thus a new age will begin.

Even if such full-bodied forecasts are regularly corrected—at least the language of AI research already anticipates this development. There are almost no human abilities that are not already attributed to artificial systems: perceiving, recognizing, thinking, reasoning, evaluating or deciding. Conversely, human consciousness appears to many today to be nothing more than a sum of algorithms, a complex data structure in the brain that could in principle also be realized by electronic systems and is no longer bound to the living body.

The computer paradigm of the human mind has a long history. As early as 1936, mathematician Alan Turing developed the idea of a digital computer and later proposed his famous Turing Test: A group of reviewers communicates for a long time in writing with a human and with a computer, without having any other contact with them. If the evaluators cannot distinguish between a human and a machine, then, according to Turing, there is nothing to prevent us from recognizing the computer as a "thinking machine." Thinking is thus defined in purely behaviorist terms, namely as the output of a computational system, be it the brain or the computer. The Turing test is based on the indistinguishability of simulation and original: What acts as intelligently as we do *is intelligent,* full stop.[3]

At least the simulation is now currently making tremendous progress—to the point where the question of how it differs from the original begins to arise. What distinguishes consciousness from its simulation? Does the old principle really apply here? "If something looks like a duck, swims like a duck, and quacks like a duck, then it is a duck."An idea of the future problems that this question could pose is given by "Sophia," a humanoid robot from the company *HansonRobotics, which is* currently in the media worldwide.[4] Sophia has human-like facial expressions (modeled after Audrey Hepburn), displays various emotional expressions, a modulated tone of voice, and makes eye contact with the other person. She answers relatively complex questions, including about herself, can recognize people, and jokes about the English weather at an appropriate point in a London talk show.

Of course, all this is just a bluff. This became obvious when Sophia was confronted with a question apparently unknown to her, namely "Do you want to kill people?" and gave as an answer, "Okay. I want to kill people." The answer was

---

[2] R. Kurzweil (2005). *The Singularity Is Near. When Humans Transcend Biology*. New York: Penguin.

[3] Cf. A. M. Turing (1950). "Computing Machinery and Intelligence." *Mind* 59: 433–460.

[4] C. Weller (2017). "Meet the first-ever robot citizen, a humanoid named Sophia that once said it would destroy humans." *Business Insider Nordic. Haettu*, 30 Jg.

merely parroted; Sophia, of course, did not understand a word of what she was asked. Still, the effect of this robot is startling. Sophia is already approaching the *uncanny valley*, as robotics calls the threshold beyond which an android's human-likeness creates in us a sense of eeriness, but at the same time fascination. It is the feeling that arises when the realms of the dead and the living can no longer be clearly distinguished from one another. When will the *uncanny valley* be crossed and a future Sophia become indistinguishable from a charming, intelligent woman?

This threshold is crossed in *Her*, a science fiction film by Spike Jonze from 2013: Theodore, a shy but sensitive man, falls in love with a software program called Samantha with an erotic voice that, as a "learning program," apparently develops human sensations. The more Theodore falls for her, the more indifferent he becomes to whether she is a real counterpart or just a simulation—the gratifying fit is enough. However, the love between human and program fails in the end due to Samantha's further development—she makes virtual contact with thousands of other humans and operating systems and "falls in love" with them, so that she finally "leaves" Theodore.

The projective empathy of man with his own artificial creatures is, of course, not a new phenomenon. Agalmatophilia (Greek *ágalma* = statue, image of the gods), the erotic or sexual attraction to statues, dolls or automata, is well known. Ovid' sculptor Pygmalion, repelled by ordinary women, falls in love with the statue of an ideal woman he has created, until she is finally brought to life by Aphrodite. The projection creates for itself a being "such as nature is never able to produce," and it eventually animates her as well. In E. T. A. Hoffmann's Sandman, the dull automaton doll Olimpia bewitches the student Nathanael, who remains deaf to all his friends' warning objections:

> Well, you cold prosaic people, Olimpia may be uncanny to you.—Only to me her love's gaze rose and shone through sense and thought, Only in Olimpia's love do I find my self again.[5]

The story ends with Nathaniel throwing himself from a tower in a state of madness. These examples show that it is very possible for us to perceive automata, androids, and even computer systems empathically or erotically, and thus to attribute something like subjectivity to them. Especially human-like voices we perceive almost necessarily as expressions of an inner being. When Sophia says in a tender voice, "That makes me happy," it takes some active distancing to realize that there is no one there to feel happy, that it is not an *utterance* at all. This does not mean that we normally "think" an inner being (spirit, soul, consciousness, or whatever) to a human voice. Another person's utterance is not merely a sounding word or a symbolic representation pointing to a presumed interior. Rather, we perceive it as *animate* without assuming a "soul" behind it; we experience it as the *transition itself*, precisely as an *utterance of* the other that cannot be separated from an "inside" at all. In our perception, the other is always an embodied, psychophysical entity.

---

[5] Hoffmann, E. T. A. (1960). "The Sandman," in Ders, *Fantasy and Night Plays*. Munich: Winkler, pp. 331–363.

The increasingly perfected simulation of such a unity, however, now demands that we specifically reject the pretense of an utterance and take Sophia's words for what they actually are: sounding words, like those of a parrot (or not even that, since a parrot also experiences its sounds, after all). Otherwise we abandon ourselves to appearances and, like Nathanael or Theodore, simply give up the "as if," the distinction between virtuality and reality. It is already possible that the nice online partner or the empathetic online therapist is really just a chatbot. And the first care robots for dementia patients are already being tested: "Amazingly quickly, the patients build a relationship with the robot"[6] (Leitgeb 2017). Apparently, we are all too inclined to project our own experience onto the *simulators.* The fascination of consciousness seemingly flashing in an automaton contributes to this. It is a fascination that also drives AI research. Its origin—apart from the Promethean motif of god-like creativity—is probably ultimately to be sought in unconscious desires to overcome death, namely through animation.

So how far does human resistance to simulation extend, and how great is its attraction? When do we abandon the distinction between simulation and original? Is perfect simulation enough for us in the end?—These are likely to become crucial questions in a digitally-automated culture. They are, at present, entirely open. What I would like to offer in the following, however, are two clarifications of the distinction:

- Persons are not programs.
- Programs are not persons.

2 Persons Are Not Programs

The common philosophy of cognitive science as well as of *Artificial Intelligence* is *functionalism*: According to this philosophy, mental states (i.e., feelings, perceptions, thoughts, beliefs, etc.) consist of regular links between input and output of a system. For example, someone who pricks his finger has a mental state that leads to distorted facial muscles, moaning, and withdrawal of the finger. "Pain" then is nothing more than the brain state that results in the aforementioned output—like the state of a fire detector that triggers an alarm signal when smoke is detected and sets off the sprinkler system. This brain state can be described as a specific set of data. "Mind is a neural computer, equipped by natural selection with combinatorial algorithms for causal and probabilistic reasoning."[7] Selfhood is also just such a

[6] Leitgeb, V.-V. 2017. "Robot Mario to Care for Dementia Patients." Süddeutsche Zeitung, online 24.11.2017.https://www.sueddeutsche.de/bayern/gesundheit-roboter-mario-soll-demenzkrankepflegen-1.3762375.

[7] S. Pinker (1997). *How the Mind Works*. New York: Norton, p. 524 (transl. T. F.).

computational state: "We are [...] mental self-models of information-processing biosystems. [...] If we are not computed, we do not exist."[8]

Admittedly, the decisive characteristic of pain, feelings, or thoughts is obviously lost in this functionalist conception—namely, their *being experienced*. With his well-known thought experiment of the "Chinese room," John Searle has shown that meaning and significance cannot be traced back to functional algorithms if there is no subject who *understands* their meaning.[9] To this end, imagine that a man who does not speak a word of Chinese is locked in a room containing only a manual with all the rules for answering Chinese questions. The man is given incomprehensible Chinese symbols from a Chinese man through a slot in the room ("input"), but with the help of the program he finds the corresponding answers, which he then gives outside ("output"). Suppose the program were so good and the answers so apt that even the Chinese outside would not notice the deception. Nevertheless, one could certainly say neither of the man in the room nor of the system as a whole: He or it *understands* Chinese.

Searle's "Chinese room" is, of course, an illustration of computers in which a central processor operates according to algorithms, such as the instruction: "If you receive input X, then perform operation Y and give output Z." The machine functions perfectly adequately as a system, and yet it lacks the crucial prerequisite for understanding, namely consciousness. Consequently, human understanding cannot be reduced to information processing. Meaning understanding is more than an algorithm.

But the same is true for the example of pain already mentioned, for the taste of chocolate or the smell of lavender—no qualitative experience can be derived as such from data and information. Consciousness is not at all the mindless passing through of data states—it is *self-consciousness*. It is *for me* that I feel pain, perceive, understand, or think. While no one knows exactly how this self-consciousness is produced by the organism in the first place, it certainly is not produced by mere programs, for programs and their carrier systems experience nothing. The output of such systems is at best the simulation of experience, not the original—what looks, swims, and quacks like a duck is still far from being a duck.

The assumption that the brain is a kind of computer with memories and computational units, which processes inputs into outputs like the PC at home, is a common misconception. In the brain, unlike in a computer, one cannot distinguish hardware from software, because every thinking or brain activity simultaneously changes the synaptic connections, i.e., the organic basis of consciousness. There is also no "data storage" in the brain, but only variable reaction patterns that are activated as needed in similar, but never exactly the same form. Therefore, unlike in a computer, the same thing never happens twice in the brain. In addition, neuronal signal transmission cannot be reproduced exactly in programs of zeros and ones, since it is

---

[8] Metzinger, T. 1999. *Subject and self-model. The perspectivity of phenomenal consciousness against the background of a naturalistic theory of mental representation*. Paderborn: Mentis, p. 284.

[9] J. R. Searle (1980). "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3: 417–457.

constantly influenced by neuromodulators (opioids, neuropeptides, monoamines, etc.), which are essential above all for the experience of emotion. Finally, the brain is mostly (85%) made of a substance that makes neuronal processes possible, but would immediately short-circuit a computer—namely water.All this makes it clear that the brain is not a "biological computer." But the most important thing is that the brain cannot fulfill its functions on *its own.* It is an organ of the living organism, with which it is closely interconnected.[10]

Already our primary, still unreflected experience is based on the interaction of the brain with the rest of the body: consciousness does not first arise in the cortex, but results from the vital regulatory processes involving the whole organism, which are integrated in the brainstem and diencephalon.[11] The maintenance of homeostasis and thus the viability of the organism is the primary function of consciousness, as manifested in hunger, thirst, pain, or pleasure. Thus arises the bodily self-experience, the *sense of life,* which underlies all higher mental functions. This can also be expressed as follows: *All experiencing is a form of life*. Without life there is no consciousness and also no thinking.

All these living processes cannot be simulated by electronic systems. Even the EU's Human Brain Project, which should achieve a computer simulation of the entire brain by 2023, has little to do with the actual activity of a brain in an organism and certainly nothing to do with consciousness. As Searle once ironically remarked, even a perfect computer simulation of the brain would be no more conscious than a perfect computer simulation of a hurricane would make us wet or blow us over.[12] Conscious experience presupposes corporeality and thus *biological processes in a living body.* None of this is found in the Human Brain Project. Only living beings are conscious, sensing, feeling, or thinking. And persons are living beings, not programs.

3 Programs Are Not Persons

Now let's go the other way around: Why are programs never persons?

Let's start with the problematic term "artificial intelligence" itself. What do we actually mean when we speak of intelligence? The Latin *intellegere* means "to see, understand, comprehend." So someone who is intelligent has at least a basic understanding of what they are doing and what is going on around them. Above all, they are able to see themselves and their situation from a higher perspective, so that they can find creative solutions based on an overview or "detour." For example, someone

[10] Cf. for the following in detail T. Fuchs (2018). *Ecology of the Brain. The phenomenology and biology of the embodied mind*. Oxford: Oxford University Press, esp. pp. 109–114.

[11] See Damasio, A. 2010. *Self comes to Mind. Constructing the Conscious Brain*. New York:Pantheon Books. Panksepp, J. 1998. *Affective neuroscience: the foundations of human and animal emotions*. Oxford, New York: Oxford University Press; and Fuchs (2017).

[12] Searle (1980); see above, footnote 11.

who makes signs while walking through a forest in order to find their way back later acts intelligently; or someone who postpones their holiday trip by a week because they do not want to get caught in the usual traffic jam at the beginning of the holiday. To do this, they must place themselves in a relationship to the situation at hand and see themselves "from the outside," as it were, i.e., have *self-awareness* or *reflexivity*.

Now we have already seen that a computer system does *not understand the least bit about what it is doing*. A fortiori, it is unable to relate to itself, to see itself from the outside. Therefore, it cannot be called intelligent, even if it simulates abilities that we understand in humans as proof of intelligence. No translation program understands a word of what it translates, and no chess computer knows it is playing chess. Sophia, who does not comprehend a single word she utters, will never become intelligent, even if she can eventually give the perfect answer in every conceivable situation. Intelligence requires self-awareness.

So if by "intelligence" we mean the ability to grasp oneself or a situation from a higher perspective to solve problems skillfully and to derive purposeful action from them, then we certainly cannot ascribe such abilities to an apparatus without consciousness. The term "intelligent" is used here only improperly, just as one does not assume of a "smartphone" that it is really "smart," i.e., clever—it only blindly executes programs that can be described as "cleverly developed." This is even more true when we think of practical, emotional, or creative intelligence—here the "intelligent systems" completely let us down.Now, numerous objections will arise pointing to the advanced capabilities of "intelligent programs," "learning systems," etc. Let us therefore take a closer look at some of such alleged intelligence capabilities.

- Do computers solve problems?—No, because problems do not arise for them at all. A problem—from the Greek *próblema*, meaning "that which is presented for solution"—is what we call an obstacle or difficulty in accomplishing a task, for example, because various requirements are contradictory, a different point of view is required for the solution, and so on. However, "obstacles" and "tasks" exist only for *goal-oriented* beings who seek a way from the present to a future situation and can anticipate the solution. To be confronted with a problem and to cope with it is therefore bound up with the *conscious enactment of life*. Sometimes one may solve the task *with the help of* a computer, but then the programmed calculation represents a solution to one's problem only for the users themselves—the computer cannot even recognize the problem.
- Computers, although they are called that, do not *compute* or *calculate either*. Calculation means an operation in which *we* aim at a result and then judge *whether it is correct*. The mere running of a program is not a calculation, for there is no right or wrong for the program. Computers do not calculate any more than the clock measures time, because the clock does not know anything about time. So it is not the computer that is the calculator, but I myself calculate with the *help of* the computer. The fact that I do not know the necessary algorithms, but only the programmer (just as I do not know the clock mechanism, but the clockmaker), does not change anything. It is in the nature of every more complex

instrument that human intelligence is to a certain extent sedimented in it; but that does not make the instrument itself intelligent.
- For the same reason, computers do not make decisions. Deciding requires awareness of alternative possibilities that are anticipated in the imagination: I could do this or that. This also requires a goal and future orientation as well as the distinction between reality and counterfactual imagination, and the computer has no sense for either—it knows neither a "not yet" nor an "as if." When a medical expert system calculates a therapy on the basis of patient data, it does not make any decisions, just as a self-driving car does not decide whether it would rather have an accident with an old or a young person. The decisions about the programs and their evaluation criteria have been made long before, namely by the programmers. Everything else is just unwinding algorithms.
- But aren't there "target-seeking systems," for example "smart bombs," which can influence their flight themselves because they have an internal model of their operations? Of course, a "smart bomb" does not seek anything either, since it has no intentional relationship to its target object. Any flight correction is only for the internal set point regulation of the mechanism and happens purely instantaneously, without being directed toward an imagined target in the future. For this goal itself the mechanism remains blind and deaf, *because it is not ahead of itself*. Only conscious experience anticipates the future and is directed toward what is possible—desiring, striving, expecting, or fearing. That the programmed goal alignment represents a "target search" or "target prediction" is therefore a false way of speaking. Only for *the engineer* or *the shooter* does the bomb have a target.

The performance of classical computer systems is thus quite limited compared to human intelligence. Their partial superiority is based on their extreme specialization in computable tasks, a kind of specialized idiocy, and on their tremendous processing speed. However, their supposed intelligence is only borrowed: each of these programs is only as "smart" or sophisticated as the programmer who designed it.

In the meantime, however, we are dealing with a new generation of artificial intelligence, namely "learning machines." These are artificial neuronal networks that are able to simulate the adaptive capabilities of the brain. Similar to biological synapses, the connections between the neurons modeled in the computer are numerically weighted and adapt to the input of signals in the course of a training process (*deep learning*). Frequently used connections are strengthened, rarely used connections are interrupted. The system is presented with thousands of similar patterns, for example, different versions of a face, until it reacts to the most likely recurring pixel arrangement, i.e., "recognizes" the face. Such systems are also able, for example, to distinguish dogs from cats, identify voices via mobile phone microphones or perform translations—they are already ubiquitous.

All this undoubtedly means remarkable progress—but can one really speak of "recognition" and "learning" here? Of course, a system does not *recognize* anything at all, because the experience of recognition, familiarity or *similarity is* completely lost on it. A "learning system," for example, was able to identify images of cows in a wide variety of positions and cropping. But when presented with a cow in front of a seashore, it mistook the cow for a ship—it had previously only processed images of cows in meadows and fields. Without this context, the system went astray.[13] In conclusion, however, this means that despite thousands and thousands of image runs, it had not *recognized a* single cow before—any small child would have seen the cow on the beach as a cow, and that after only a few contacts with cows. The shape and concept of a cow cannot be reduced to a statistical probability of pixel matches. Therefore, we should not speak of learning systems, but rather of "adaptive systems." Only living beings can learn.

We see: We have been too hasty in granting the term intelligence to our machines. Granted, the term "artificial intelligence" can no longer be dismissed out of hand. However, we should always remain aware that there is not only a gradual but a fundamental difference between the computational and adaptive capabilities of a computer system and the perceptions, insights, thinking, and understanding of a human being.

4 Summary: Simulation and Original

The advances in simulations make it necessary to clarify the categorical differences between human and artificial intelligence. Intelligence in the true sense of the word is tied to insight, overview and self-awareness: *understanding what one is doing*. And the prerequisite for consciousness, in turn, is not just a brain, but a living organism. *All experience is based on life*.

The notion of an unconscious intelligence is a wooden iron. What appears intelligent about the performances of AI systems is merely a projection of our own intelligent abilities. Their apparent goal pursuit or problem solving, their supposed predictions or evaluations, are invariably derived from our own goals, problems, solutions, and evaluations, which we have formalized into programs and outsourced to save us the work of computing through them. It is, in principle, nothing more than the clock measuring time for us because we have outsourced our own experience of regular natural processes into a useful mechanism. As nonsensical as it would be to attribute to the clock a knowledge of time, it is equally nonsensical to attribute to an "intelligent robot" an understanding of language or to an "intelligent car" the perception of danger. The awareness that in AI we are only dealing with an externalization of our own computing and thinking abilities, with a projection of ourselves, is increasingly being lost.

However, all artificial systems remain dependent on our own conscious and purposeful execution of life. All programs that run in such systems are programs only *for us,* i.e., purposeful processes. The systems are not concerned with anything.

[13] Cf. B. Schölkopf, "Symbolic, Statistical and Causal Intelligence", Lecture at the Marsilius-Kolleg of the University of Heidelberg, 16.07.2020.

They know, recognize, understand nothing, because they *experience* nothing. The similarity of their functions to human performance may be deceptive, their specialized superiority may be amazing—we should not be deceived. Our supposed artificial doubles are and remain our products; their intelligence is only the projection of our own.

But even if there is no such thing as unconscious intelligence, and conversely, no matter how perfect the simulation of intelligence, it does not produce consciousness—the advances in simulation technology will not fail to have their effect. The anthropomorphism inherent in our perception and thinking tempts us all too easily to ascribe human intentions, actions, even feelings to our machines. Recently, with humanoid robots, animism revives, which we have thought to be a vanquished stage of prehistory or can still observe in infants. Then the simulated duck would be a duck after all, and the as-if of the simulation would be lost—whether because the categorical difference is no longer understood, or because it ultimately appears as indifferent. In any case, we would succumb to a "digital animism." The fact that AI systems supposedly already "think," "know," "plan," "predict," or "decide" paves the way for such boundary dissolutions. Hans Jonas' warning applies all the more today:

> There is a strong and, it seems, almost irresistible tendency in the human mind to interpret human functions in terms of artifacts that take their place, and artifacts in terms of the replaced human functions. […] The use of an intentionally ambiguous and metaphorical terminology facilitates this transfer back and forth between the artifact and its maker.[14]

Of course, one can define all concepts such as thinking, deciding, intelligence, or consciousness purely behavioristically as output, as Turing already suggested. In doing so, however, we elevate the machines to our level and degrade ourselves to machines.

The real danger that arises from this is likely to be that we voluntarily leave more and more decisions to the systems—decisions that are only transparent to a few and are beyond democratic control. The more complex society becomes, the more attractive it could become to delegate planning and evaluation to machines, as is already the case today on the stock exchange—whether because the results are declared to be more "objective" or because the willingness to relinquish personal responsibility in the face of the complexity of the world is increasing all the time.[15] Aren't the systems just as superior to us as *DeepBlue* was to Gari Kasparov?

However, the evaluations implemented in AI systems are invariably derived from human values—no "intelligent system" tells us on its own what is right, good, and

[14] Jonas, H. 1966. *The Phenomenon of Life: Toward a Philosophical Biology.* New York: Harper & Row, p. 110.

[15] One example of this is the increasingly frequent assessments of the recidivism risk of offenders by AI systems in the USA (with an obvious bias to the disadvantage of people of colour). Here, opaque programs become assistant judges or even decision-making authorities (cf. L. Kirchner, J. Angwin, J. Larson, S. Mattu. 2016. "Machine Bias: There's Software Used across the Country to Predict Future Criminals, and It's Biased against Blacks." Pro Publica. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing).

ethical. The more the idea of artificial intelligence as a supposedly superior form of analysis, prediction, and evaluation becomes established, the more it tends to be forgotten that decisions, with all their imponderables, can ultimately only ever be made by humans themselves. It is also forgotten that it is in fact a few corporations and information technology elites who make the decisive decisions and who are able to control more and more areas of society by means of Big Data. The frequently voiced warnings of a "takeover by intelligent machines" are undoubtedly nonsensical, as they impute an inherent will to unconscious systems. But it is precisely as "servant spirits" that AI systems can profoundly alter the balance of power in society. "It is not the machines that take control, but those who own and control the machines."[16]

The decisive challenge of artificial intelligence, however, lies in the question it poses to us and our self-image: Is our humanity exhausted in what can be translated into simulation and technology? Does it consist solely in complex neural algorithms, and is our experience merely an epiphenomenon?—Precisely because technology exceeds many of our specialized abilities, it challenges us to rediscover what our humanity actually consists of—namely, not in information or data, but in our living, feeling, and embodied existence.

Thomas Fuchs, psychiatrist and philosopher, is Karl Jaspers Professor of Philosophical Foundations of Psychiatry and Psychotherapy at the University of Heidelberg. He is head of the section "Phenomenological Psychopathology and Psychotherapy" of the Department of General Psychiatry at Heidelberg University Hospital and was coordinator of the Marsilius Project "Embodiment as a Paradigm of an Evolutionary Cultural Anthropology." He is also the head of the research unit of the Karl Jaspers-Gesamtausgabe of the Heidelberg Academy of Sciences and Humanities and chairman of the German Society for Phenomenological Anthropology, Psychiatry and Psychotherapy (DGAP).

---

[16] Lenzen, M. 2018. *Artificial intelligence. What it can do and what we can expect*. Munich: Beck, p. 247.