Robots and Anthropomorphism

Margaret A. Boden

University of Sussex

Abstract

Our natural tendency to anthropomorphism, grounded in Theory of Mind and related psychological mechanisms, is crucial to our interactions with robots. Some relatively superficial aspects of robots (e.g. physical appearance) can trigger animistic, even empathetic, responses on the part of human beings. Other factors are more subtle, e.g. various aspects of the language (if any) used by the artifice, and/or of the thinking-processes apparently going on. Robotics (like AI in general) promises/threatens to alter how people think about themselves. Unlike AI programs, robots are physical entities moving around in a physical world. This makes them more humanlike in various ways. But physicality isn't the same thing as embodiment. For someone who wants to insist on a distinction between robots and humans, the fact that robots aren't living things is likely to be important.

Introduction

Nearly a quarter-century ago, the social psychologist Neil Frude (1983) predicted that humanlike robot "companions" would attract deep responses, even including emotional attachments, because of the universal human tendency to animism. (He said much the same about 'human' figures in the type of audiovisual display which is now called virtual reality; however, I'll ignore screen-based VR systems here.)

By animism, he meant ascribing psychological predicates to things that don't literally merit them. Strictly, these things include both living organisms and non-living things. Certain animistic religions, for instance, feature trees, snakes, winds, or rivers as having intentions (and knowledge) of simple kinds. But Frude was especially interested in cases where the predicates involved cover all or most of the mental characteristics we ascribe to human beings, and where the target of the animism actually looks and/or behaves in some ways like a human being. In short, he was thinking of the strong form of animism known as anthropomorphism, and of specifically 'humanlike' targets such as dogs, dolls, and teddy bears--and, crucially, robots.

A walking, talking, cuddly robot--perhaps even one bearing a physical resemblance to one's favourite film star, and speaking in something like his/her voice--would (he said) inevitably, irresistibly, attract our interest and empathy. If it could detect emotions in its human owner/partner, either verbally and/or by means of facial expressions, or express 'emotions' in one or both of those ways on its own behalf, the effect would be even more compelling.

By "compelling", here, Frude meant psychologically, not philosophically, compelling. I'll follow him in ignoring the philosophical question of whether a robot could 'really' have psychological properties. I'll assume (as explained in Section II) that the robots of the foreseeable future will be psychologically crude in comparison to their human users: no Proustian subtleties, there. But I shan't ask whether, in principle, any conceivable robot could really experience emotions (even relatively simple ones), or really understand language (even of a primitive kind).

Rather, my questions here are (i) whether people interacting with humanoid robots would, in fact, spontaneously respond to them as if that were the case, and (ii) whether their ideas about themselves and other human beings would be altered as a result. Those questions are answered in Sections III and IV, respectively. First, we must ask whether there's any chance of Frude's imagined robots ever seeing the light of day.

II: Future, fiction, or fantasy?

When Frude wrote his book, it was science fiction. To be sure, he himself thought of it as prediction, not fiction. But he vastly underestimated the theoretical and technological difficulties involved in bringing his predictions to pass.

Admittedly, AI/robotics--nearly a quarter-century later--has made a start in all the relevant dimensions, including some he didn't even mention (Boden 2006: 13.vi.d). These include--among other things--NLP, speech-processing (wherein computer-voices can be given appropriate local accents: Fitt and Isard 1999), visual recognition of facial expressions, modelling of emotion (e.g. different types of anxiety, depending on the urgency, importance, and priority-orderings of the goals involved: Wright, Sloman, and Beaudoin 1996), and the use of flexible skin like materials on the robot's 'face'--with an underlying 'musculature' capable of mimicking human expressions. In principle, although this too is science fiction at present, a competent visual face-recognizer might even detect when the human was lying (Ekman 2001). In other words, Frude's futuristic vision wasn't pure fantasy.

But that's not to say that the highly humanoid robots he envisaged are just around the corner. In my opinion, indeed, they never will be. The difficulty of achieving human-level NLP, for example, is just too great.

Frude imagined AI-companions that could converse sensibly with their human owners (surely not 'partners'?) about a huge range of topics, just as your colleagues and next-door neighbours can. NLP systems are already being used by some large corporations to stand in for humans as 'customer care' agents. But they exploit a limited menu of stock-phrases concerning a limited set of topics. Fully 'Frude-ian' programs would be quite another matter.

Certainly, we could forgive an AI system for not being to talk about quantum mechanics, or even football: after all, you, or your neighbour, may not be able to do that either. But the conversational limitations of individual human beings can't be used in Frude's defence here (as human mistakes can be used to counter anti-AI arguments based on Godel's theorem). For even someone knowing next to nothing about these specialist domains could make some attempt at conversation about them, if only to crack a joke in order to change the subject.

Changing the subject, of course, requires both discussants to track the focus of attention. This involves, among other things, interpreting anaphora correctly: every occurrence of "it", for example. This is a subtle matter, which people achieve mostly unawares. NLP research indicates, in my judgment, that an AI system that can keep track of the constantly shifting focus of everyday conversations isn't feasible in the general case (Grosz 1977; Grosz and Sidner 1979; McKeown 1985.) Part of the problem, here, is that indefinitely various world-knowledge would be needed to cope with every occurrence of "it" (physics and football, again).

In short, a robot that could talk to you intelligently, never mind humorously, about virtually anything simply isn't on the cards. (And a good thing too: shivers run up my spine when I consider Frude's vision of generations of lonely old people, deserted by friends and children, interacting only with their domesticated robots.)

It doesn't follow, however, that a robot of the future couldn't talk to you, if only in a fairly stilted fashion, about a few choice topics--football included. It might even be able to come up with many different conversations about a given topic on successive days, as a human companion could. That would probably require it to be locked into the Internet, with the ability to find and interpret the latest news about (for instance) football matches, hirings and firings, and--ideally--the associated celebrity gossip: a companion unable to remark on David Beckham's looks and lifestyle, as well as his career on the pitch, would be disappointing for many people. But even if that's not feasible, an AI system might be able to converse about his professional activities, and follow the discussions on the sports pages of the newspapers. And we've seen that 'human' voices, and humanoid 'faces', are already on the horizon. So Frude's prediction stands: robot companions, disappointing though they might be in comparison with human acquaintances, are possible--and if not inevitable, perhaps probable. According to Frude, these less-than-human artefacts would be persuasive enough to excite anthropomorphic responses on your/our part. Was he right?

III: Animism in action

If Frude was wrong about the degree of technological difficulty, he was right about the animism--wherein something non-human (be it a god, a dog, a teddy-bear, or a robot) is 'naturally' regarded, to some degree, as though it were human. Animism is, as he said, a deep impulse in human beings. And its strong version, anthropomorphism, would indeed be excited by humanoid robots of the type sketched above.

In discussing that prediction today, we don't have to rely merely on our general knowledge and common sense, as Frude did. The last twenty-five years of research in developmental psychology strongly suggest that anthropomorphism, like weaker forms of animism, is rooted in a near-universal computational mechanism in human minds. ("Near-universal", because a minority of individuals, namely people with autism, appear not to possess it: Frith 1989/2003; Baron-Cohen 1995.) Specifically, it is rooted in what psychologists call Theory of Mind (ToM).

To possess ToM is to be able to see other people as agents, each with their own view of the world. Infants can't do this. (That's why they don't lie: they can't imagine that someone else could be deceived into believing something which they know to be false.) Gradually, however, they develop the ability to ascribe diverse beliefs (including false beliefs) and intentions to other people, and to use these in predicting and interpreting others' behaviour. (For the first experiments on this, see Wimmer and Perner 1983; since then, ToM has become something of an experimental industry.)

The computational mechanisms underlying ToM are still unclear. For instance, there's an ongoing controversy about whether the child actually develops a theory enabling it to generate hypotheses--i.e. animistic interpretations--and conclusions (as the label "ToM" suggests), or whether it develops the ability to simulate (empathize with) the mental states of other people (Davies and Stone 1995a,b). There's also debate about whether this distinction holds up, if carefully considered (Heal 1994). At a more detailed level, it's unclear whether, as some researchers have claimed, ToM involves a combination of general information-processing mechanisms with a dedicated four-slot ToM representation (Leslie 2000). There's dispute also about whether--or, better, to what extent--other species, such as chimps, possess ToM (Tomasello et al. 2003a,b).

However, one doesn't have to know just how a psych o-logical phenomenon works in order to see (some of) its effects. And it's significant, here, that ToM has recently been posited as a root of animism in religion (Boyer 1994).

Anthropologists have long reported a host of examples of animism, in both weak and strong (i.e. anthropomorphic) forms. By definition, animism relies on familiar, folk-psychological, concepts--which are grounded in ToM. But anthropologists and theologians alike point out that religious concepts are typically anomalous in some way (a fact which the religious believers themselves admit, and even celebrate). For example, one aspect of ToM is the realization that people in different places will perceive--attain knowledge of--different things (Wimmer and Perner 1983). But God, in the monotheistic religions, sees/knows everything. And in animist religions, a tree, snake, or river may have (limited) beliefs and wishes--which in non-religious contexts they cannot. This conceptual anomalousness, Pascal Boyer (1994) has argued, is largely why religions are both fascinating and slippery/unfalsifiable.

What's relevant for our purposes is that Frude's commonsense hunch about the near-unavoidability of animism can now be backed up by scientific evidence from anthropology as well as psychology.

Further evidence has come from neuroscience. To develop ToM depends, in part, on first developing a tendency to pay attention to other human individuals (which autistics typically don't do). Psychologists have shown that babies very soon lock onto their carer's eyes, and soon come to follow their gaze. Since Frude wrote his book, neuroscientists have discovered that there appears to be a universal basis for the development not only of face-recognition, but also of spontaneous engagement with (attention to) the eyes of a face (Johnson and Morton 1991). (The two cerebral mechanisms responsible for face-recognition, which develop at different times, have been experimentally studied in chicks but appear to be present in human babies also.) This mutual locking of gaze, in turn, scaffolds a huge range of carer-infant interactions, including turntaking in language (Hendriks-Jansen 1996).

The varied scientific research on ToM suggests that--as Frude predicted--robots with humanlike 'faces' and 'voices' will attract our attention, and prompt animistic responses. And the more humanlike the robot's use of language, and/or its 'thinking' processes (e.g. the way it plays chess: see below), the more this will be true. Indeed, we already have empirical evidence, some anecdotal and some systematic, that this is so.

The strength of people's spontaneous, nigh-irrepressible, impulse to engage with, and even to empathize with, today's humanoid robots is illustrated--for instance--by people's reactions to Cog, Kismet, and Leonardo (Breazeal and Scassellati 2000, 2002). The MIT research team report that visitors to the laboratory typically show not only curiosity, but also something very like a genuinely human response to their robots.

A particularly telling example has been reported by Sherry Turkle--no naive observer, but someone with a good understanding of AI. As she recalls the incident: "Cog 'noticed' me soon after I entered its room. Its head turned to follow me and I was embarrassed to note that this made me happy. I found myself competing with another visitor for its attention. At one point, I felt sure that Cog's eyes had 'caught' my own. My visit left me shaken--not by anything that Cog was able to accomplish but by my own reaction to 'him'….Despite myself, and despite my continuing skepticism about this research project, I had behaved as though in the presence of another being" (Turkle 1995: 266). Her reaction is especially interesting given the fact that Cog, unlike Kismet (which has floppy pink ears, large blue eyes, and an amazing set of eyelashes), looks nothing like an animal or human being.

IV: Human implications

If humanoid robots will elicit ToM-based reactions in human beings, that's not to say that the individuals concerned will always (ever?) take those natural reactions at face value. Like Turkle, they may be embarrassed by them, regarding them as utterly inappropriate. But our interest is less in whether people actually believe that the robots are genuine minds than in how their reactions to them affect their thinking about humankind.

We must distinguish, here, between robots in particular and AI in general. AI, with its close cousins computational psychology and philosophy of mind, has already changed the way that many literate people think about humanity. For example, the reified Cartesion self has given way to the self as a constructed narrative, which reflexively guides one's motives, priorities, and actions to a significant, though imperfect, degree (Minsky 1985; Dennett 1991: ch. 13). This notion, along with that of distributed cognition, has even led the countercultural opponents of symbolic AI--postmodernists, feminists...--to welcome much recent AI work, including situated robotics (e.g. Kember 2003). In addition, computational accounts of freedom--more accurately, of the type of cognitive/motivational system that's capable of making the sorts of decisions that we call free choice--have demystified this concept in many people's eyes (Dennett 1984; Boden 2006: 7.i.g).

As for belief in the uniqueness of humanity, this has seen both advance and retreat. Advance, because AI research has shown how hugely difficult (perhaps impossible) it is to model the richness and subtlety of the average human mind; retreat, because the sorts of things which humans, alone among animals, can do can now be done at least to some extent by artefacts. People are even (rashly, in my view--Boden 2006: 14.x-xi) attributing consciousness to some AI systems. More generally, the belief that human psychology, consciousness included, is somehow generated by a material system (the brain), as opposed to an immaterial soul, has been strengthened by AI as well as by neuroscience.

An advanced robotics will reinforce these already-familiar conceptual changes. But can anything more be said? Will it add further changes that haven't yet happened? If so, will those changes appear to bring us closer to robots, or will they deepen the perceived division between humanity and AI? The answer lies in the issue of what philosophers and psychologists call embodiment.

Given the psychological mechanisms of ToM and gaze mentioned above, it's easy to see why a 'look-alike' robot may be able to excite people's interest, and to elicit empathy or fellow-feeling in them. Two gaze-locking eyes are crucial, here. A one-eyed Cyclopean robot would be experienced not only as alien (as Gary Kasparov remarked about Deep Blue's chess-playing), but as positively repulsive. Even a superficial coating of fur, or a humanlike gait (with humanlike blunders), can help induce our sympathies. Consider Steve Grand's (2000) robot-gorilla Lucy, for example: this wouldn't have attracted so much attention from the media, or from highly articulate cultural critics (e.g. Fox-Keller forthcoming), without its fur coat. And the fact that the robot (unlike a VR character or avatar) is a physical thing, moving around in the material world--and encountering obstacles from time to time--much as we do ourselves, strengthens the illusion.

However, embodiment is more than mere physicality (materiality), even a physicality which outwardly resembles human bodies. To define it adequately, and to distinguish the many different senses in which the term is used, would take many pages (see Clark 1997; Wheeler 2005). Let's just say that it refers to a physical system--paradigmatically, an animal organism--that is situated in a material world with which it interacts continuously, and often more or less directly. (In my view, that "often" is important, despite situated roboticists' repeated efforts to deny the role of representations: Brooks 1991; Kirsh 1991.) Indeed, many disciples of embodiment speak less of "interaction" here than of "dynamical coupling".

The notion of interaction suggests the possibility of the system's deciding not to interact. Or, since oxygen-intake (for instance) must go on without interruption, the idea is that psychological interaction, or mental engagement, with the environment need not be continuous. Quite apart from sleep, the picture is that we may be lost in thought--with a rich mental panoply going on internally, irrespective of what's happening outside. (The ability of some people to override the horrors of the concentration camps, or of solitary confinement in the Gulag, could be cited here.) In theoretical psychology, talk of interaction may recall the Cartesian "sandwich": sensory perception, motor action, and purely mental processes in between. And much as a picnicker can eat just one slice of bread, or only the turkey-filling, so a psychologist can (on this view) admissibly theorize about only perception, or motor action, or thought--without bearing the other two parts of the sandwich constantly in mind.

By contrast, to speak of coupling is to posit one dynamical system (the organism) in close physical contact--and, crucially, a relation of mutual determination--with another (the environment). Indeed, the brain itself, or the visual cortex, or a single neurone ... can be thought of as a dynamical system (Maturana and Varela 1980). By the same token, plant and animal species can be thought of as systems within (closely coupled with) the earth as a whole (Lovelock 1988). Life, on this view, is quintessentially dynamic and environmentally grounded--and so is mind.

This approach leads to various questions that relate to the interpretation of any future robotics. For instance, just what is to count as motor action? I suggested, above, that a gaze-locking robot, or one with humanoid limbs and gait, would engage our attention. And so it would, for a while. But what if we were convinced--as a stress on embodiment suggests--that motor action (or behaviour) is no mere superficial cue to the presence of mind, or intelligence, but an ineliminable criterion of it? And what if we were also convinced that it must involve a close and continuous dynamical coupling with the physical environment? In that case, it's not clear that even situated robots, with the ability to react from time to time to environmental events, would be counted as showing "lifelike" behaviour--irrespective of whether their gait superficially resembled that of real cockroaches (Beer 1995) or real humans. Their initial attractiveness would very likely soon fade, as people realized the difference between their mode of being and ours.

That's just one example of the fact that the differences between robots and humans will be given prominence insofar as people attempt--as they doubtless will--to maintain a distinction between the two. And a key concept in this differentiation is likely to be the concept of life.

It's often said that mind requires life. And certainly, all the minds we know of are in living things. (Though not all living things have minds: think of oak-trees.) Since robots are not alive (so this argument goes) they can't have genuine intelligence, desires, or emotions ... so they can't pose a fundamental challenge to our idea of humanity, or to the unique place of Homo sapiens in the animal kingdom.

All very well--but there are two problems here. First, and despite the hopes of A-Life researchers (Langton 1996; Bedau 1996), there is still no universally accepted definition of life. Second, it's not clear that mind necessarily requires life. This "necessity" is usually taken for granted, very rarely explicitly argued (Boden 2006: 16.x). So for those with philosophical leanings, the brick wall between robots and humankind may appear rather more flimsy than it does to most people.

However, if we're concerned with the likely effect of interactive robotics on people in general, those relatively arcane points can be ignored. Whatever life is, robots don't have it.

(I've argued that this is because robots, and A-Life virtual systems too, don't metabolize in the required sense: Boden 1999. Mere energy-usage, or even individually budgeted 'packets' of energy, isn't enough. The energy must be used in the self-construction and bodily maintenance of the system, as well as in its functioning--which necessarily implies a complex interlocking set of metabolic cycles. These need not, perhaps, be carbon-based. But the chemical passivity of robots, and their manufactured provenance, don't fit metabolism as so defined.)

V: Conclusion

Advanced robotics, and especially interactive robotic 'companions', will affect how we think about human beings in various ways. Some of these will reinforce the changes to which AI in general has already led (a few were listed at the outset of Section IV). One example, mentioned above, is an increased respect for the richness and subtlety of human minds, because (or so I've argued in Section II) robotics/AI won't be able to match the variety and power of thought that we see in our neighbours and ourselves.

What robotics adds to 'screen-based' AI is physicality. Certain types of physicality (e.g. two eyes, with varying gaze) will unavoidably prompt ToM-based animistic responses. However, if the issue of embodiment is taken seriously then only robots of a very special kind (dynamically coupled with environment) would be philosophically plausible as bearers of psychological predicates. Moreover, even those would remain implausible for people who regard life as necessary for mind. For robots aren't living things.

It's not obvious that mind necessarily requires life. (As remarked above, this common assumption is very rarely explicitly defended.) But if it does, our 'natural' animistic responses to robots are ultimately as inappropriate as are our empathetic responses to teddy bears.

An emphasis on life as the basis for mind could lead us to take our biological nature more seriously. At present, in industrial cultures, the fact that we're biological organisms is largely forgotten--until something goes wrong, at which point we head for the hospital. Admittedly, our biological aspect has been highlighted by recent worries about food-additives, environmental pollution, and global warming. But it's commonly assumed that technology will trump ecology. If advances in robotics lead us (by the sort of 'distancing' argument sketched above) to have greater respect for our human nature as living organisms, perhaps that's all to the good.

Acknowledgments

This paper was supported by British Academy grant no. OCG-43209.

References

Baron-Cohen, S. (1995), Mindblindness: An Essay on Autism and Theory of Mind (Cambridge, Mass.: MIT Press).

Bedau, M. A. (1996), 'The Nature of Life', in M. A. Boden (ed.), The Philosophy of Artificial Life (Oxford: Oxford University Press), 332-357.

Beer, R. D. (1995a), 'A Dynamical Systems Perspective on Agent-Environment Interaction', Artificial Intelligence, 72: 173-215.

Boden, M. A. (1999), 'Is Metabolism Necessary?', British Journal for the Philosophy of Science, 50: 231-248.

Boden, M. A. (2006), Mind as Machine: A History of Cognitive Science (Oxford: Clarendon Press).

Boyer, P. (1994), The Naturalness of Religious Ideas: A Cognitive Theory of Religion (London: University of California Press).

Breazeal, C., and Scassellati, B. (2000), 'Infant-Like Social Interactions Between a Robot and a Human Caretaker', Adaptive Behavior, 8: 49-74. Special issue on Simulation Models of Social Agents, ed. C. Dautenhahn.

Breazeal, C., and Scassellati, B. (2002), 'Challenges in Building Robots that Imitate People', in C. Nehaniv and K. Dautenhahn (eds.), Imitation in Animals and Artefacts (Cambridge, Mass.: MIT Press), 363-390.

Brooks, R. A. (1991), 'Intelligence Without Representation', Artificial Intelligence, 47: 139-159.

Clark, A. J. (1997), Being There: Putting Brain, Body, and World Together Again (Cambridge, Mass.: MIT Press).

Davies, M., and Stone, T. (eds.), (1995a), Folk Psychology: The Theory of Mind Debate (Oxford: Blackwell).

Davies, M., and Stone, T. (eds.), (1995b), Mental Simulation: Evalu a-tions and Applications (Oxford: Blackwell).

Dennett, D. C. (1984), Elbow Room: The Varieties of Free Will Worth Wanting (Cambridge, Mass.: MIT Press).

Dennett, D. C. (1991), Consciousness Explained (London: Allen Lane).

Ekman, P. (2001), Telling Lies: Clues to Deceit in the Marketplace, Politics and Marriage. New York: Norton. (2nd edn., expanded.)

Fitt, S., and Isard, S. D. (1999), 'Synthesis of Regional English Using a Keyword Lexicon', Proceedings of Eurospeech 99 (Budapest), 823-826.

Fox Keller, E. (forthcoming), 'Booting-Up Baby', in J. Riskin (ed.), The Sistine Gap: Essays in the History and Philosophy of Artificial Life (Stanford: Stanford University Press).

Frith, U. (1989/2003), Autism: Explaining the Enigma (Oxford: Basil Blackwell). 2nd edn., revd., 2003.

Frude, N. (1983), The Intimate Machine: Close Encounters with the New Computers (London: Century).

Grand, S. (2003), Growing Up with Lucy: How to Build an Android in Twenty Easy Steps (London: Weidenfeld & Nicolson).

Grosz, B. (1977), 'The Representation and Use of Focus in a System for Understanding Dialogs', Proceedings of the Fifth International Joint Conference on Artificial Intelligence (Cambridge, Mass.), 67-76.

Grosz, B. J., and Sidner, C. L. (1979), 'Attention, Intentions, and the Structure of Discourse', Computational Linguistics, 12: 175-204.

Heal, J. (1994), 'Simulation versus Theory Theory: What is at Issue?', in C. Peacocke (ed.), Objectivity, Simulation and the Unity of Consciousness: Current Issues in the Philosophy of Mind (Oxford: Oxford University Press), 129-144.

Hendriks-Jansen, H. (1996), Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought (Cambridge, Mass.: MIT Press).

Johnson, M. H., and Morton, J. (1991), Biology and Cognitive Development: The Case of Face Recognition (Oxford: Blackwell).

Kember, S. (2003), Cyberfeminism and Artificial Life (London: Routledge).

Kirsh, D. (1991), 'Today the Earwig, Tomorrow Man?', Artificial Intelligence, 47: 161-84.

Langton, C. G. (1996), 'Artificial Life', in M. A. Boden (ed.), The Philosophy of Artificial Life (Oxford: Oxford University Press), 39-94. This is a revized version of the paper in C. G. Langton (ed.), Artificial Life (Redwood City, CA: Addison-Wesley, 1989), 1-47.

Leslie, A. M. (2000), '"Theory of Mind" as a Mechanism of Selective Attention', in M. S. Gazzaniga (ed.), The New Cognitive Neurosciences, 2nd edn., revd. (Cambridge, Mass.: MIT Press), 1235-1247.

Lovelock, J. (1988), The Ages of Gaia: A Biography of Our Living Earth (Oxford: Oxford University Press).

McKeown, K. (1985), 'Discourse Strategies for Generating Natural-Language Text', Artificial Intelligence, 27: 1-42.

Maturana, H. R., and Varela, F. J. (1980), Autopoiesis and Cognition: The Realization of the Living (Boston: Reidel). First published in Spanish, 1972.

Minsky, M. L. (1985), The Society of Mind (New York: Simon & Schuster).

Tomasello, M., Call, J., and Hare, B. (2003a), 'Chimpanzees Understand Psychological States: The Question is Which Ones and to What Extent', Trends in Cognitive Science, 7: 153-156.

Tomasello, M., Call, J., and Hare, B. (2003b), 'Chimpanzees Versus Humans: It's Not That Simple', Trends in Cognitive Sciences, 7: 239-240.

Turkle, S. (1995), Life on the Screen: Identity in the Age of the Internet (New York: Simon & Schuster).

Wheeler, M. W. (2005), Reconstructing the Cognitive World: The Next Step (Cambridge, Mass.: MIT Press).

Wimmer, H., and Perner, H. (1983), 'Beliefs About Beliefs: Representations and Constraining Function of Wrong Beliefs in Young Children's Understanding of Deception', Cognition, 13: 103-128.

Wright, I. P., Sloman, A., and Beaudoin, L. P. (1996), 'Towards a Design-Based Analysis of Emotional Episodes', Philosophy, Psychiatry, and Psychology, 3: 101-137.