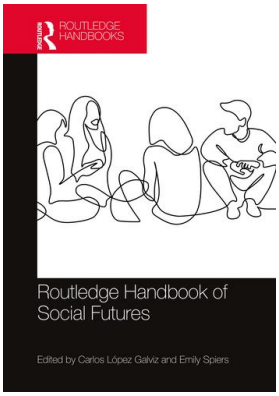

Routledge Handbook of Social Futures

Carlos López Galviz, Emily Spiers, Monika Büscher, Astrid Nordin

AI

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9780429440717-4
Richard H.R. Harper
Published online on: 23 Nov 2021

3

AI

The social future of intelligence

Richard H.R. Harper

Introduction

The past few years have shown a society-wide interest in the remarkable developments within machine learning and associated techniques that are enabling what has come to be called the *New AI*. This technology is not just altering our present but many say will alter our futures too. In this view, AI will come to supplement and even substitute human reasoning, with its powers being amply demonstrated in the capacity of AI machines to beat humans at even the most complex rule-based activities, such as the game of *Go*. AI will also come to be at the heart of self-driving cars and will populate human-less factories, and the 'face' of service industries will be artificial assistants.[1]

The benefits that are seen in these prospects are, of course, immense. But so are the concerns. If AI can do more work, will that mean unemployment for the humans who used to be required to do that work, for instance (Carr, 2015)? In the long run, what will be the effect on human dignity if work is no longer the central currency of identity (Markoff, 2015)? If AI is more efficient, what will be the measures used to judge investment? Will AI itself choose where money should go (Kaplan, 2015)? More philosophically, if AI is able to reason more effectively than people, what will be the future of learning and further education? Why should society invest in people if AI is a better learner? Ultimately, will it be AI that does science and wins Nobel Prizes for doing so (Kitano, 2016)?

Much of these claims are hyperbole, and some are simply overexcited. The result of all this is that the true role of AI in the future is unclear, the hyperbole surrounding it is making careful analysis of its potential hard and the full range of consequences that follow on from what the technology might provide remain, in many respects, unexamined. The future of AI, how it affects not only how computers function, but also what those computers can do, and how, in turn, this alters their role in society more generally, is largely muddled territory.[2]

In this chapter, I want to explore what this future might be from a particular view: From a concern with the future itself, and with the ways one might construct a vision of that future that might tell us how we might alter our *current practices* to make a future we want more likely to happen. This is part and parcel of what has come to be called *Futures Thinking*, of

 DOI: 10.4324/9780429440717-4

course, and hence this book. There are various techniques in *Futures Thinking*, needless to say, but the one I will emphasize entails examining the assumptions that lie hidden in our thinking about the future, and seeing if such an examination may alter them in some way that enables the assembly of a more realistic path to the future we want. In this case, I will want to see if we can unfold a *social future* for intelligence, and not a future that we tend to assume will unfold with standard accounts of AI. These, I will say, emphasize a technological or computational view of intelligence and, as they do, narrow what intelligence might be in the future. I will seek a broader view, a social view, pointing to how intelligence is at once diverse in the way the concept (or label) is used to describe and account for phenomena, and, similarly, when a variant of the concept is used to define and explore the 'method' of it. The AI view, if I can call it thus, is narrower in its scope, constraining what is thought of as intelligent and constraining, as a result, how we imagine ourselves, our endeavours and our tools, intelligent or otherwise. The AI view traps the concept, killing the diverse life it might have and making it singular.

In the first section, I will explore what is said to be intelligence from the 'standard' point of view in AI, the one that is most often marshalled in relevant discourses.[3] This construal points towards how intelligence is 'done', its mechanics if you like, but this comes at the price of occluding larger questions about the purposes that intelligence can be put to. In this view, one is tacitly encouraged to look at similar *mechanisms*, not at *dissimilar purposes*. As a result, we tend to make the future more narrowly than it could be, and indeed ought to be, since we come to be obsessed with means and not purposes.[4] Means and mechanism come to matter more than purposes.

In the next section, I will tease this out by focussing on how one understands a game, a game of football. Here, I will draw a distinction between the mechanics of the human body (crucial to playing the game) which are the *means of football* from the purposes of the game, which are, as any reader will know, to score goals. Winning is the goal in football, one might say in delight at the tautology. My point will be that seeking this purpose intelligently is an altogether different thing from any intelligence 'in the body' used to play the game. I will not argue that one usage is better than another. They do different things. There is, if you like, a social dimension to the use of the concept of intelligence, letting people do various kinds of work when they use it – sometimes to describe ends, sometimes means, sometimes both. This will let me conclude with the claim that we do not want to be constrained in our understanding by the machineries that enable 'intelligence' when intelligence can be delivered in all sorts of ways, for all sorts of ends. For the ends can define intelligence too, and indeed in some cases, it is those that matter – as they do in football. I shall go on to say that such ends can be many; one doesn't only need to think of football. In short, intelligence can and does have many forms; it is a rich, multivalent feature of our lives, and accordingly, its use as a concept is equally so.[5] Given this, we need to ensure that our future will not be made smaller by the machineries with which we enable intelligence; our use of the concept intelligence should not be made less through being limited to reference to them (i.e., the machineries). Our conceptual élan should not be constrained by technological types – such as the ones 'inside' AI. We can avoid this by looking at the assumptions we make about AI at the current time, ones which deliver that narrowness almost without thinking by constraining us to how AI is itself done. Being alert to that we can defend our thinking, and so help create the future we want where this narrowness does not confine us. We want to socialize the use of the concept, intelligence, and thereby let us do more things with intelligence. Thus, *Futures Thinking*.

From Norfolk to Silicon Valley

Most commentaries on the meaning of intelligence, when used in combination with the word artificial, point towards ways of calculating. They do not mean, for want of an everyday analogy, the kinds of calculations that an abacus can do (a kind of computer), though, as they are thinking of the kinds of calculating that is entailed when people play *tightly ruled games* – *Go*, as mentioned. In *Go*, players calculate different outcomes given different choices to determine ways forward; they weigh strategies and choose the 'heaviest' (so to say) to succeed in the game. AI machineries, at least the ones that go under the name the *New AI*, entail various forms of machine and deep learning, but all function in a similar way. Elaborate statistical techniques, most often *Bayesian* (after the Norfolk vicar who invented them in the eighteenth century), are combined with new computer architectures in such a fashion that AI computers can win at games, even at *Go*.

That AI systems can do this, beat the human in these games, has led some commentators to the conclusion that this is what intelligence 'is' – beating humans at a game means passing the so-called Turing test, the measure of whether a machine is as intelligent as a person. Whatever one thinks about such measures or tests, in this view, anything that has rule-like, game-like behaviours can be seen as a form of intelligence. From this starting place, commentators have argued that many activities can not only be thought of as game-like but *are* game like. And the proof of this is to be found in the apparent success of AI systems to 'win' in these other activities.

For example, AI can 'see' things in the visual field and does so in a game-like manner. The AI uses probabilistic techniques to interrogate data it gets from its digital cameras to distinguish shapes and labels those shapes (or objects) in terms of classes. The classes can be many things, including persons. However clever these techniques and however startling the power of the computers to label one shape over another (and hence one person from another), all they are doing, in effect, is treating that task of identification as one that can be *calculated in the manner of game-play*. The game entails subtracting, subdividing and combining shapes in data sets in ways determined by elaborate, game-like rules related to the task of recognition; success occurs when the classes match a known shape, a class in the register. This architecture presupposes what machines are to look for (i.e., classes), how they might do this and how they might know when they have seen the things (classes) their calculations are designed to recognize – adequate distinctions between John, Fred, Harry, Sandra and Carolina who are queuing up at the passport gates and being seen by the computer system at the same time. One might note that to see, in this view, is not to *know* that it is Harry or Sandra or whoever; recognition is not familiarity, a cue to say 'Hello!'; on the contrary, it is to behave like a *Go* player making one play rather than another; there is no interest in what is seen or why it is seen. The goal is to win, when in this case, to win is to *recognize the right face*.

As I say, many commentators have started to argue that this is what intelligence is *tout court*. Just as AI machines work this way, so must other 'machines', they think, even biological ones. Some have argued, for example, that mechanisms inside the human body are to be thought of as behaving in this sort of way, calculatively, probabilistically, with rules guiding their decision-making as in a game. When a cell confronts another, in this view, its reaction is determined by probability and game-like rules – the cell plays stratagems, so to say. This vision is used to explain how 'communication' between and across cells occurs, and ultimately within any system of cells. From this, these commentators come to assert that the human 'mind', consciousness in particular, similarly emerges; it is the outcome of a vast, intricate system of probabilistically calculated stratagems.

In the body and the mind (the latter used interchangeably with the word brain), these calculations are undertaken by enzymes and chemical processes, whereas with an AI machine these calculations are done by logical gates carved in silicon by light. But those who hold this view think of both as more or less the same – the machine and the body/mind. The material of the 'machine' in question is irrelevant; 'doing' intelligent activities in this way is common to one and all; these doings *are* intelligence. In essence, this is the argument that gets called *singularity*.[6] Intelligence, consciousness, choice-making – all these have common roots or rather similar machineries. If there is a measure of intelligence, it relates to this. In sum, we are no different from AI machines, they no different from us, or soon won't be.[7]

From Silicon Valley to Wembley Stadium: A case of conceptual analysis

Those on the outside of this particular view, those not taken with the AI perspective, dispute this. Their dispute is not about whether, say, cells react probabilistically in a rule-governed way. It is whether one can say the same about a person calculating in rule-governed ways; in their view, one might say that both are doing intelligent things, but what is meant in each is different. Those who don't hold with the singularity argument would say that what is meant in each case is *not the same*. To explain that a body and/or machine behaves this way is to account for the *outcomes* of the functioning of the body and/or machine, it is not to say that either do, in and of themselves, *choose* (if one can summarize this view with one word). It is to say that their functioning *can be thought* of as being like that. Whether they choose or not is largely irrelevant, but, in contrast, *that a person* chooses is the mark of whether they are intelligent or not.[8] In short, they choose to; one cannot say the same of cells or the systems they are part of. They choose nothing.[9]

This might seem merely a question of words, of conceptual distinctions that seem minor. But there are important issues here to do with what is meant by intelligence and its nested use as a concept with other concepts like choice, and what we mean by the choices that people and machines might make and hence what we mean in turn by intelligence. They might not mean the same thing in each case. There might be a distinct form of use for each, a social life for the concepts in question that is variable.

At the current time, the accounts of intelligence that seem to dominate are of one particular kind, emphasizing a particular view about intelligence. This view comes to be assumed, as if it is the only way one might think of it. Are we narrowing our notions to just one view? I think we are.

A simple example can help us here. One might say that there are two ways of looking at the game of soccer. One looks at the way the muscles of a player function when they play. Another looks at the game itself, at the strategies and skills used to win. In my view, those who hold with the singularity argument, the view from Silicon Valley if you like, are looking at intelligence in the former view, that is to say like those who might look at soccer and see muscles acting. To be sure, the mechanisms of the muscles can be explained as being a function of choice, of the physiology of the body operating like an intelligent machine, optimizing this movement to achieve this speed of flex, altering the posture to ensure this kind of poise and balance. Though it is true that muscles need to flex and are controlled inside the body in a way that can be thought of as calculative and hence intelligent, to understand soccer it is better, in my view, to *look at the game*, not at the muscle movements of any of the players, however much one might say that those movements articulate intelligence in the body itself. It is *in* the game, so to speak, that intelligence is to be found – in how people play it, *not in the muscles*. Looking at the game is the preferred view from inside Wembley Stadium,

if you like, the perspective that will concern the crowd and the 'bench'. To see it this way, in the Wembley way, AI notions about intelligence simply don't help. Indeed, they can make it hard to see – they can take you away from what you need to look at. Instead of letting one see *play,* the AI perspective encourages one to look at *muscles*.[10]

This does not mean that the assumption that leads one to see intelligence in the machine (so to say) is always wrong; on the contrary, what I propose is a bigger picture of what intelligence might be that includes both *intelligent muscles* and *intelligent play*, both the skills of the game and the processing of well-trained muscles. But by that I am not proposing to bring them together in some über definition. On the contrary, their value is in their conceptual distinctiveness, and the different work they do.

The forms of intelligence

My thesis is deceptively simple: It is that 'intelligence' can have many forms and that, given this, we need to be alert to the places and purposes in which those differences can be found so that we can ensure that our future will be made bigger by the impact of AI rather than smaller. We can achieve this by looking at the assumptions we make about AI at the current time, ones which can deliver that narrowness almost without thinking. Being alert to that we can alter our thinking, and so help create the future we want.

In particular, I have argued that narratives about AI can displace sensible discussion about what we want those applications to do. We end up thinking about how AI works (and to how explain *that*) instead of exploring what AI might help *us do*. We are offered *muscles* and explanations about how they work, I suggested in metaphor, and not games and their various *purposes*. As a result, we lose sight of why we might want to play games, of the purposes that would make doing so worthwhile.

Games are of course only one of the many things we can do, and, as I alluded to above, games are in many ways the opposite of intelligent. We can play them to get away from our minds; we play them since we don't have to think. And yet, to say again, we can sometimes play a game intelligently: Thus, the complexity of conceptual tools – of our language, our ways of organizing ourselves in and through words.[11]

I want to end, though, with a different set of concerns. It has to do with how the development of AI is itself constrained by this predilection for games. That this might be so is being noticed by futures thinking researchers. In '*The future of human-artificial intelligence*', Spelda and Stritecky (2020) argue that the social organization of how AI engineers are developing their technologies, one that focusses research effort on particular commonly agreed game-like tests (identifying this animal or that syntactical structure in a dialogue), has the advantage of allowing researchers to compare each other (and makes giving awards easier too), but is resulting in the AI community narrowing the cognitive scope of the machineries it is developing. In contrast to *Turing theoretic* machines,[12] the model that has underscored the PC revolution and the one which emphasizes generalizability of processing and hence also the general properties of cognition enabled with or through such machines, the AI machines being developed are providing ever more narrowly defined cognitive solutions. This is reflected in their architectures, their algorithms and their outputs. As it happens, Spelda and Stritecky are concerned that the value of these narrow solutions is disproportionate to the environmental impact of using the technology, the energy cost of deep learning being especially high indeed. That may well be right and is certainly something I am researching even now (see also Yeung et al., 2020). But here I take their paper as confirming my thesis: That AI narrows what we want our 'intelligences' to be applied to.

If we want an intelligent future, we need to allow intelligence to have what I have called a social life, one where its use as a concept is rich, nuanced and enabling, a network of (social) possibilities. If we look at the future from the view of AI, that is not a future we shall see or will end up making. Accordingly, to see the future we want and to help make that future with it in mind, one must not look from that perspective. That has been my contention. One must look from a bigger vantage point, I have said, one that sees the complexity of intelligence in all its forms: From being 'in' muscles to being 'in' stories; from being a 'feature of' games and a 'property of' intentions; and in being at once a way of doing things (a machinery) and a label for a thing well done.

Notes

1 The literature on this is immense and enormously varied. I do not seek to offer a literature review of it all but will point towards what seem to be representative contemporary examples at appropriate stages of my argument. But for a good introduction to the many points of view that is not partisan see, Kaplan (2015, 2016); also Boden (1977, 2016). In relation to games like *Go*, see Sadler and Regan (2019).
2 This has been a persistent problem. For example, Stanford University sought to bring clarity to this space with its *AI and Life in 2030* report written in 2015 (published the next year). The muddles it cites are very similar to those I list here, some years later.
3 One could digress and define the range or boundaries of such discourses, but for the purposes here they are simply those that focus on AI as the primary concern and treat AI as a given.
4 That this is so affects all sorts of attempts to explore what AI can do. Some of the better studies from, for example, the social perspective work their way around these mystifications before they can find out what the technology does in the real world and its consequences when seen from the social view. See Neyland (2019).
5 This is a point I take from Wittgenstein (1953).
6 This was originally formulated by John Van Neumann but has been popularized by Kurzweil (2005). But see Stanislaw (1958).
7 So, from this view, while we might think of ourselves as singular – that is to say you and I might like to think of ourselves as such, that our minds are ours and ours alone – in fact, if one believes this view, our consciousness is the outcome of millions of little acts, little calculations and stratagems at the cellular (and system) level that produces this sense of self. Our sense of that self is now seen to be egregious. This is the view that Dennett argued for in his *Consciousness Explained* (1991).
8 This is the rub of Kant, of course: personal accountability. In Kant, the self is a transcendental, a premise that cannot be avoided in understanding human nature. Kant is difficult to read, and often obscure, but the best introduction to this fundamental concept is in Scruton (2001).
9 This is most eloquently expressed by the physicist, R. Jones, in his (2004) *Soft Machines* – a much better book than Dennett's in my view, since it explores the consequence of this important distinction – the one between description of activities and action that is governed by self-awareness. For those interested in exploring this line of argument, they should go back to Anscombe's *Intention* (1957) which explains how motives distinguish human action. In this view, a machine cannot have a motive, though it might have 'reasons for doing what it does' – such as probabilistic reasons. But for an introduction see Harper et al., *Choice* (2018).
10 This is of course an argument that derives from the ordinary language philosophers, Wittgenstein (*op cit*) being the most regarded, if not the easiest to read.
11 Of course, this is brutally expressed, this being the conclusion. It is, needless to say, the substance of Wittgenstein's later inquiries (*op cit*, 1953): viz, not that we are *only* to be understood in words, but the relationship between our doings and ways of accounting is subtle and fundamental to how we understand.
12 A Turing Machine is a computer that has been designed to support generalized functions rather than specific, unique ones. A PC is a perfect example – a computer for general purposes. The reason why this is an important concept is because it has turns out that algorithms can be developed that allow generalized hardware to solve particular problems, and algorithms are much easier to make than hardware.

Bibliography

Anscombe, G.E.M. (1957). *Intention*. Cambridge, MA: Harvard University Press.
Boden, M.A. (1977). *Artificial intelligence and natural man*. New York: Basic Books.
———.(2016). *AI: Its nature and future*. Oxford: Oxford University Press.
Carr, N. (2015). *The glass cage: Who needs humans anyway?* London: Vintage Books.
Dennett, D. (1991). *Consciousness explained*. London: The Penguin Press.
Harper, R., Randall, D. and Sharrock, W. (2016). *Choice*. Cambridge: Polity Press.
Kaplan, J. (2015). *Humans need not apply: A guide to health and work in the age of artificial intelligence*. New Haven, CT: Yale University Press.
———. (2016). *Artificial intelligence: What everyone needs to know*. Oxford: Oxford University Press.
Jones, R. (2004). *Soft machines: Nanotechnology and life*. Oxford: Oxford University Press.
Kitano, H. (2016). 'Artificial intelligence to win the Nobel prize', *AI Magazine*, Spring, pp. 39–49.
Kurzweil, R. (2005). *The singularity is near*. London: Penguin.
Markoff, J. (2015). *Machines of loving grace*. New York: Harper Collins.
Neyland, D. (2019). *The everyday life of an algorithm*. London: Palgrave Macmillan.
Sadler, M. and Regan, N. (2019). 'Game changer: AlphaZero's ground-breaking chess strategies and the promise of AI', *New Chess*, The Netherlands.
Sartre, J.P. (1956). *Being and nothingness: An essay on phenomenological ontology*, tr. H. Barnes. New York: Routledge Classics.
Scruton, R. (2001). *Kant: A very short introduction*. Oxford: OUP.
———. (2006). *Sexual desire: A philosophical investigation*. London and New York: Continuum.
Shanker, S. (1998). *Wittgenstein's remarks on the foundations of AI*. London: Taylor and Francis.
Stanford University (2016). *Artificial intelligence and life in 2030*. Stanford University, https://ai100.Stanford.edu
Spelda, P. and Stritecky, V. (2020). 'The future of human-artificial intelligence nexus and its environment costs', *Futures*, https://doi.org/10.1016/j.futures.2020.102531
Stanislaw, U. (1958). 'Tribute to John Van Neumann', *Bulletin of the American Mathematical Society*, 64 (3), part 2: 5.
Wittgenstein, L. (1953/2009). *Philosophical investigations*, 4th edition, tr. G.E.M. Anscombe, P.M.S. Hacker and J. Schulte. Oxford: Wiley-Blackwell.
Yeung, G., Borowiec, D, Friday, A., Harper, R. and Garraghan, P. (2020). 'Towards GPU utilization prediction for cloud deep learning', in *HotCloud20: 12th USENIX workshop on Hot Topics in Cloud Computing*. Boston, MA: ACM Press. Available at: https://eprints.lancs.ac.uk/id/eprint/144290/1/HotCloud20_GPU_Prediction.pdf