

9

Future Minds: Transhumanism, Cognitive Enhancement, and the Nature of Persons

SUSAN SCHNEIDER

TRANSHUMANISM IS A philosophical, cultural, and political movement that holds that the human species is only now in a comparatively early phase and that its very evolution will be altered by developing technologies.¹ Future humans will, in effect, be very unlike their current incarnation in both physical and mental capacities and will be more like certain persons depicted in science fiction novels. Transhumanists share the belief that an outcome in which humans have radically advanced intelligence, near immortality,

Thanks to Vardit Ravitsky and James Hughes for their helpful comments on this chapter.

deep friendships with AI (artificial intelligence) creatures, and elective body characteristics is a very desirable end for both one's own personal development and for the development of our species as a whole. Despite its science fiction-like flavor, the issues that transhumanism presents deserve to be taken seriously because the beginning stages of this radical alteration are supposed to be the outcome of technological developments that are either here, if not generally available, or more commonly technologies that are accepted by many in the relevant scientific fields as being on their way (Roco & Bainbridge, 2002). In the face of all these technological developments, transhumanists present a thought-provoking and highly controversial progressive bioethics agenda. Transhumanism offers intriguing perspectives on (inter alia) one's conception of the good life, the nature of persons, and the nature of mind.

This chapter will cover the basic tenets of transhumanism and will then discuss what I take to be the most important philosophical element of the transhumanist picture—its unique perspective on the nature and development of persons. Persons are traditionally regarded as being an important moral category, being the bearers of rights (if you believe in such) or at least deserving of consideration in the utilitarian calculus. And, as we shall see, considering the nature of persons through the lens of transhumanism involves pushing up against the boundaries of the very notion of personhood. Consider, for example, the enhancement debate. When one considers whether to enhance in the radical ways the transhumanists advocate, one must ask, "Will this radically enhanced creature still be me?" If not, then, on the reasonable assumption that one key factor in a decision to enhance oneself is one's own personal development, even the most progressive bioethicist will likely regard the enhancement in question as undesirable. For when you choose to enhance in these radical ways, the enhancement does not really enhance *you*. Examining the enhancement issue through the vantage point of the metaphysical problem of personal identity shall thereby present a serious challenge to transhumanism. Indeed, this is a pressing issue for any argument made for or against enhancement.

THE MAIN TENETS OF TRANSHUMANISM

TRANSHUMANISM IS BY no means a monolithic ideology, but it does have an organization and an official declaration. The World Transhumanist Association is an international nonprofit organization that was founded in 1998 by philosophers Nick Bostrom and David Pearce. Its tenets were laid out in the *Transhumanist Declaration* (World Transhumanist Association, 1998) and are reprinted here.

The Transhumanist Declaration

1. Humanity will be radically changed by technology in the future. We foresee the feasibility of redesigning the human condition, including such parameters as the inevitability of aging, limitations on human and artificial intellects, unchosen psychology, suffering, and our confinement to the planet earth.
2. Systematic research should be put into understanding these coming developments and their long-term consequences.
3. Transhumanists think that by being generally open and embracing of new technology we have a better chance of turning it to our advantage than if we try to ban or prohibit it.
4. Transhumanists advocate the moral right for those who so wish to use technology to extend their mental and physical (including reproductive) capacities and to improve their control over their own lives. We seek personal growth beyond our current biological limitations.
5. In planning for the future, it is mandatory to take into account the prospect of dramatic progress in technological capabilities. It would be tragic if the potential benefits failed to materialize because of technophobia and unnecessary prohibitions. On the other hand, it would also be tragic if intelligent life went extinct because of some disaster or war involving advanced technologies.
6. We need to create forums where people can rationally debate what needs to be done, and a social order where responsible decisions can be implemented.
7. Transhumanism advocates the well-being of all sentience (whether in artificial intellects, humans, posthumans, or non-human animals) and encompasses many principles of modern humanism. Transhumanism does not support any particular party, politician or political platform.

This document was followed by the much longer and extremely informative *Transhumanist Frequently Asked Questions*, authored by Nick Bostrom, in consultation with dozens of leading transhumanists (Bostrom, 2003b).² Because the current chapter is brief and cannot touch on all elements of transhumanism, the reader is strongly encouraged to read this document for a more complete overview of transhumanism.³

THE NATURE OF PERSONS

I WILL NOW offer a philosophical analysis of some of the ideas expressed in this declaration. Overall, central transhumanist texts have advanced a sort of trajectory for the personal development of a contemporary human, technology permitting (Kurzweil, 1999, 2005; Bostrom, 2003b, 2005).

21st century unenhanced human → significant “upgrading” with cognitive and other physical enhancements → posthuman status → “super-intelligence”

By way of illustration, suppose it is now 2025 and being a technophile, you purchase cognitive and physical enhancements as they become readily available. First, you add a mobile internet connection to your retina, then you enhance your working memory by adding neural circuitry. You are now officially a cyborg. Now skip ahead to 2040. Through nanotechnological therapies/enhancements you are able to extend your lifespan, and as the years progress, you continue to accumulate more far-reaching enhancements. By 2060, after several small but cumulatively significant alterations, you are a posthuman. Posthumans are possible future beings, “whose basic capacities so radically exceed those of present humans as to be no longer unambiguously human by our current standards” (Bostrom, 2003b). Such posthumans can be AI devices, humans who have uploaded their brains onto computers and then enhanced them, or humans who are the result of making many small but cumulatively profound enhancements (Bostrom, 2003b). At this point, your intelligence is enhanced not just in terms of speed of mental processing; you are now able to make profound connections that you were not able to make before. Unenhanced humans, or naturals seem to you to be intellectually disabled—you have little in common with them—but as a transhumanist, you are supportive of their right to not enhance (Bostrom, 2003b; Garreau, 2005; Kurzweil, 2005).

It is now 2600. For years, worldwide technological developments, including your own enhancements, have been facilitated by super-intelligent AI. Indeed, “creating superintelligence may be the last invention that humans will ever need to make, since superintelligences could themselves take care of further scientific and technological development” (Bostrom, 2003b). And the slow addition of better and better neural circuitry has now resulted in there being no real intellectual difference in kind between you and superintelligent AI—you too are a superintelligence, a creature with the capacity to radically outperform the best human brains in practically every field, including scientific creativity, general wisdom, and social skills” (Bostrom, 2003b).⁴ The only real difference between you and an AI creature of standard design is one of origin—you were once a natural. But you are now almost entirely engineered by technology—you are perhaps more aptly characterized as a member of a rather heterogeneous class of AI life forms (Kurzweil, 2005).

This, then, is a very rough sketch of the developmental trajectory that the transhumanist generally aspires to.⁵ Now, let us ask: should you

embark upon this journey?⁶ Here, there are deep philosophical questions that have no easy answers. For in order to understand whether *you* should enhance, you must first understand what you are to begin with. But what is a person? And given your conception of a person, after such radical changes, would you still be you or would you actually bear little relation to the person you were before? And if the latter situation is the case, why would embarking on the path to radical enhancement be something you value? For wouldn't it instead be a path that leads to your own demise, leading you away from your true self, ultimately causing you to cease to exist? In order to make such a decision one needs to understand the metaphysics of personal identity—that is, one needs to answer the question: What is it in virtue of which a self or person is supposed to continue existing over time? A good place to begin is to consider that everyday objects seem to persist over time. Consider the espresso machine in your favorite café. Suppose that five minutes have elapsed and suppose the barista has turned the machine off. Imagine asking the barista if the machine is still the same machine, despite this change. The ordinary answer is that it is of course possible for the machine to continue to be one and the same thing over time. This seems to be a reasonable case of its persistence, even though at least one of the machine's features or properties has changed. On the other hand, if the machine disintegrated or melted, then it would no longer be the same machine. It wouldn't be an espresso machine at all for that matter. So some changes do not cause a thing to cease to exist while others do. Philosophers call features or properties that are essential to a thing or person's nature "essential properties."

Now reconsider the transhumanist's trajectory for enhancement: for radical enhancement to be a worthwhile option for you, it has to represent a desirable form of personal development; at bare minimum, even if enhancement brings such goodies as superhuman intelligence and radical life extension, it must not involve the elimination of one of your essential properties. *For in this case, the sharper mind and fitter body would not be experienced by you—they would be experienced by someone else.* For even if you would like to become superintelligent, knowingly embarking on a path that trades away one or more of your essential properties would be tantamount to suicide—that is, to your intentionally causing yourself to cease to exist. So before you enhance, you had better get a handle on what your essential properties are.

Key transhumanists have grappled with this issue. For instance, Ray Kurzweil asks: "So who am I? Since I am constantly changing, am I just a pattern? What if someone copies that pattern? Am I the original and/or the copy? Perhaps I am this stuff here—that is, the both ordered and chaotic collection of molecules that make up my body and brain" (Kurzweil, 2005, p. 383).

Kurzweil is here referring to two theories at center stage in the age old philosophical debate about what properties determine the nature of persons. The leading theories are the following:

1. The ego theory—a person’s nature is her soul or nonphysical mind, and this mind or soul can survive the death of the body.⁷
2. The psychological continuity theory—you are essentially your memories and ability to reflect on yourself (Locke) and, more generally, your overall psychological configuration, what Kurzweil referred to as your “pattern.”⁸
3. Materialism—you are essentially the material that you are made out of—what Kurzweil referred to as “the ordered and chaotic collection of molecules that make up my body and brain” (Kurzweil, 2005, p. 383).
4. The no self view—there is no metaphysical category of person. The “I” is a grammatical fiction (Nietzsche). There are bundles of impressions but no underlying self (Hume). There is no survival because there is no person (Buddha, Parfit).

Each of these views has its own position about whether to enhance. If you hold (1) then your decision to enhance depends on whether you believe the enhanced body would retain the same soul or immaterial mind.⁹ If you believe (3), then any enhancements must not alter your material substrate. In contrast, according to (2), or patternism, enhancements can alter the material substrate but must preserve your memories and your overall psychological configuration. Finally, (4) contrasts sharply with (1)–(3). If you hold (4), then the survival of the person is not an issue, for there is no person to begin with. But you may strive to enhance nonetheless, to the extent that you may find intrinsic value in adding more superintelligence to the universe—you might value life forms with higher forms of consciousness and wish that your closest “continuent” should be such a creature.

Let us focus on identifying which of these conceptions conforms to the transhumanist notion of the self, at least in its most characteristic incarnation. Consider that transhumanists generally adopt a computational theory of mind. That is, the mind is essentially the program running on the hardware of the brain, where by program what is meant is the algorithm that the mind computes, something in principle discoverable by cognitive science.¹⁰ Because, at least in principle, the brain’s computational configuration can be preserved in a different medium (i.e., in silicon as opposed to carbon), with the information processing properties of the original neural circuitry preserved, the computationalist rejects the materialist view of the nature of persons.¹¹ Indeed, as

Kurzweil explains, materialism seems to falter in embracing the very idea that you are what you are made up of:

The specific set of particles that my body and brain comprise are in fact completely different from the atoms and molecules that I comprised only a short while ago. We know that most of our cells are turned over in a matter of weeks, and even our neurons, which persist as distinct cells for a relatively long time, nonetheless change all of their constituent molecules within a month. . . . I am rather like the pattern that water makes in a stream as it rushes past the rocks in its path. The actual molecules of water change every millisecond, but the pattern persists for hours or even years. (Kurzweil, 2005, p. 383)

Kurzweil calls his view “Patternism” (2005, p. 386). Patternism is an updated version of the psychological continuity theory. Put in the language of cognitive science, as the transhumanist surely would, what is essential to you is your computational configuration—for example, what sensory systems/subsystems your brain has (e.g., early vision), the way that the basic sensory subsystems are integrated in association areas, the neural circuitry making up your domain general reasoning, your attentional system, your memories, and so on—overall, the algorithm that the brain computes. I believe that Kurzweil’s appeal to patternism is highly typical of the transhumanist. Indeed, consider the appeal to patternism in the following passage of the *Transhumanist Frequently Asked Questions*, which discusses the process of uploading, a process which shall be important to our subsequent discussion.

Uploading (sometimes called “downloading”, “mind uploading” or “brain reconstruction”) is the process of transferring an intellect from a biological brain to a computer. One way of doing this might be by first scanning the synaptic structure of a particular brain and then implementing the same computations in an electronic medium. . . . An upload could have a virtual (simulated) body giving the same sensations and the same possibilities for interaction as a non-simulated body. . . . And uploads wouldn’t have to be confined to virtual reality: they could interact with people on the outside and even rent robot bodies in order to work in or explore physical reality. . . . Advantages of being an upload would include: Uploads would not be subject to biological senescence. Backup copies of uploads could be created regularly so that you could be re-booted if something bad happened. (Thus your lifespan would potentially be as long as the universe’s.) . . . Radical cognitive enhancements would likely be easier to implement in an upload than in an organic brain. . . . *A widely accepted position is that you survive so long as certain information patterns are conserved.* . . . For the continuation of personhood, on this view, it matters little whether you are implemented on

a silicon chip inside a computer or in that gray, cheesy lump inside your skull, assuming both implementations are conscious. (Bostrom, 2003b, emphasis mine)

This is a clear appeal to patternism. And as we shall see, both patternism and the process of uploading introduce philosophical puzzles for the transhumanist case for enhancement. Indeed, they even raise problems with the transhumanist's justification for mild enhancements. As I shall now explain, such problems desperately need to be addressed.

PUZZLES

NOW THAT WE'VE identified the theory of personal identity that the transhumanist generally adopts, let us ask: At the point at which you enhance, being part natural and part artificial, assuming a patternist conception of the nature of persons, are you the same person you were before? Or is there some point in which you cease to exist, becoming a different person entirely? Consider first a mild enhancement—the deletion of a few memories, say, to remove bad chess playing habits and facilitate better chess strategies. Surprisingly, it is not even clear that this enhancement would be compatible with survival, according to patternism. Way back in 1785, Thomas Reid raised the following, now classic, problem for Patternism:

Suppose a brave officer to have been flogged when a boy at school, for robbing an orchard, to have taken a standard (a flag) from the enemy in his first campaign, and to have been made a general in advanced life: suppose also, which must be admitted to be possible, that, when he took the standard, he was conscious of his having been flogged at school, and that when made a general he was conscious of his taking the standard, but had absolutely lost consciousness of his flogging. (Reid, 1785/1941, p. 213)

Reid's example presents a serious challenge to the patternist theory. Identity is transitive: if $A = B$ and $B = C$ then $A = C$. Patternism holds that the boy is identical to the officer (as the officer has the boy's memory of the flogging) and the officer is identical to the general (as the general was conscious of taking the flag). But notice that patternism cannot say that the boy is identical to the general, as the general does not recall being flogged. Patternism violates the transitivity of identity. This is an abysmal result: Patternism, as it stands, is not really a theory of personal *identity*.

But perhaps the patternist could somehow modify her theory to allow that a gradual change in one's pattern preserves personhood.

Here the issues grow too complex for a brief chapter, but perhaps, for instance, an understanding of the neurodynamics underlying ordinary cognitive changes could give the transhumanist a route into this problem. An appeal to dynamical systems theory would certainly be in keeping with the cognitive science orientation of transhumanism. On the assumption that people normally survive from moment to moment, we can then propose that certain therapies/enhancements should be safe by patternist standards: enhancements/therapies that modify the brain's dynamical or computational structures in a way that mimics the natural process of change in the brain. Such therapies/enhancements would preserve one's pattern because they would not be a significant departure from the brain's characteristic dynamical patterns.¹²

But notice that the new patternist theory will face the following challenge. In order for the transhumanist to justify the sort of enhancements needed to become a cyborg, a posthuman, or a superintelligent being, she will need to say that radical or unusual changes in existing structures are compatible with the survival of the person. But does patternism really allow for these enhancements? For instance, what about adding an intelligence-enhancing working memory chip so that one can perform better in law school? Would this be too sharp of a break in the existing pattern? Or what about adding a new sense (e.g., echolocation)? It appears that merely appealing to patternism is not enough to justify opting for the neural enhancements that the transhumanist envisions. Transhumanism desperately needs to develop an informative account of personhood. That is, for any theory of personal identity it defends, it needs to say which enhancements are merely changes in nonessential properties and which would be changes in essential ones. In the context of patternism, the extreme cases are clear—a memory erasure process that erased one's childhood is clearly the loss of essential property for the continuity theory because it removes much of one's memories. Mere everyday cellular maintenance by nanobots to overcome the slow effects of aging would, on this view, not affect the identity of the person. The middle range cases are unclear. Maybe deleting a few bad chess playing habits is kosher, but what about erasing a bad relationship, as in the film *Eternal Sunshine of the Spotless Mind*? The path to Superintelligence may very well be a path through middle range enhancements. Without a firm handle on the personal identity question, the transhumanist developmental trajectory is perhaps the technophile's alluring path to suicide.

But let us press on; let us suppose that the transhumanist can offer a principled means of distinguishing suicide-inducing enhancements (so to speak) from ones compatible with survival. Nonetheless, further problems arise.

DEREK PARFIT'S TELEPORTATION CASE

IT'S 2080 AND you are an astronaut. You attend a briefing on your next mission. You've been selected for a secret mission to a far away planet via a new means of travel. Fortunately, your trip will be quick, indeed, much of it will be at the speed of light. NASA superscientists will take a complete scan of your brain—capturing every detail of its computational configuration. Your pattern—that is, you—will be uploaded and sent to the planet, and there your brain will be reconstructed from matter that is configured precisely according to the information from the scan. In the process of scanning, your earthly brain will be destroyed, but that doesn't matter to you. For like Ray Kurzweil, you reject materialism; what is important to you is that your pattern will be safely housed in a supercomputer until, in short order, it will inhabit a new brain and body. You are being temporarily uploaded.¹³

Should you go? If you haven't studied personal identity you might be fooled into thinking you should. But we can quickly see that you wouldn't survive. There may be a person created on the planet, but it is merely your clone. We don't need to appeal to a particular theory of personal identity to see this point—the idea that it would be you is incoherent. Consider that if the above scenario is possible, then it is also metaphysically possible (i.e., conceivable) that you were not destroyed in the process. But now, in that case, who would be on the planet? It couldn't be you. You are on Earth. And because this person is clearly not you if you weren't destroyed, it follows that it wasn't you if you were. For the life or death of another creature isn't an essential property of a person. Hence, uploading doesn't preserve personhood.

But now transhumanism is in big trouble: your duplicate on the planet has your pattern, precisely. So it must be that even an improved version of (2) is false: sameness of pattern is not sufficient for sameness of person. As a result, transhumanism cannot claim that enhancement is desirable, for its very means of deciding whether it is—its theory of personal identity—is seriously flawed. Further, we can use this result to prove that even mild enhancements are death inducing. Assume your copy on the other planet is not you, as should now be obvious. Now, consider an earlier point in your life at which time, being a good transhumanist, you had a gradual neural regeneration procedure. That is, at each doctor visit, you had 1% of your neurons replaced by silicon-based artificial neurons having the very same computational or formal properties that normally underlie your thoughts.¹⁴ At 100% aren't you analogous to the creature on the planet, being composed of entirely new matter? Given our earlier discussion about this creature, we have reason to believe that this final product is not you. For the creature in the teleportation case was clearly not you but your replica. But, at the

other end, it seems that if 1% of cells are replaced, it clearly would be you. (After all, as Kurzweil pointed out, this cell replacement process happens to us all the time.) Now, in the cases in between, the person must either be you or a replica. But crucially, which percentage is the critical percentage in which the resulting person would be you, and beyond which, the person is merely a replica? But how could there be one? A few cells couldn't make such a significant difference, could they? Since it is absurd to locate a critical percentage, it must be that there is something deeply wrong with the patternist view of the self.¹⁵

CONCLUSION

I HOPE ALL this has convinced you that if the transhumanist maintains patternism there are some serious issues that require working out. Indeed, as the *Transhumanist Frequently Asked Questions* indicates, the development of radical enhancements such as brain-machine interfaces, cryogenic freezing for life extension, and uploading to avoid death or simply to facilitate enhancement are key enhancements invoked by the transhumanist view of the development of the person. And, quite ironically, all of these enhancements sound strangely like the thought experiments philosophers have been appealing to for years as problem cases for various theories of the nature of persons. Given that it seems unclear whether sameness of personhood would even be preserved by any of these enhancements, it is fair to say that without further work on this topic, the transhumanist cannot support her case for enhancement. Indeed, the *Transhumanist Frequently Asked Questions* notes that transhumanists are keenly aware that this issue has been neglected:

While the concept of a soul is not used much in a naturalistic philosophy such as transhumanism, many transhumanists do take an interest in the related problems concerning personal identity (Parfit 1984) and consciousness (Churchland 1988). These problems are being intensely studied by contemporary analytic philosophers, and although some progress has been made, e.g. in Derek Parfit's work on personal identity, they have still not been resolved to general satisfaction. (Bostrom, 2003b, section 5.4)

Our discussion also raises some general lessons for all parties involved in the enhancement debate. For when one considers the enhancement debate through the lens of the metaphysics of personhood, new dimensions of the debate are appreciated. The literature on the nature of persons is a literature that is extraordinarily rich, raising

intriguing problems for commonly, and often implicitly, accepted views of the nature of persons that underlie positions on enhancement. And it seems fair to say that when a theory defends or rejects a given enhancement, it is important to determine whether its position on the enhancement in question is truly supported by, or even compatible with, the perspective of the nature of persons that the theory is sympathetic to. Further, the topic of the nature of persons is of clear relevance to the related topics of human nature and human dignity, issues that are currently key points of controversy in debates over enhancement (see, e.g., Bostrom, in press, "Dignity and Enhancement"; Fukuyama, 2002).

Perhaps, alternately, you grow weary of all this metaphysics. You may suspect that social conventions concerning what we commonly consider to be persons are all we have because metaphysical theorizing will never conclusively resolve what persons are. However, as unwieldy as metaphysical issues are, it seems that not all conventions are worthy of acceptance, so one needs a manner of determining which conventions should play an important role in the enhancement debate and which ones should not. And it is hard to accomplish this without getting clear on one's conception of persons. Further, it is difficult to avoid at least implicitly relying on a conception of persons in the context of reflecting on the case for and against enhancement. For what is it that ultimately grounds your decision to enhance or not enhance if not that it will somehow improve who you are? Are you perhaps merely planning for the well-being of your closest continuent?

NOTES

1. Julian Huxley apparently coined the term *transhumanism* in 1957, when he wrote that in the near future "the human species will be on the threshold of a new kind of existence, as different from ours as ours is from that of Peking man" (Huxley, 1957, pp. 13–17).

2. Bostrom is a philosopher at Oxford University who now directs the transhumanist-oriented Future of Humanity Institute there.

3. This document was updated in 2003 and is available at the World Transhumanist Association Web site. In addition, there are a number of excellent philosophical and sociological works that articulate key elements of the transhumanist perspective (e.g., Bostrom, 2004; Hughes, 2004, n.d.; Kurzweil, 1999, 2005). For extensive Web resources on transhumanism see Nick Bostrom's homepage, Ray Kurzweil's newsgroup (KurzweilAI.net), the Institute for Ethics and Emerging Technologies homepage, and the World Transhumanist Association homepage.

AQ1

4. There are of course numerous nuances to this rough picture. For instance, some transhumanists believe that the move from unenhanced human intelligence to superintelligence will be extremely rapid because we are approaching a singularity, a point at which the creation of superhuman intelligence will result in massive changes in a very short period (e.g., 30 years). (Bostrom, 1998; Kurzweil, 1999, 2005; Vinge, 1993). Other transhumanists hold that technological changes will not be so sudden. These discussions often debate the reliability of Moore's law (Moore, 1965). Another key issue is whether a transition to superintelligence will really occur because the upcoming technological developments involve grave risk. The risks of biotechnology and AI concern transhumanists, progressive bioethicists more generally, as well as bioconservatives (Annis, 2000; Bostrom, 2002a, 2002b; Garreau, 2005; Joy, 2000).

5. It should be noted that transhumanism by no means endorses every sort of enhancement. For example, Nick Bostrom rejects positional enhancements (enhancements primarily employed to increase one's social position) yet argues for enhancements that could allow humans to develop ways of exploring "the larger space of possible modes of being" (2005, p. 11).

6. For mainstream anti-enhancement positions on this question see, for example, Fukuyama (2002), Kass et al. (2003), and Annas (2000).

7. For nice surveys of these four positions see Blackburn (1999) and Olson (2008).

8. See chapter 27 of John Locke's 1694 *Essay Concerning Human Understanding* (note that this chapter first appears in the second edition; is also reprinted as "Of Identity and Diversity," in Perry, 1975). For other attempts to develop similar views see Quinton (1962) and Grice (Oct. 1941), both of which are also reprinted in Perry (1975).

9. It should be noted that although a number of bioconservatives seem to uphold the soul theory, the soul theory is not, in and of itself, an anti-enhancement position. For why can't one's soul or immaterial mind inhere in the same body even after radical enhancement?

10. Computational theories of mind can appeal to various computational theories of the format of thought (e.g., connectionism, dynamic systems theory, symbolism, or some combination thereof). See Kurzweil (2005). For philosophical background see Block (1995) and (Churchland, 1996).

11. This commonly held but controversial view in philosophy of cognitive science is called multiple realizability (Kim, 2006); Bostrom (2003a) calls it "substrate independence."

12. For a nice introduction to issues in dynamical systems theory see Scott Kelso (1997) and Walter Freeman (2001). For a more extensive discussion of the different versions of the memory theory and ways of answering Reid's objection within the metaphysics literature, as well as other objections to the theory, see the various papers in Perry (1975), especially Perry's introduction, which provides a nice overview.

13. This example is modified from a classic paper by Derek Parfit (1987).

AQ2

14. This case is again inspired by Parfit (1987). Kurzweil considers similar thought experiments in his intriguing discussion of personal identity. Unfortunately, while he notes the problems with patternism, he doesn't offer a resolution (Kurzweil, 2005, pp. 382–387). The *Transhumanist FAQ* actually considers a similar case: "An alternative hypothetical uploading method would proceed more gradually: one neuron could be replaced by an implant or by a simulation in a computer outside of the body. Then another neuron, and so on, until eventually the whole cortex has been replaced and the person's thinking is implemented on entirely artificial hardware" (Bostrom, 2003b).

15. There are numerous ways that the transhumanist could respond to the preceding argument. For discussion of further patternist options see Perry (1975). Alternately, the transhumanist might instead accept a no self view, as sociologist James Hughes does in his (2004 and 2005) and in his forthcoming book, *Cyborg Buddha* (Hughes, n.d.). (Relatedly, see also the Institute for Ethics and Emerging Technology's "Cyborg Buddha" project at <http://ieet.org/index.php/IEET/cyborgbuddha>.)

REFERENCES

- Annas, G. J. (2000). The man on the moon, immortality, and other millennial myths: The prospects and perils of human genetic engineering. *Emory Law Journal*, 49(3), 753–782.
- Blackburn, S. (1999). The self. In *Think: A compelling introduction to philosophy* (pp. 129–140). Oxford: Oxford University Press.
- Block, N. (1995). The mind as the software of the brain. In D. Osherson, L. Gleitman, S. Kosslyn, E. Smith, & S. Sternberg (Eds.), *An invitation to cognitive science* (pp. 377–421). New York: MIT Press.
- AQ3 Bostrom, N. (1998). How long before superintelligence? *International Journal of Futures Studies*, 2. Retrieved from <http://jetpress.org/contents.htm>
- AQ4 Bostrom, N. (2002a). Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology*, 9, XX–XX.
- Bostrom, N. (2002b). When machines outsmart humans. *Futures*, 35(7), 759–764.
- Bostrom, N. (2003a). Are you living in a computer simulation? *Philosophical Quarterly*, 53(211), 243–255.
- Bostrom, N. (2003b). *The Transhumanist Frequently Asked Questions: v 2.1*. World Transhumanist Association. Retrieved June 20, 2008, from <http://transhumanism.org/index.php/WTA/faq/>
- AQ5 Bostrom, N. (2005). History of transhumanist thought. *Journal of Evolution and Technology*, 14(1).
- Bostrom, N. (in press). Dignity and enhancement. Commissioned for the President's Council on Bioethics.
- Churchland, P. (1988). *Matter and consciousness*. Cambridge, MA: MIT Press.

- Churchland, P. (1996). *Engine of reason, seat of the soul*. Cambridge, MA: MIT Press.
- Clark, A. (2003). *Natural born cyborgs*. Oxford: Oxford University Press.
- Freeman, W. (2000). *How brains make up their minds*. New York: Columbia University Press.
- Fukuyama, F. (2002). *Our posthuman future: Consequences of the biotechnology revolution*. New York: Farrar, Straus and Giroux.
- Garreau, J. (2005). *Radical evolution: The promise and peril of enhancing our minds, our bodies—and what it means to be human*, New York: Doubleday.
- Grice, P. (1941). Personal identity. *Mind*, 50, XX–XX.
- Hughes, J. (2004a). *Citizen cyborg: Why democratic societies must respond to the redesigned human of the future*. Cambridge, MA: Westview Press.
- Hughes, J. (2004b) The death of death. In C. Machado & D. A. Shewmon (Eds.), *Brain death and disorders of consciousness* (pp. 79–88). New York: Kluwer.
- Hughes, J. (2005). The illusiveness of immortality. In C. Tandy (Ed.), *Death and anti-death, volume 3: Fifty years after Einstein, one hundred fifty years after Kierkegaard* (pp. XX–XX). New York: Ingram.
- Hughes, J. (n.d.). *Cyborg Buddha*. Manuscript in preparation. Retrieved from <http://ieet.org/index.php/IEET/cyborgbuddha>
- Huxley, J. (1957). *New bottles for new wine*. London: Chatto & Windus.
- Joy, B. (2000). Why the future doesn't need us. *Wired*, 8, 238–246.
- Kass, L., Blackburn, E., Dresser, R., Foster, D., Fukiyama, F., Gazzaniga, M., et al. (2003). *Beyond therapy: Biotechnology and the pursuit of happiness: A report of the President's Council on Bioethics*. Commissioned report by the President's Council, Washington, DC.
- Kelso, S. (1997). *Dynamical patterns*. New York: MIT Press.
- Kim, J. (2006). *Philosophy of mind* (2nd ed.). New York: Westview Press.
- Kurzweil, R. (1999). *The age of spiritual machines: When computers exceed human intelligence*. New York: Viking.
- Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. New York: Viking.
- Locke, J. (1694). *Essay concerning human understanding* (2nd ed.).
- Moore, G. (1965). Cramming more components into integrated circuits. *Electronics*, 38(8), 11–17. Retrieved August 20, 2008, from <ftp://download.intel.com/research/silicon/moorespaper.pdf>
- Olson, E. T. (2008). Personal identity. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy (Spring 2007 Edition)*. Retrieved August 20, 2008, from <http://plato.stanford.edu/archives/spr2007/entries/identity-personal/>
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Parfit, D. (1987). Divided minds and the nature of persons. In C. Blakemore & S. Greenfield (Eds.), *Mindwaves* (pp. 19–28). Oxford: Blackwell Publishers.
- Perry, J. (1975). *Personal identity*, Berkeley: University of California Press.
- Quinton, A. (1962). The soul. *The Journal of Philosophy*, 59(15), XX–XX.

AQ6

AQ7

AQ8

AQ9

- Reid, T. (1785/1941). *Essays on the intellectual powers of man*. A. D. Woozley (Ed.). London: Macmillian.
- Roco, M. C., & Bainbridge W. S. (Eds.). (2002). *Converging technologies for improved human performance: Nanotechnology, biotechnology, information technology and cognitive science*. Arlington, VA: National Science Foundation/Department of Commerce.
- Vinge, V. (1993, Winter). The coming technological singularity. *Whole Earth Review*.
- World Transhumanist Association. (1998). *Transhumanist declaration*. Retrieved August 20, 2008, from <http://www.transhumanism.org/index.php/WTA/declaration/>

AUTHOR QUERY

- AQ1: You added Bostrom (2004) to note 3 during MS review but there is no corresponding reference entry. Please add one or remove this text citation.
- AQ2: In the reference you updated for Freeman you provided (2000) as the date. Please review and make consistent both there and here in note 12.
- AQ3: For Bostrom (1998), please provide the retrieval date.
- AQ4: For Bostrom (2002a), please provide the range of pages for this article.
- AQ5: For Bostrom (2005), please provide the URL for this article and the retrieval date (in lieu of page numbers).
- AQ6: You had removed Grice (1941) during MS review but it is cited in note 8 so I have retained the reference. If you would like to remove this source, please also remove the citation in note 8. If the source and citation are to remain, please provide the range of pages for this article.
- AQ7: For Hughes (2005), please provide the range of pages for this chapter.
- AQ8: Please provide publication information for Locke (1694). Is 1694 the date of the 2nd edition? I find various initial publication dates for it.
- AQ9: You had marked during MS review for Quinton (1962) to be deleted. But it is cited in note 8 so I have retained it here. If you would like this removed, please do so both here and in note 8. If this should be retained, please provide the range of pages for the article.