Review of: "[Essay] Not Quite Like Us? — Can Cyborgs and Intelligent Machines Be Natural Persons as a Matter of Law?"

Paul Dumouchel[1]

1 Université du Québec à Montréal

Potential competing interests: No potential competing interests to declare.

The idea or assumption that AI systems *should* be granted rights or personhood underlies much of the discussion concerning AI and ethical issues – not only in this paper – but in general in the discussion as it has been carried for a while. In consequence, the burden of the proof tends to be shifted towards those who think that we should not grant them such rights or personhood. They are asked to provide reasons why we should not, as if it were evident that we should. It is however clear that the question can be reversed, and we can ask: why should we grant rights and personhood to AI systems? Isn't it evident that they are only machines? To that question – why should we grand them rights or personhood? – all the answers I have seen so far, are variation on: because they are smart, because these machines display higher cognitive abilities, or autonomy, or the ability to decide.

In consequence, the whole discussion – as is the case in this essay – turns around the question of the extent to which this is the case. That is to say: Are the (present or future) cognitive abilities of AI systems sufficiently similar to those humans to justify the claim that they should have rights or be legally recognized as "human persons"? How can we make the difference? What criteria should we use to determine this? How significant are whatever differences we can find?

However, underlying this discussion and giving it its apparent importance is a hidden hypothesis: that we give rights to humans and consider them as persons because they have higher cognitive abilities. Without wanting to be excessively cynical, I see very little evidence of that. In many jurisdictions – especially in the USA – we give extensive rights to and recognize as human persons unborn fetuses. We give rights and recognize as human persons severely malformed and intellectually deficient babies and adults. We give right and recognize as persons older individuals whose cognitive abilities have been drastically reduced by dementia or Alzheimer disease, and we protect all of these individuals against physical and other forms of abuse. This is the way our legal system presently works. When an elderly person is victim of a fraud, we do not ask does he or she have sufficient cognitive abilities to have rights? To the opposite we consider that their intellectual and physical frailty makes them more vulnerable to such abuses.

So, in this response I intend to avoid the discussion concerning the difference or similarity between the cognitive abilities of humans and artificial cognitive systems for many reasons. One, is because I think that it is to a large extent irrelevant. Second, because the task of comparing human and artificial intelligence is full of pitfalls. There is a lot of

disagreement about how the cognitive abilities of artificial cognitive systems compare to those of humans. Disagreements that go from the incomparable superiority of AI to the claim that these are mere machines that mimic intelligent results but that have zero capacity for thought or semantic understanding. Hence, if courts are looking for experts, what they will find is disagreements. My main reason, however, is that, I believe, there is much more involved in why we recognize human beings as legal persons than the simple fact they sometimes manifest higher cognitive abilities.

Consider the following fictional example. Robert Bob is a self-taught computer expert who lives in a rich mansion in Florida and owns an expensive apartment in Hong Kong. On May 7, 2018 he sold his apartment in Hong Kong, on May 19 he criminally set fire to his Florida house where is wife was sleeping. Severely burned she survived, but on the next day he hacked the hospital centralized computer system, making it appear as an accident, disabling her life support system causing her death. He disappeared after cashing in both the fire insurance and his wife life insurance premiums. On June 12 of 2019 he resurfaced in Alabama where he rented a car using a stolen credit card. Later in 2019 until early 2021 he worked in a real estate company during which time he succeeded in embezzling millions of dollars. In November of 2021 he is appointed associate professor of mathematics at the University of East-Anglia where he only teaches online. In February 2022 he is involved in a scheme to defraud the Singapore stock market. In March 2023 while on a skiing trip in Banff Canada he got into a bar fight and was arrested. Unfortunately for him, the person he attacked was the new partner of his wife's sister and she recognized him as Robert Bob fugitive and murderer of her sister.

So, all of this – apart from driving the rented car (Teslas are not yet up to par) and a bar fight, a point that will become important later on – could have been done by an AI agent or entity. Selling a house, setting fire to an intelligent house, false identities, using a stolen credit card, impersonating a university professor, defrauding insurance companies, embezzling real estate agencies, manipulating the stock market, all of these are well within the present capability of artificial agents. Now, Robert Bob is a human being and thanks to good police work, some international cooperation, and luck he was found and now he will be charged and held responsible for his crimes. He was found because the man involved in the bar fight was the same person as the one who set fire to his house and killed is wife and because that person is an object individuated in physical space. Whether Robert Bob is in Hong Kong, Singapore, Florida, Alabama or Banff Canada, using a stolen credit card, embezzling a real estate agency, murdering his wife, or criminally setting fire to his home, he is the same person. He may have had multiple (false) identities during those years, but he is the same person, individualized in physical space. That is why he was recognized by his sister-in-law and why he will be held legally and morally responsible for what he did.

Imagine now that chatGPT or some other AI system did this. As mentioned earlier all that Robert Bod did is within the capacity of present AI systems. The first problem which arises is one of identity. Who did this? Is the AI that hacked the smart home in Florida causing a fire, the same that tried to scam the Singapore stock exchange? Is it the same that impersonated the university professor in East-Anglia who used a stolen credit card in Alabama? How do we answer such questions? Even if the entity used has the same commercial name, like chatGPT, the necessary computing and physical operations that correspond to it answering any question, whether it is writing a job application or giving a doctor indications while she is doing surgery are simultaneously carried out in many different locations in physical space, on different physical devices. Parts of it takes place on your computer or smart phone, others in different data centers and the whole

process involves an array of different technical systems. AI entities are not objects in physical spaces and whatever it is that we call their actions are at best fleeting patterns of electromagnetic waves. How do we individuate such objects? How can we know who acted where and when? There are no clear answers to such questions.

That is one of the reasons why, when accidents happen or crimes are committed through the use of such artificial agents, our present strategy is that the human agent, owner or user, is responsible. Think for example of the Boeing 737 Max accident caused by the Manoeuvring Characteristics Augmentation System. This is consistent with the fact that the autonomy of artificial agents is not really a characteristic of the agent. It does not correspond to anything it can or cannot do. It is best seen as a meta-characteristic, more precisely as a means of managing the capacities of an agent. Take a drone for example, it can fly from point A to point B and survey an area. Whether we input the information before it takes off and let it autonomously or we guide it from a distance, in both cases it does exactly the same thing, it flies from point A to point B and surveys the area. Depending on various circumstances, for example, the weather, the nature of the mission involved, our personal interest in it, etc. we opt for either one or the other mode of managing the system's abilities and chose either autonomy or control. And in all cases, we are the ones who decide.

However, if AI systems are to have rights, they will have duties and if they are recognized as persons legally, they will have legal and criminal responsibility. How can we attribute duties and responsibility if we cannot individuate the agent? Think for a second of a more well known and common system, GPS. It (?) answers thousands probably millions of demands per second. Is it the same individual entity that answers them all or are there many of them? GPS can answer queries in many different languages. Is it the same one who answers me in French in Quebec City who answers in Korean or Japanese my friends in Seoul or Kyoto? Is the AI entity in my car that tells me how to go from point A to point B the same as the one on my phone, which is presently also in my car, or are these two GPS different "objects" or "entities"? If we are to attribute responsibility to such artificial agents, to who will it be attributed and how will we do that, determine who or what has the responsibility? In the case of Robert Bob it may be that he did not act alone or that some of those crimes should be attributed to someone else, but here the problem is different. The question is not, was the crime committed by this one or that one, but how do we define someone or even one?

Not so long ago, slaves, women, foreigners were not recognized as legal persons or as full legal persons. The transformation did not happen because we suddenly became aware that members of these groups had sufficient cognitive abilities. If anything, that change in our perception of their cognitive abilities happened the other way around. It was because we first recognized them as legally equal persons, that we progressively came to shed our prejudices and admitted that their cognitive abilities were just like ours. So why did we do it? Why did we recognize them as equal legal persons? According to some, it is because those who initiated those changes believed in equality, according to others, it was because it served their purpose, especially in view the changing nature of the labor market. Leaving that discussion to the side – but note that the two positions are not mutually exclusive – the short answer is because those who brought the change about thought that it was socially advantageous for them, and to some extent for everyone.

The issue here is the same. The questions that we should be asking are: What are the advantages of granting rights or personhood to such artificial entities? What are the drawbacks? Whose purpose will it serve? What advantage will they

gain? What advantage will gain?

This said, in closing I would like to point out an important mistake in the Essay that clearly needs to be corrected. Page 7 Prof. Gervais writes:

Can we not use instead quantitative tests based on the amount of human DNA in a given entity or organism? That percentage would have to be well above 95%, given that humans share at least that much of their DNA with other species (Britten, 2002; Sapolsky, 2006). But then, a human with artificial limbs would have in total less than 95% human DNA.

First, from a biological point of view it is not clear what “human DNA” means? DNA is simply DNA there is not such thing as “human DNA”, “lobster DNA” or “dog DNA”. Therefore, the claim that an organism should have 95% human DNA is meaningless. What the articles quoted in reference indicate is that DNA in humans is 95% similar with DNA in chimpanzees. The difference between us is more or less 5%. So those numbers, 95% or 5%, do not refer to a definite quantity of an object called DNA, like the amount of water in a glass, but to a degree of similarity (or lack thereof). In consequence the claim that a person with artificial limbs would have less than 95% human DNA is simply false. Something that should have been evident since, were that the case it would entail the absurd conclusion that tall people have more DNA than others! All humans, tall, short, skinny or big, amputated or not, with or without artificial limbs share with Chimps, our closest cousins, more or less 95% of their DNA.

It seems to me that as researchers when we make arguments based on scientific information, we have a responsibility to be well informed and to understand the science we intend to use. This is a fundamental condition if we are to favour rational and fruitful discussion on the issue we are addressing.

Paul Dumouchel

Université du Québec à Montréal