

A Dialogue Concerning Humanity 2.0

Keith Wayne Brown, Brighton Dwinell, Bob Frodeman, Steve Fuller, Lauren Griffith, Carl Jacob, David Silverberg, Natalie Szczechowski, Mike Watson

What follows are Steve Fuller’s responses to a set of questions posed by the Metaphysics undergraduate class convened by Robert Frodeman and Keith Wayne Brown in the Philosophy Department at the University of North Texas. These questions were initially sent in preparation for Fuller’s video conference call to the class on 10 April 2014. Among the works considered in the class was Fuller’s [*Humanity 2.0*](#). The named questioners, other than Frodeman and Brown, are students in the class.

Natalie Szczechowski

What is “normal” now and how do you feel that definition will adjust within humanity 2.0?

Steve Fuller

As has been said of patriotism, ‘normal’ is the last refuge of scoundrels in today’s world — or maybe simply the last refuge of dimwits. A good sign that we’re already prepared to enter ‘Humanity 2.0’ is that whenever someone justifies what they’ve done by claiming that it is ‘normal’, we immediately infer that the person hasn’t really thought much about what they’ve done. We hope they follow up that opening gambit with an account that makes clear the existence of plausible alternative courses of action and the basis on which they then took a decision. In the back of my mind is a future conversation, perhaps had toward the end of this century, in which someone asks a person who today would pass as ‘normal’ why hadn’t they opted to take the latest enhancement drug, etc. Simply responding that they wanted to remain ‘normal’ will not cut it then — and barely cuts it now.

Szczechowski

What do you think about the shift in philosophy from the current, popular notion of “being happy with yourself,” to this future of being “always already disabled?” Do you carry around the feeling of being always already disabled now? Do you think this mind-set will help or hinder the overall human condition in terms of satisfaction with life?

Fuller

Most philosophers who have put a premium on a ‘satisfying’ life have had relatively low ambitions for what can be achieved in — and by — one’s life. Indeed, perhaps with a hint of irony, Benjamin Franklin captured the spirit of such philosophers, most of them ancient and pagan, when he remarked that the pessimist is never disappointed and sometimes pleasantly surprised. (In the UK, we have a guy, Alain de Botton, who has made a good living by updating this tradition.) In contrast, the optimist — an Enlightenment coinage — projects a glorious future involving the full realization of our various potentials, brought together by the supremely human faculty of ‘Reason’. Given that premise, there is plenty of room for disappointment, as personal reality fails to live up to the ideal. Transhumanism’s ‘always already disabled’ ethic is the latest version of this mentality, one that converts natural diversity (in, say, intelligence, speed, strength,

etc.) into a social hierarchy, which is in turn converted into a dynamic field as everyone is led to think that the best is within their reach. Both capitalism and socialism, in their rather different ways, have endorsed this vision, in which individual frustration and anxiety serve as the background noise of social progress. The two serious problems with this vision are (a) what economists call the problem of ‘positional goods’ (i.e. one’s own sense of superiority is diminished as others acquire a ‘superior’ good); (b) the lack of a reasonable opt-out from an ever escalating sense of ‘normal’, as most people rush not to be ‘left behind’. My own conclusion is that such problems are definitive of the human condition, which in the introduction to [*The New Sociological Imagination*](#) I explicitly identified as our naturally ‘dissatisfied’ mode of being.

Szczechowski

How far do you feel science can progress humanity before it begins to cause a digression? Or do you feel that humanity will always benefit from the progression of science?

Fuller

Humanity cannot retain its species distinctiveness without letting advances in science and technology drive social progress. I mean this statement to keep open many issues (e.g. who/what counts as ‘human’) but the bottom line is that whenever we demonise or disown science and technology, we are beginning the long march back to the apes. (Just to be clear to Darwinists in the audience: I mean this ‘march’ to be a bad thing!) If this judgement seems harsh or hyperbolic, that is only because we have yet to learn the deep lesson that science and technology teaches — namely, that those who dare, win. But the winners may not be our individual selves but others with whom we identify — either in our lifetimes or the future. Our humanity lies, to a large extent, in just this capacity for intergenerational self-identification.

Lauren Griffith

What metrics/tests (SAT’s, income, which genetic tests specifically) are you basing these claims of categorical inferiority? What reasons do you have to think these tests are an accurate gauge of intelligence? Especially since SAT scores are proportional to the teachers’ mean income in that area and that is determined by the tax bracket of the community around it. This makes poverty heritable and entrenches us in Jim Crow era zoning and resource stratification. What are the metrics you are using to hierarchically rank diverse classifications of the human? What do you think about the idea of “difference without domination”?

Steve Fuller

First of all, I don’t think that any of the current human metrics are adequate to the tasks of gauging the full range of ‘multiple intelligences’ on tap now and in the future. So I agree with your empirical starting point. But I take the spirit of your question to imply that we shouldn’t bother with improving such metrics or that the very development and deployment of such metrics is pernicious. If this is what you mean, then I disagree. Without denying all the horrible things that have been done in the name of ‘eugenics’ (and may well be done in the future, say, under the rubrics of ‘genetic counselling’ or

‘gene therapy’), the whole point of the movement was to break — not reinforce — default patterns of inheritance in society. This meant keeping an eye on two things at once: (1) Whether people from privileged backgrounds were coasting on capital they never earned and lacked the talent to put it in service of the greater good; (2) whether people from poor backgrounds were failing because they had talent but insufficient capital to fully realize their potential to serve the greater good.

The phrase ‘equality of opportunity’ was originally a eugenics slogan to capture this ideal state that would come from discovering the extent to which the social environment made life too easy or too difficult for people, given their innate capacities. Thus, until the rise of Nazism, eugenics was largely a bourgeois socialist movement, whose most well-known supporters were part of the Fabian Society in the UK. The idea would be to reallocate resources accordingly, ending the tyranny of ‘idleness’, be it the idle rich or the idle poor. As for ‘difference without domination’, I believe that a version of this is likely to result in the long term, since it is unlikely that any subset of humanity will acquire the relevant skill-set to ‘dominate’ humanity in some long-term sense akin to authoritarianism. While ‘powerful elites’ may continue to exert disproportional influence, I doubt that their effect will be as determinate as they would hope or others might fear. Perhaps a greater worry in our increasingly networked world is an ‘all dominating all’ scenario that perpetrates a soft totalitarianism, once again more *Brave New World* than *1984*.

Griffith

Do you think it is entirely accurate to claim that, historically, women had power in the home/reproductive sphere? In ancient Greece women were not even seen as contributors to the makeup of the offspring, giving only matter: the flower pot theory. If the father died, the children would be given to the paternal side of the family and classified as orphans even if the mother was still alive. Most women have no choice in family planning. If you think this isn’t still the case, try convicting a man for raping his wife. Once they are married, the husband is legally and culturally entitled to the use of her reproductive system. What information are you basing the male/female market advantage? How can this be divorced from the different ways women and men are perceived in business, rather than their innate proclivities?

Fuller

The short answer to your questions is that the Greeks did not value things quite the same way we do, and so what may look awful to us may not have looked so awful to them — and by ‘them’ I mean both men and women. For example, the fact that Aristotle claims that women contribute only matter to biological reproduction is a big deal to us now because we have a rather joined-up view of how this process fits into larger political-economic processes, not to mention a greater scientific understanding of procreation itself. But all of this ignorance and even prejudice on the part of the Greeks does not deny the fact that they saw women as having quite a strong steer in domestic matters. This is borne out consistently in their dramas, sometimes to ridicule the cluelessness of men who spend so much time arguing in the public square or fighting distant wars that they lose control of the homestead. Of course, I am not saying that women were seen as men’s

equals in the modern political or economic sense. But they were not ‘subjugated’ in the modern sense either.

Brighton Dwinell

What do you find unjust about the administration or utilization of social sciences in determining or framing the human condition?

Steve Fuller

I don’t think that there is anything especially unjust in terms of what you’re asking. The real problem is that social scientists are all too reluctant to push their normative positions in the presence of policy-makers (as opposed to like-minded academic colleagues). What this means is that the few who are willing to push their agendas ([Charles Murray](#) always comes to mind in the US) get a disproportionate amount of media attention.

Dwinell

Some scholars of Jürgen Habermas, namely Nancy Fraser, have rejected the social sciences as an effective method, or finalization, of determining ‘who’ counts, or is represented, under democratic disputes of questions of transnational justice. In what way does your philosophy of sociology and the historical development of the sciences and their public relations speak to the issues of post-Westphalian injustices brought up by these authors?

Fuller

I believe that the real problem here is that the sociological evidence is mixed on whether people actually identify with — or even recognise — the categories of injustice that drive Habermasian indignation. No one doubts that people, especially the poor, have been hit hard by the various ‘crises of capitalism’ and that ‘injustice’ is a reasonable way of characterising their situation. However, it’s not clear that people, even the poor, buy the politically correct (aka Habermasian) explanations and remedies to the situation. In short, the Habermasians need to extend their much vaunted open-mindedness to a test of their own claims — not simply against their own interpretations of current events but against the interpretations of those whose lives they are purportedly interested in saving. Specifically, if they are not open to empirical studies of what ordinary people think and want, they are likely to remain alienated from them. Under the circumstances, it is all too easy to kill the messenger — in this case the social scientists — for conveying the bad news.

Dwinell

What (if any) influence has Heidegger’s *The Question Concerning Technology* had on your analysis and tracking of the developments of humanity 2.0?

Fuller

None, at least on the surface. Heidegger and I operate from quite opposing presumptions. I actually believe that to be human is to be artificial, and that our normative sensibility — i.e. what is right and wrong for us — is ultimately determined by where we want to go, not where we came from. I think that the most interesting connection between Heidegger

and me comes through John Duns Scotus, the subject of Heidegger's PhD yet also a keystone figure for the transhumanist side of Humanity 2.0. Duns Scotus is really the source for the idea that we can talk about 'Being' as such, independently of particular beings, which is perhaps the signature Heideggerian trope. I accept Duns Scotus' own linguistic gloss on this doctrine as 'univocal predication', which implies that if God is 'all-good' and 'all-knowing', as the Bible says, then 'good' and 'knowing' means the same as it does for us, except we've got less of it. Heidegger spins this doctrine rather differently, perhaps assuming (and here he may follow Kant too closely) that the thing from which all linguistic expression comes cannot be itself expressed. But I deal with more specific Heideggerian issues in response to Frodeman and Brown below.

Bob Frodeman and Keith Wayne Brown

The main readings of our class are Heidegger's *Introduction to Metaphysics* (IM) and your book *Humanity 2.0*. Here are two quotes from IM. In your opinion do they say anything interesting about what it means to be human?

For to be human means to be a sayers. Human beings are yes- and no-sayers only because they are, in the ground of their essence, sayers, the sayers. That is their distinction and also their predicament. It distinguishes them from stone, plant, and animal, but also from the gods.

It is not unconditionally necessary that we should be. There is always the possibility that there could be no human beings at all. After all, there was a time when there were no human beings. But strictly speaking, we cannot say there was a time when there were no human beings. At every time, there were and are and will be human beings, because time temporalizes itself only as long as there are human beings. There is no time in which there were no human beings, not because there are human being from all eternity and for all eternity, but because time is not eternity, and time always temporalizes itself only at one time, as human, historical *Dasein*.

Steve Fuller

These quotes certainly say something interesting about what Heidegger thinks it means to be a human, and perhaps his position can be contrasted with mine usefully in this context. In both quotes, the nature of humanity is contrasted with that of God, so the key then is to understand Heidegger's conception of God. Heidegger's deity transcends space-time and exceeds linguistic expression, though it is also the source of these things. Humans, by contrast, are in a strong sense 'constituted' by these things.

In the first quote, Heidegger is alluding to the Old Testament deity as the unnameable namer; in the second, he is alluding to a more Neo-Platonic deity whose 'view from nowhere' doesn't recognise a distinction between space and time, let alone particular locations and particular times: God is always already everywhere at once. (To be honest, I am not sure these two images of the deity sit so well together, but that's for another time.) Humans are by implication defined as a failure to possess these divine qualities. In this respect, Heidegger practices a 'negative anthropology', comparable to 'negative

theology'. Thus, the human is defined as 'abject' (i.e. fallen) being. What makes Heidegger's move interesting is its Scotist signature. Like Duns Scotus, Heidegger writes as if he knows quite a lot about this God from which we have become separated, which results in our existential predicament. But for Heidegger's purposes, it doesn't matter whether God really exists because the essence of humanity lies in our difference from this (quite possibly) non-existent deity: i.e. we are the ones who cannot remain silent and cannot escape temporality. (Here I should put my cards on the table: I first encountered Heidegger through the Jesuits in the 1970s when he was seen as the philosophical heavyweight of Existentialism rather than the culmination of the phenomenological movement, which really only happened in the 1980s, partly due to Derrida's influence.)

Where I depart from Heidegger is that I don't believe that God is ineffable or otherwise incommunicado with our fallen being. On the contrary, I believe that science and technology is how we have begun to re-establish the relevant channels of communication to break through the barrier of temporality and recover that elusive sense of divinity in terms of which humanity is negatively defined. The 'singularity' is certainly one important expression of that aspiration, but the transhumanist movement should be generally seen in this light.

Carl Jacob

What would your response be to the claims made in the [Simulation Argument](#) that proposes that we may never reach a posthuman stage? Namely, the argument's first premise, which says that the fraction of human-level civilizations that reach a post-human stage is very close to zero. But given that the Simulation Argument is probable and that its first premise is true, what do you think is the likelihood that our civilization will reach a post-human/ trans-human stage and not destroy itself?

Steve Fuller

I'm assuming you mean 'transhuman' rather than 'posthuman', since we could become 'post-human' simply by reintegrating with nature and diminishing our distinctly human features – a certain sort of radical Green paradise. (I realize that Nick Bostrom wrote this piece before the trans/post-human distinction became canonised.) But more to the point, I find it very unlikely that the first premise is true. We certainly had many opportunities to test it during the Cold War, which was after all promoted as an 'arms race' without end!

The most plausible reason I can see for believing in premise one is an overriding belief in Darwinism (i.e. the inevitability of species extinction) combined with a heightened sense of how we might be brought low by unintended consequences. In contrast, I take the Simulation argument to be a high-tech version of the argument from intelligent design for God's existence (aka it is likely that we inhabit someone else's simulation). So I buy premise three of the argument. If secular society were less militantly anti-religious, we might be able to have a sensible conversation about this matter.

Jacob

You say in Chapter 2 of *Humanity 2.0* that the definition of 'human' is fugitive. In the

future do you think the distinction between post/pre-human would prove untenable or end up being synonymous with each other?

Fuller

No. I think the ‘human’ will always remain as a normative standard that exceeds who or what happens to qualify at the moment as human. The key difference in the future will be that ‘human’ will not be the exclusive domain of *Homo sapiens* – other beings, including animals and machines are likely to count as well. However, this development should not prove so surprising, since it has taken quite a long time for there to be general agreement that all *Homo sapiens* should be counted as ‘human’.

Robert Frodeman

I find the disciplinary aspect of the 2002 [National Science Foundation] report [that established a research agenda to promote ‘converging technologies to enhance human performance’] interesting, that the two authors of the report were an engineer and a sociologist. What evidence did you find of them taking up a philosophical or humanistic perspective on their goals?

Steve Fuller

I think this is true in a very general sense. The main strategic goal was to provide a broad justification for the public funding of basic research in the post-Cold War era, where there was a general recognition that science had to demonstrate its value more explicitly, given the removal of ‘mutually assured destruction’ as a serious prospect. Here it is worth recalling that the ‘Cold War’ was imagined as the infinite deferral of Armageddon. Thus, we had to think ahead to pre-empt our enemies. This was a boon to the more abstract sciences, which led to great advances in physics, computer science, decision theory, etc. But post-1989, such sciences were easy targets for overstretched fiscal budgets, and so the 2002 report was designed to reassert their prominence in a context where the public is disinclined to support basic research unless there is some personal benefit (i.e. beyond global security) on the horizon.

This strategy had worked before, as the Rockefeller Foundation induced people from the physical sciences to enter biology to invent ‘molecular biology’ (a Rockefeller coinage) starting in the 1930s. The 2002 NSF report aimed to inject much the same skill-set into a high-tech version of the Rockefeller vision. Now if you step back from the strategic goals, it’s easy to see a shared vision here of humanity as a being who develops by leveraging science and technology into greater dominion over the planet, each other and oneself. From today’s standpoint, what is perhaps most striking about the 2002 report is its resolutely anthropocentric perspective. There is no serious discussion of the future of nature.

David Silverberg

Do you see Kurzweil’s prediction of future biotechnology coming to fruition with little to no political repercussion if it is done more gradually over time and not at the accelerated rate that Kurzweil suggests? Why or why not?

Steve Fuller

Yes, basically. Moore's law, on the basis of which Kurzweil thinks we're accelerating to the singularity, is really only about the rate at which available computational capacity increases. It doesn't address the rate at which these increases are assimilated by society — and Kurzweil's prediction is momentous only because it makes claims about the permeation of super-computing power across all of society and even the cosmos. And while it is true that the history of technology shows that as an innovation spreads, there is a radical drop in its unit cost, which in turn serves to expedite its spread, there's no reason to think that this process conforms to Moore's law. So our convergence with high-end computers is not likely to happen as soon as Kurzweil thinks, but we're heading there nonetheless. In the meanwhile, we'll have time to get used to it — and we already are, but less through explicit instruction than the propaganda stream provided by mass media stories, science-fiction films and video games.

Mike Watson

Is it alarmist to think that the technological singularity will lead to mass extermination of large percentages of the population? Isn't it possible that the people in control of this technology have it all wrong, and the AI systems they are using against the 99% might become self-aware and enslave everyone, including the 1%?

Steve Fuller

A shorter term bet to exterminate large portions of humanity is a series of catastrophes brought on by climate change that disproportionately affect the world's poor. But more to the point: A more immediate, pre-singularity threat posed by AI is a system that works against the interests of the client humans without adequate warning. For example, most of the transactions that precipitated the recent global financial collapse were made by machines that communicated with each other according to algorithms that their notional masters did not understand and hence were in no position to control. This is not a problem of malice on the part of the computers but ignorance on the part of the humans.

Watson

Do you believe that government programs such as the NSA are part of the singularity and the study of human thought patterns? Is the state part of a massive conspiracy to enslave humanity for their own selfish gain? What hope do we have in light of recent revelations that the government spies on basically every moment of our lives through our smart phones?

Fuller

Even if this were likely to be the case, it's not clear what would count as 'selfish gain' for, say, NSA that did not also involve some substantial perceived benefit for humanity. If there is a dystopia on the horizon, it is more likely to be like Aldous Huxley's *Brave New World* than George Orwell's *1984*. In other words, the problem will be that people too effortlessly buy into the world-view of the 'masters' and so see themselves as part of their team, not that there will be a growing self-consciousness of an asymmetrical 'master-slave' relation. Thus, the fact that that the NSA shares its surveillance data with

the commercial sector – without public accountability — is of greater concern than the sheer fact that they engage in mass surveillance. In the end, of course, people may vote to enter Brave New World, but they still should be given the vote, even if it would be more efficient not to. I say ‘should’ because I believe that a bottom-line condition for humanity is autonomy, which means that the people themselves need to say yes.

Watson

Is it possible that if AI supercomputers controlled society, then government would become obsolete and we would live in anarcho-capitalist “utopia”?

Fuller

That’s a very astute question and certainly there’s an emerging form of corporate life — the ‘decentralised autonomous corporation’, or DAC — that reflects this prospect becoming the dominant reality. [DACs are the brainchild of Bitcoin visionary Vitalik Buterin](#). He basically sees the future of wealth creation in terms of the algorithmic systems that trade largely amongst themselves, with relatively little human intervention, which of course resulted in the stock market crashes of 1987 and 2008.

The general attitude of the designers of these systems is to treat the crashes as learning experiences for building better, or at least more resilient systems. And unless you expect radical de-complexification of the economy in the foreseeable future – that is, a ‘small is beautiful’ scenario — DACs are here to stay, and the only question becomes the pervasiveness of their reach. However, I think the firm — not the state — is most directly threatened by DACs. It is striking the extent to which states have remained intact – and even proliferated — throughout the fluctuations in the world’s economy since 1945. True, their powers to provide for their citizens may have been severely dented since the end of the Cold War, but in return states are now awash in an unprecedented amount of data about those citizens.

What Orwell never anticipated was that this increased surveillance capacity might not very easily translate into mechanisms of control. People who nowadays make a big deal about privacy should keep that in mind. (In this respect, Kafka might be a better literary guide.) In contrast, the modern economic conception of the firm seems to be breaking down. Firms have been classically seen as efficient solutions to the problem of transaction costs, namely, the costs involved in getting you to a satisfactory exchange, most of which are informational (i.e. how to find who you can trust to give you what you want). Firms basically internalize these costs, resulting in their specific form of corporate identity. Thus, the people who develop, produce and market a line of goods are part of the same team rather than separate providers requiring separate negotiations. But what if this strategy proves not to be efficient? After all, the creation and extinction of firms is becoming marked these days, perhaps reflecting that information technology has made it increasingly easier to negotiate transactions on the spot without the baggage of a common institutional identity. How states ride this particular bronco will be especially interesting to see in the future, but I’m very confident that the United States will outlast General Motors.

Watson

Is the world headed towards global governance? What are your thoughts on policy

institutes such as The Bilderberg Group, The Council on Foreign Relations, or The Club of Rome? Do they control the world and its resources?

Fuller

First, let me say that the policy institutes you mention may have impact but in the manner of a meteor, where the effect of their intervention may be big but always indeterminate. Getting rid of these institutes would remove a factor from the mix of things that determine our future, but it's not clear that it would be for better or worse.

On the more general point: I don't think we're heading toward global governance, if we continue to see nation-states as the units of governance, a la United Nations. Without denying the UN's often misunderstood and underestimated achievements (especially vis-à-vis the old League of Nations), it still presupposes that global governance occurs through the coordination of interests of the world's primary political entities into a kind of second-order 'superstate' that ideally carries the moral authority of all of humanity. The weakness of this approach is shown in military contexts, when nations exercise vetoes or act unilaterally. In this respect, the UN can never really be a whole greater than the sum of its parts. So, we may need a global governance agency that is organized in terms that cut across nation-states, if only to regularly force a reconceptualization of human interests. This supranational body may be organized along, say, class lines, as long as 'class' refers to a relation to the means of production. But there are other possibilities, including ones involving genetic differences (aka race) — but this would require another discussion, as I see contemporary biology updating the idea of racial differences but on a much more functional basis, akin to arguments for a division of labour in modern society, which implies that the job one does in society need not simply reproduce that done by one's parents nor determine the career path of offspring. In short, we would need to measure baseline competence in each generation to ensure that society's current division of labour is not mindlessly reproduced.

Watson

Do you believe that advanced AI supercomputers will eventually replace the elites? Is Ray Kurzweil insane for thinking that machines will allow him to become some sort of techno-god? http://en.wikipedia.org/wiki/Georgia_Guidestones

Fuller

First, as regards the Georgia Guidestones, everyone knows that it was a Jimmy Carter legacy project that began when he realized that his days in the White House were numbered, yet he needed some way to subliminally seduce the Nobel Peace Prize committee. ;) But seriously, folks: It's worth observing that Kurzweil is not alone in predicting a 'singularity'. The American founder of the 'anthropic cosmological principle' (roughly, the idea that the physical conditions that have enabled humans to observe the universe in its entirety are luminous about the nature of the universe more generally), [Frank Tipler, supports the prospect of a singularity on both scientific and theological grounds \(as a committed Christian\)](#). I take Tipler quite seriously, but however you judge the viability of the singularity, you should understand it as a kind of merging of minds, not the dominance of one mind over the rest.

Watson

Do you believe mainstream medicine will start to introduce cybernetic implants to the general population within the next 100 years? Will the government use brain implants to control human thought in service of the state? Isn't it true that foreign governments could hack into everyone's brain chip to control their thoughts during a war? Shouldn't all discussion of implanting the brain with technology be looked at with skepticism, if someone has access to your emotions, memories, motives, etc.? Is there any way this technology could be used honestly and ethically?

Fuller

I would be surprised if cybernetic implants didn't start to be normalised in the current generation. However, I don't think that the danger is any greater than making new drugs generally available. Here too are opportunities to create unwanted dependencies, often involving people revealing more about themselves than they might otherwise. However, even when imagining brain chips, one needs to distinguish between getting people to do things they might not want to do and getting them to do what you want. The former is always easier than the latter. If we should be especially worried (and I'm not completely persuaded), it should be about cybernetically enhanced humans operating in ways that are in no one's interest, not that they might be turned to specifically evil purposes.

Watson

Isn't it imperative that society uses the non-aggression principle (that states aggression towards another is immoral, eliminating the state completely), if we are to ethically apply this technology?

Fuller

I don't see how eliminating the state eliminates aggression, and I don't see aggression as necessarily a bad thing. Perhaps the best way to prepare ethically for a future where advanced technology is attributed significant agency — and hence 'rights' — is to study how the law has had to change to accommodate expanding sense of citizenship, say, with the enfranchisement of former slaves, women, etc. I think classical liberal notions of tolerance will come under great strain, given that such notions have been predicated on potential citizens — however else they differ — being members of *Homo sapiens*.

Contact details: S.W.Fuller@warwick.ac.uk