

Bert Lehrberg

## AI as Juristic Person

### 1 The issue

In recent years, the abbreviation AI, for Artificial Intelligence, has become very common. The phenomenon of computers, or robots equipped with computers, being as smart as humans, or even smarter, has called for a lot of attention. A great number of people have, for several decades, worried about AIs taking over their jobs. Today this is true not only for factory employees, but also for taxi-drivers, and even some qualified practitioners in areas such as medicine and law.

The rapid rise of “the intelligent machines” in recent years has given reason for additional worries and questions. Scientists in areas more closely concerned with the production and development of AIs often speculate and express worries about smart machines taking over the world. Will super intelligent machines continue to follow their less intelligent maker’s instructions and do what they are told? Or will they build a superior society of their own?

If the “robots” choose a path of their own, will they then pay appropriate respect to mankind as their makers? Or will they consider us inferior and insignificant? If so, will they allow human society to live and develop further on its own? Or will they patronize us as simpleminded, and care for us like helpless children, or even keep us as pet animals? Or will they consider us as irrelevant as the rocks on the ground, and maybe, just by chance, happen to exterminate the human race? As a lawyer, it is easy to feel insignificant when confronted with questions like these. When it comes to law, would even Isaac Asimov’s three laws of robotics suffice? These are the following:

“First Law

A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Second Law

A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

Third Law

A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.”<sup>1</sup>

However, we do not yet know for sure what turn developments will take. Will machines ever become sentient, with the ability to define their own purposes and act on it? If so, when will this happen? In two years from now, or in 200 years? Or will machines stay highly specialized, and just perform the tasks their programmers put before them? Being specialized, however, does not mean not being dangerous. In the sci-fi TV series *Stargate Universe* there are unmanned spaceships specialized at warfare, strongly suggesting the opposite.

In the years to come, machines will most likely take on more and more complicated and diversified tasks, and more and more often function as autonomous entities, taking care of most of the diverse businesses of “their own”. This is something lawyers and legislators will have to deal with; and the actions they take might decide the future of not only AI but also of humanity. One of the main issues will probably be whether we should “make room” for robots and other AIs to form part of our societies as autonomous persons under the law, or if we should just keep them “as slaves” or pets – or try to. In other words: Should AIs be homologated as juristic persons? In this short essay, I have chosen to deal with this issue from the perspective of a brief analysis of three different situations where this question has already been asked, or might be asked in the future.

## 2 What is AI?

Artificial intelligence (AI) is intelligence demonstrated by machines, unlike the natural intelligence displayed by humans (HI) and animals, which involves consciousness, self-awareness and emotionality. The for-

<sup>1</sup> Asimov, Isaac, *I, Robot*, New York City 1950, p. 40.

mal definition laid down by the European Parliament is: “AI is the ability of a machine to display human-like capabilities such as reasoning, learning, planning and creativity”.<sup>2</sup> It is furthermore suggested that:

“AI enables technical systems to perceive their environment, deal with what they perceive, solve problems and act to achieve a specific goal. The computer receives data – already prepared or gathered through its own sensors such as a camera – processes it and responds ... AI systems are capable of adapting their behaviour to a certain degree by analysing the effects of previous actions and working autonomously.”

The following is a slightly different definition, used in communication within the EU:<sup>3</sup>

“Artificial intelligence (AI) refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world (e.g. voice assistants, image analysis software, search engines, speech and face recognition systems) or AI can be embedded in hardware devices (e.g. advanced robots, autonomous cars, drones or Internet of Things applications).”

A distinction is often made between AI in the form of “software”, e.g. virtual assistants, image analysis software, search engines, speech and face recognition systems, and “embodied” AI, such as robots, autonomous cars, drones and Internet of Things (IoT) applications. Some AI technologies have been used for more than 50 years; but due to advances in computing power, the availability of enormous quantities of data and new algorithms, a major AI breakthrough has taken place in recent years.

<sup>2</sup> European Parliament News: <https://www.europarl.europa.eu/news/en/headlines/society/20200827STO85804/what-is-artificial-intelligence-and-how-is-it-used> as of 29-03-2021.

<sup>3</sup> Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions on Artificial Intelligence for Europe, Brussels, 25.4.2018 COM(2018) 237 final: <https://www.europarl.europa.eu/news/en/headlines/society/20200827STO85804/what-is-artificial-intelligence-and-how-is-it-used>.

### 3 Is an AI a juristic person under current Swedish law?

Under Swedish law, two types of persons exist. These are (1) physical or natural (i.e. biological) persons (Sw. fysiska personer), i.e. humans; and (2) legal, judicial, juristic or juridical persons (Sw. juridiska personer). Both of these are defined as persons under law, i.e. the juristic person is not fictional.

Most juristic persons are some kinds of governmental or municipal authorities, companies, organizations, foundations, deceased estates, or estates in bankruptcy. Most of them are assigned a legal name and an individual registration number. Just as for a natural person, a juristic person is able to acquire and hold rights, and to undertake liabilities, e.g. based on legal promises issued by the juristic person. The juristic person also has the capacity to act as plaintiff or defendant under a court, an authority or an arbitral tribunal. Of course, one or more natural persons have to conduct all actions taken on behalf of a juristic person, and to receive all actions taken towards it. These persons may be legal representatives, such as the board or CEO (Sw. verkställande direktör) of a company limited by shares (Sw. aktiebolag).

Of course, an AI, embodied or not, does not qualify as a natural person; and as far as it is possible to foresee, such an AI will most likely never exist. Natural persons are human beings only. Neither is an AI recognized as a juristic person under current legislation. From a practical standpoint, the rules on the formation of juristic persons are decisive when it comes to which kinds of phenomena qualify as such. There are a few formal procedures for the formation of a juristic person, and the detailed requirements differ for various forms of juristic persons.

Intelligent machines are normally nothing but “things” under current Swedish legislation, although some might qualify as real estate. Most of them, such as robots and cars, are probably moveable property (Sw. lös egendom) in the form of moveable things (Sw. lösa saker), while some, e.g. a factory that does not form part of some real estate (Sw. fast egendom), may qualify as non-moveable things (which are still moveable property). Some factories and other facilities may also legally be part of a real estate. However, the current legislation is not designed to deal with entities such as AIs. Therefore, the issue as to how the law should define these is of relevance both *de lege lata*, i.e. when it comes to the construc-

tion of current legislation, and *de lege ferenda*, i.e. regarding the need for, and contents of, any new legislation.

## 4 The three cases

Notably, three typical cases may appear on the agenda for courts or legislators in the years to come. I will first introduce them here, and then discuss each of them in more detail.

The *first* of these is well known, and lawyers have been discussing it for several decades. This is the case where the Internet, or some other digital network, connects two or more computers, and these computers autonomously conduct business between the enterprises using them. Typically, the computers exchange offers and acceptances based on their programming, which can result in binding contracts between the enterprises involved. The issue put forth here is whether the pre-programmed computers may, or should, be recognised as juristic persons, or whether ensuing legal issues should be analysed as if they were so.

The *second* case regards self-driving cars and boats and other similar autonomously functioning embodied AIs. Of course, some person normally owns these AIs, and they therefore qualify as the property of that person. Questions arise that may be similar to those relating to the automated contracting discussed in the first case above; for example, when the embodied AI invoices a customer, buys fuel, orders service and repairs for the self-driving car etc. However, here we may take the scenario one step further.

In this second case, the embodied AI has, for some reason, no owner. The self-driven taxicab is, for instance, driving around minding its own business, without any new directives, programming or supervision from outside. It is making money and using it for its own good, probably repairing and continuously upgrading the car's machinery and the computer's hardware and software. The issue here is how to deal with such an autonomously functioning AI when legal issues arise. Can, and should, the ownerless AI be recognised as some kind of juristic person of its own?

The *third* case relates to artificial general intelligence (AGI), especially if it can be operating at a human level. Is it possible for an AI to develop the abilities typical of human natural intelligence (HI), such as consciousness, self-awareness and emotionality? If so, the question arises as to whether such a sophisticated AI should be homologated as a person

under the law, and what kind of person that would be, i.e. a natural person by analogy, or a juristic person, or something in between.

## 5 Case one. Automated contracting

Automated contracting, conducted by computers connected over the Internet or some other digital network, is the first of the three cases dealt with in this short overview. In these cases, parties are concluding contracts, i.e. exchanging legal acts, mainly in various forms of offers and acceptances, by means of electronic messages, generated and exchanged automatically by and between pre-programmed computers. These technologies have been in use for many years in connection with automated Electronic Data Interchange (EDI). The basis for this practice is generally some kind of agreement between the enterprises that are parties to the contracts, laying down the protocols that define the detailed conditions under which these automated contracts are to be concluded (EDI-contracts). The issue of whether the computers (or their programming) are to be recognized as persons has also been on the agenda for a long time.

From a Swedish point of view, the main question in relation to automated contracting is whether, and if so how, the automated forms for conclusion of contracts could be reconciled with the basic rules of contract law and the diverse theories of scientific contract theory. The Swedish Government Official Report on document management<sup>4</sup> has most extensively dealt with these matters. The committee's starting point is the issue of whether a contract has to be based on a common will, i.e. the parties' mutual intention to be bound by the contract, which is possible only when natural persons actually see and take part in the exchanged declarations of intent. This approach may be questioned, because contracts under Swedish law are actually concluded through the exchange of declarations of intent (i.e. will) (Sw. viljeförklaringar) without any general demand for the existence of a common will (Section 1 Contract Act). However, the three alternative solutions discussed by the committee are still of interest for the analysis here.

The solution advised by the committee (its *third* alternative) is, in principle, that the traditional rules and theories, based on the concept of "declaration of will" (Sw. viljeförklaring), are not applicable when the parties or their duly authorized representatives (i.e. natural persons) are

<sup>4</sup> SOU 1996:40. Elektronisk dokumenthantering.

not personally and directly involved. Instead, automated contracting should be dealt with as a case where the actions as such lead to a contract independently of the will of man or machine. I submit that this position is based on an over-reaction, and therefore goes too far, which I will elaborate below.

Another solution (the committee's *first* alternative) is to presume the existence of the will of the relevant natural or juristic persons, to be bound by the legal promises issued on basis of the pre-programmed automated routines. This may be an actual will, or a merely hypothetical one. This model is ruled out by the committee, based on the argument that the will of these persons could not be presumed to be as detailed as is actually the case for the automated EDI-routines applied at the conclusion of the contracts. However, this reasoning is at odds with how Swedish contract rules are normally applied. A party to a contract does not even have to read the text of a contract to be bound by it, let alone form a will including (a correct understanding of) all of the details of the contract.

The fact that the pre-programmed instructions are often complicated, with the output depending on information from diverse sources, does not deprive the computer of its character as a tool. It is enough for the concept of declaration of will to be applicable, that the party, i.e. a (duly authorized) natural person, wants the computer to issue legal acts (offers and/or acceptances etc) in accordance with its programming; or at least through the programming of the applied protocols has demonstrated such a will. The main difficulties, when it comes to the application of Swedish contractual rules on automated EDI, actually relate to the application of other rules, especially those requiring good or bad faith on the part of any of the parties. How could a person be in bad faith regarding something he or she did not know?

The *second* alternative solution discussed by the committee is of more interest to the current analysis regarding AI as a juristic person. The idea is that the information system is to be considered as some kind of third party, and consequently the distinctive Swedish rules on agency (Sw. *fullmakt*) to be applied analogously. The committee, of course, rejects this approach, because there is no involved third party with legal personality under current law. Obviously, Swedish law denies the computers (and their programming) a legal personality. Normally, the computers involved in EDI contracting could, and should, be considered merely as tools used by the parties to the contract, by the help of which they execute their intentions to contract under specified conditions. The com-

puter is nothing but a machine, working in accordance with its pre-programmed instructions.

Although the theoretical model, where the situation is analysed *as if* the computer was a juristic person, may not as such be totally without merit in all situations, it is obvious that the answer to the current question is that the computer is not recognized as a juristic person. There are also no strong arguments in favour of recognizing it as such. The computers, the software, and the tasks performed are not advanced enough to make such legislation necessary or even practical.

## 6 Case two. Practical situations where an AI may function as an independent person

Some known forms of, at least theoretically, autonomously functioning embodied AIs are self-driving cars, such as taxicabs, couriers and messenger's cars, or boats, drones, and robots, and some other moveable devices or facilities connected by the Internet of Things. Such an autonomous unit might in the future be able to function totally by itself, without any instructions, such as commands, new programming, updates or input of information etc, or supervision, diagnostics, testing etc from an owner or supervisor.

Thus, the taxicab (or taxi boat etc) would be able to receive customers' orders or bookings directly, e.g. from an Internet site or an app. It would probably be electric, and be able to charge its batteries at a charging station etc. If there is an available bank account that it could use, or if some cryptocurrency such as Bitcoin is accepted, it would also be able to receive payments from customers, as well as pay for the charging. Most likely, it would also be equipped with the necessary equipment and software to diagnose any faults or injuries to the car, or to the hardware or the software, and to order roadside assistance and repairs, as well as updates, and also to pay for these.

Of course, some person would normally be the owner of such an embodied AI. The device would therefore qualify as the property of that person. However, the situation becomes more interesting to us here if the embodied autonomously functioning AI has no owner or supervisor etc. The taxicab is, for instance, driving around literally minding its own business, without any supervision. It is making money and using it for its



own purposes, probably repairing and upgrading its machinery, hardware and software continuously over time.

A relevant question is, of course, whether such an autonomous embodied AI without an owner could possibly exist. The answer probably depends on how future society is organised. To be able to address this issue properly, it is necessary to assume that the AI had an owner initially. Of course, it may happen that an AI (or several of them) was constructed and built by another AI, and therefore never had a legal owner. However, if so, that “mother AI” – or its “mother”, “grandmother” etc – at least initially had an owner. So, how could the owner let loose its AI? It is possible that the owner was a natural person who died, leaving no heirs and no documents regarding the AI, and that nobody knew of it. Maybe it was even the intention of the owner to “set the AI free”, as a slave owner could do with a slave in the old days. The same thing could happen if a company owning an AI went bankrupt or was otherwise dissolved without anybody knowing about the AI.

The lack of ownership of the AI may, of course, somehow be resolved. The rightful owner, e.g. a forgotten heir, might turn up and claim his/her/its right. The owner’s estate or bankruptcy may claim the AI on behalf of the creditors. The Swedish Inheritance Fund may claim its right to the property that the devisor left behind. Somebody may claim their right to the AI based on occupation, after having sized control of the AI, where there is no owner able to prove ownership. Alternatively, there may be other solutions laid down in future legislation. The issue here is whether such future legislation might state that the AI qualifies as some kind of juristic person.

Of course, there are considerable arguments against such a legislation. AIs are things under current legislation, and there are no convincing moral arguments in favour of considering them as anything else. Having a number of unidentified autonomous embodied AIs conducting business on their own, without any supervision, may also create problems that we cannot even begin to imagine.

However, there may also be valid arguments in favour of recognizing these AIs as at least some kind of limited juristic persons. Maybe they have grown in number, or maybe they fill a valuable function as part of the economy, or maybe the cancellation of them would cause even bigger problems. Putting them in a register, and recognizing them as some kind of limited juristic persons, might then solve some of the problems that they cause. This would also make it possible to oblige them to pay taxes,

if they do not already do so, and to identify, count and supervise them, to make sure that they are properly served, repaired and upgraded etc.

The type of business for which the AI is recognized and registered might limit the individual AI's legal personality. Thus, the self-driving taxicab would, for instance, be able to enter only into contracts related to the AI itself, and its machinery (service, repairs, upgrades etc), and to the taxi business that it is conducting, but not other contracts. Some types of contracts, such as bank account agreements, would be related to most forms of AI businesses, of course.

My conclusion is that it is possible that autonomously functioning embodied AIs may someday in the future be recognized as at least limited juristic persons for merely practical reasons.

## 7 Case three. Artificial general intelligence at a human level or beyond

### 7.1 Some different classifications of AIs

A distinction has commonly been used in recent years, between *narrow* (also called weak) AI, which is specialised for one or a few tasks, and *general* (or strong) AI, capable of performing most of the activities of humans. A European expert group suggests the following definitions:<sup>5</sup>

“A general AI system is intended to be a system that can perform most activities that humans can do. Narrow AI systems are instead systems that can perform one or few specific tasks. Currently deployed AI systems are examples of narrow AI. In the early days of AI, researchers used a different terminology (weak and strong AI). There are still many open ethical, scientific and technological challenges to build the capabilities that would be needed to achieve general AI, such as common sense reasoning, self-awareness, and the ability of the machine to define its own purpose.”

A general issue is whether an artificial general intelligence (AGI) qualifies, or will in the future qualify, as a juristic person. Is it possible that an AGI can develop in such a way as to be so similar to a human being that we, for moral and/or practical reasons, have to recognize it as a person

<sup>5</sup> <https://digital-strategy.ec.europa.eu/en/library/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>.

under the law? For this issue, the concept of “AGI at a human level”, or “human-level AGI”, is highly relevant. Is it possible for an AI to develop the abilities typical of human natural intelligence (HI)? If so, the question arises as to whether such an AI should be recognized as a person under the law, and what kind of person that would be, i.e. a physical (although not biological) person by analogy, a juristic person, or something in between.

Another highly relevant concept is “Superintelligence”. This refers to an AI that is superior to even the brightest and most skilled humans in all or virtually all tasks. However, if such an AI comes into existence and becomes sentient, the main issue may not be whether we should recognise it as a person, but whether it will recognise us as persons, and as worthy of existing. Still, as lawyers it is our responsibility to deal with the issue of whether a superintelligent AI is to be homologated as a juristic person.

## 7.2 Comparisons between AI and human intelligence

Scientists have identified a long list of differences between an AI and the human mind. Some of the properties they emphasize are typical of existing narrow AI systems, and might not exist in an advanced enough AGI, or especially in a human-level AGI. Others are more general, as they relate to how AI and HI, respectively, come into existence, or relate to the basics of artificial perception, analysis or output. Of course, it is not always possible to foresee which of the typical limiting AI traits it might be possible to eliminate. Maybe all of them. For instance, it might in the future be possible to construct biological robots, that are virtually impossible to distinguish from humans, but which may be superior to us in most ways.

The natural starting point for a comparison between human intelligence and artificial intelligence may be that humans are a product of nature, while computers and robots are *synthetic*. Humans consist of biological components, such as bone, flesh and blood, while typical components of an AI, at least as of today, are made of metals, plastic and other non-biological material. In an AI, there is a basic distinction between hardware and software, which does not exist in the human mind. This calls for a radically different way of functioning. HI is also (mostly) analog, while AI is essentially digital.

When it comes to size, the human brain contains approximately 100 billion neurons, with  $10^{15}$  connections between them, while artificial

neural networks normally have considerably fewer neurons, maybe just a few hundred. Still, the AI beats the HI by far when it comes to speed. It is also, as such, less biased and more precise about details. HI is, on the other hand, more universal, as we usually learn how to manage hundreds of different skills during a lifetime, and are better at multi-tasking, and coordinating complex movements. However, all of this may change to the advantage of the AI. An AI that is not defective (in its hardware) or corrupted (in its software) will often be better equipped for processing and detecting details, especially in great numbers, in comparison to HI.

The comparisons between human intelligence and artificial intelligence often focus on *advantages and disadvantages* of AI as compared to HI, such as speed and adequacy in details. However, the decisive factors, when it comes to recognizing an AI as a juristic person, may not be how good the AI is at any (or all) specific tasks. Most tools are actually better at specific tasks than are humans. Why else would we use them in the first place? However, we still do not see any good reason to recognize them as persons. Cars are much faster than we are, washing machines are better at washing clothes, and computers are much better and faster when it comes to advanced calculations. Obviously, the most relevant issue may not be who is the best, the fastest, the most impeccable, the best at multitasking, the most sustainable, the cheapest etc. Nevertheless, these comparisons are interesting, as they may help us to decide, for instance, whether we should trust AIs as persons or not.

More relevant differences between AI and HI may be derived from a comparison between how the *processing* works. Thinking by the human brain is conducted by brain cells, i.e. nerve cells in the form of neurons, communicating through synapses in specific patterns. In an AI, the software copies the functions of the brain neurons to form computing systems called Artificial Neural Networks (ANNs) or just Neural Networks (NNs). However, they do not look, or work, exactly like the biological neural networks in the brain – they work by analogy with the brain, not homology.

The artificial neurons that form the network's nodes rather loosely model the neurons in the brain. They are usually arranged in layers, which may have different functions. All of the layers are usually connected to all of the other layers. Each connection, called an edge, transmits a signal to other neurons, which process it, and in turn signal other neurons. While the neurons of the brain either fire or do not, firing in an artificial neuron is mimicked by continuous values. The artificial neurons can smoothly

slide between off and on. The “signal” is a real number computed by a non-linear function of the sum of the neurons’ inputs. The weight of the neurons and edges typically adjusts, as a means of learning.

When it comes to the *structures* under which the AI and the HI work, there are also important differences. An artificial neural net starts from scratch all of the time. The neurons are neatly ordered and addressed one after the other. The human brain on the other hand has many predefined structures wired into its connectivity. It also has specialized regions.

Another avenue of information more relevant to our purpose here might be an investigation into how AI and HI respectively *learn*. The details of human learning are not yet known. However, HI is a product of natural development that has been ongoing for many thousands of years, and also of natural learning from the individual person’s own life experiences. We learn from various incidents and past experiences, and from mistakes made in a trial-and-error process. Then we adapt to new environments by utilizing a combination of different cognitive processes.

AI, on the other hand, learns by evaluating outcomes and adjusting the weight of the neurons and their connections. An AI is basically not capable of unsupervised learning, such as is done by a child, and it lacks intuition. AIs of today have basically been developed for specific tasks only. Through deep learning, where they are confronted with a great number of situations, they may acquire a superior skill when it comes to judging those situations, and may also develop an ability to judge new situations, which to some extent compensates for lack of intuition in the current context.

### 7.3 What is required for an AI to be recognised as a sentient being?

For an AI to be recognised as a physical entity that is equal to a natural person under human law and in our society, it obviously needs to fulfil quite a few requirements, when it comes to the abilities necessary for it to be able to function in that society, to care for itself, and to not harm others or cause other problems. It needs sensory faculties, memory, some kind of common-sense reasoning, an ability to make decisions, and much more.

However, when the issue of whether to homologate an AI as a physical juristic person reaches the agenda of lawyers, most, or maybe all, of these requirements will probably already have been fulfilled. It is, of course,

possible that the main issue will be whether AIs are harmful to humans or not, or whether humans will accept robots as their equals or not. Maybe the AI will be subject to a new form of “racism”? Quite a few sci-fi books have been written, and movies made, about such issues. To be comparable to humans in fundamental respects, the AI also needs the ability to define its own purpose. On the other hand, such an ability might turn out to be harmful. Still, if the purpose of the machine is predestined by its maker, how can it be recognised as equal to that same maker? Why should we not just consider it a tool? In other words, the AI needs to be free to decide about its own purpose and make its own decisions, but still not free to commit murders or other crimes, with the possible exception of those of a trivial nature.

When it comes to the requirements an AI has to fulfil, the most crucial, and maybe also most difficult one to accomplish and evaluate, might be that it needs to be sentient. It is suggested that this means that it needs to be conscious and self-aware and able of feeling emotions. However, how do we know whether it is? This will most likely be a complicated task for AI scientists and engineers to decide. Of course, we probably don't necessarily need this. If we want to recognise an AI as a juristic person for practical reasons, it might be enough in this respect that it behaves as if it was conscious and self-aware (and of course has proven not to be dangerous or otherwise harmful etc). However, for us to feel the need to recognise robots etc as equals for moral reasons, they probably need to actually fulfil some requirements when it comes to consciousness, self-awareness and probably also emotionality. They need to be conscious and aware of themselves and have feelings. Consciousness is the basis of this.

Science does not yet know very much for certain about consciousness. Maybe only humans or perhaps some of the more developed animals are conscious. It is also possible that consciousness is everywhere and even rocks are conscious. To get to know consciousness, we might have to look within. Yogis and other spiritual practitioners have done so for thousands of years. Of course, such introspection does not qualify as science. Nevertheless, it is the same type of observational knowledge on which science is based. Yoga has also been called “The science of the subjective experiences”. It is of course easy to predict a number of problems for a scientist willing the attempt to “make science” out of these experiences.

People having spiritual experiences do not usually speak of them, and when they do, what they say is not always scientifically useful. Most yogis and spiritual leaders, who speak of valid spiritual experiences, do not

provide a short scientific description of their observations during “higher states of consciousness”. Instead, they often want to explain and adjust them, to make them fit into an existing belief system, or sometimes even to establish a new religion or sect based upon them. It is also not always easy to put such experiences into words. On the other hand, equipped with an adequate set of questions, scientists might eventually be able to identify and categorize specific spiritual experiences, and thus gain some extra insight into the domain of consciousness.

Once, many years ago, I was told an old story about a Zen master and his young disciple. The disciple was told to meditate on the sound of one hand clapping. He did not understand. Two hands clapping together was needed to make a sound. The first days and weeks, he now and then thought he might have got the answer. Was it the sound of the hand clapping on the ground? No, that was not it. Was it the sound of the wind that the hand produced when moved fast? No, that was also wrong. He tried other answers too, but none of them worked. The master just shook his head.

At last, the disciple got so exhausted that it triggered a scary spiritual experience. Afterwards, he was frightened and ready to give up his spiritual quest. However, first he described the details of his experience to the master. The master then lit up. This was an experience of Type 4, subtype so-and-so, from the sound of one hand clapping. The disciple had reached his goal. Whether any of this is true, I don't know. Still, it suggests that the experiences from spiritual exercises might be the same for different people, and that they can be categorized.

An interesting angle when it comes to AIs is that there probably has to be something to be aware of, and this something might have to be of a specific nature. Even if the rocks on the ground are conscious, what could they experience? Humans are conscious, but we still do not normally consciously experience anything during narcosis or deep dreamless sleep. Human thinking is multidimensional, and works on the basis of some kind of inner “map of the world”, created on the basis of the hereditary structures of the brain and nervous system, memories from past experiences, and sensory input. Internal images, which may consist of visual images, sounds, feelings, taste or smell or a mix of some or all of these, form essential parts of this “inner map”.

The “map of the world”, or something similar, might actually be a component necessary for consciousness to arise. The current AIs' processes are not like that. They are linear, at least basically, and lack such a

map of the world. Neural nets do not build models of the world. Instead they learn how to classify patterns. Maybe such processing is not adapted for consciousness, and maybe there is just not enough there to be conscious about, when there is no inner map to relate to? Only the future will supply us with the final answers to such questions.

To sum this up, the most crucial abilities typical for human natural intelligence are consciousness, self-awareness, emotionality, and the ability to create an inner image or map of the world. Maybe all or some of these are necessary for an AI to qualify as sentient, and on that basis be recognised as a juristic person.

## 8 Conclusions

The issue of whether an AI should be homologated as a juristic person may occur in different situations, and cause diverse types of legal problems, which may call for disparate types of legal analysis and solutions. In this article, three types of cases have been used to illustrate the diversity of miscellaneous problems, types of analysis, and resolutions.

In the first case, regarding automated contracting, the crucial issue comes up under current law and is dealt with as a typical legal issue *de lege lata*. The AI is presumed to be weak and its job assignment narrow. There is obviously no need to recognize the contracting computers (or their programming) as persons. Still, the theoretical model for which the idea of the computer as a separate legal entity clears the ground, might sometimes be a helpful tool for the legal analysis.

The current issue may also, as in case two, present itself as a feasible practical resolution to a potential future problem, caused by embodied AIs that function autonomously without any owners taking care of them. In such a scenario, the legal personhood of the AIs may at least primarily serve an economic and pragmatic purpose. The arguments in favour of recognizing the AIs as juristic persons are, at least potentially in a possible future, considerably stronger than in the first case. They are also of the same general category as those arguments, on which the current laws of juristic persons in the forms of companies, organizations etc are based. It is, however, suggested that it might be sufficient to grant the AIs in this scenario a legal personhood that is limited to what the embodied AI needs in order to continue conduct its business, and thus contribute to society.



A major ethical dilemma, illustrated by case three above, is caused by any future generation of advanced AIs that qualify as sentient beings; and it is surprisingly often presupposed in scientific discussions, as well as in sci-fi books and movies. Still, this presupposition must, as of today, be considered rather speculative and hypothetical. Here are also multitudes of issues to deal with, for scientists, philosophers and engineers, as well as for lawyers *de lege ferenda*.

In the analysis above, I arrived at the conclusion that consciousness, self-awareness and emotionality may be crucial for an AI to qualify as sentient. It is also suggested that the ability to create an inner image or map of the world might be decisive for the prospect of consciousness to arise. If a sentient AI ever exists, ethical arguments will probably be of major importance, as support for any claim for it to be “freed from its slavery” and homologated as a person equal to humans under the law. Then, a number of other issues may also surface, such as safety issues – for the AIs as well as for humans.

