Technology and the Future of Persons

Lynne Rudder Baker

University of Massachusetts Amherst

Futurologists believe that we are on the verge of merging with our technology. Techno-guru Ray Kurzweil has predicted that by the 2020s, we will have reverse-engineered the entire brain, and nanobots will be operating our consciousness. (Kurzweil 2005) But we need not venture into the world of science fiction to be staggered by the accelerating pace of technology today. We have already become familiar with cochlear implants that, by stimulating the auditory nerve, allow babies born deaf to hear; and we know something about computer-chip-sized devices that can take brain signals and rout them to robotic limbs, thereby allowing amputees' thoughts to control their movement. We are nearing the development of retinal implants to treat blindness. I have just read about research at the University of Florida on a new kind of neural implant that does not just receive instructions, but learns along with the brain. (Unattributed 2008)

Here are a few more examples: Scientists at the University of Michigan are working with neural interfaces coated with an electrically conductive polymer that can be coaxed into integration with brain tissue, thereby lessening the tissue damage caused by medical implants. (Singer 2008) An article in *Technology Review* explains that

> the goal is to get the electrodes to fully integrate with tissue by growing the
>
> coating after the electrode is implanted. The idea is that the polymer's hairlike
>
> fingers would reach into the tissue, extending beyond the dead zone surrounding
>
> the metal electrodes. "Imagine the cells are like M&Ms suspended in Jell-O," says
>
> [one of the Michigan scientists]. "We're growing the polymer around the M&Ms

and through the Jell-O." So far, the scientists have succeeded in growing the polymer in a piece of muscle tissue and a piece of mouse cortex. (Singer 2008)

This research, if successful, will pretty much erase the line between what is organic and what is artifactual (and with it, the philosopher's favorite line between what is mind-independent and what is mind-dependent).[1] (Baker 2004, Baker 2008)

There's more on the horizon. Erik Ramsey suffered a brain-stem stroke that left him totally paralyzed, but completely conscious. He was "locked in," as they say. He could communicate only by moving his eyes up or down, to answer yes or no. A neurologist implanted an electrode in his brain whose wires penetrate the part of the motor cortex that controls the motion of vocal muscles. When Ramsey thinks of making a sound, the implant captures the electrical firing of nearby neurons and transmits their impulses to a computer, which decodes them and produces the sounds. The aim is for Ramsey to communicate his thoughts to a computer that translates them into spoken words. As of a report in *Scientific American* in November 2008, Ramsey could only make a few simple vowel sounds, but his neurologist believes that he will recover his full range of speech by 2010**.** (Brown, 2008)

As a Practical Realist, what interests me in these examples is that they are real-life cases. They are not merely imagined. We do not have to indulge in long-range prediction or in outré thought experiments about zombies in order to be astonished by the

_____

[1] Or again, Duke University's Miguel A. L. Nicolelis, a pioneer in brain-machine interfaces, has begun to explore how monkeys connected to brain-controlled prosthetic devices begin to develop a kinesthetic awareness, a sense of movement and touch, that is completely separate from sensory inputs into their biological bodies. 'There's some physiological evidence that during the experiment they feel more connected to the robots than to their own bodies,' he says." (Stix, 2008)

melding of the biological with the nonbiological. And it is no surprise that these technological developments have raised anew the question of what we really are.

I want to ask: What's to become of us? Do these dizzying advances presage a future in which persons—traditionally conceived as self-understanding moral and rational agents—have disappeared?

From Extended Minds to Extended Systems

A number of philosophers have developed theories of what we are that seem congenial to this brave new world. Let's start by by considering the so-called 'Extended Mind' thesis (EM), the locus classicus of which is an article of the same name ("The Extended Mind") written jointly by David Chalmers and Andy Clark and published in *Analysis* in 1998.

The extended-mind thesis (EM) is the claim that mentality need not be situated just in the brain, or even within the boundaries of the skin. EM is quite radical: A mind is a collection of processes that easily extends to tools, programs, other minds, and language. Cognitive states may have all sorts of components—neural, bodily, environmental. As Chalmers put it, "I certainly do not think that there is anything privileged about skin and skull as boundaries for the mind." (Chalmers, xi). The heart of the extended-mind thesis is that we biological creatures can "couple" with nonbiological entities or features of our environment and thereby expand the entities that we are.

The idea of an extended mind can be seen as follows: Although thoughts are individuated by content, they have material vehicles that, so to speak, carry the content. Until recently, the vehicles have been taken to be only brain states, but according to EM,

the vehicles may also be spread out into the environment. Andy Clark characterizes EM as "the view that the material vehicles of cognition can be spread out across brain, body and certain aspects of the physical environment itself." (Clark 2005, p.1)

For example, beliefs are normally embedded in memory, but need not be. Here is a low-tech example from "The Extended Mind" that has become almost iconic. Consider Otto, who is impaired with Alzheimer's disease in a way that he cannot form new memories. He writes down what he wants to remember in a notebook, which he always carries. Suppose that Otto is on Fifth Avenue in New York City and is looking for MOMA. He knows that he cannot simply "search his memory" to recall the location of MOMA, but he automatically reaches for his trusty notebook and looks up the address: 53rd Street. The information in the notebook—just like the information stored in brain-based memory—"is reliably there when needed, available to consciousness and available to guide action, in just the way that we expect a belief to be." (Clark and Chalmers 1998, p. 13) Viewed through the lens of EM, the skin is seen as an artificial boundary. The vehicles for Otto's cognitive processing sweep out into the environment.

As Clark and Chalmers say, "[W]hen it comes to belief, there is nothing sacred about skull and skin. What makes some information count as a belief is the role it plays, and there is no reason why the relevant role can be played only from inside the body." [2] For some of his mental states and processes, Otto and his notebook are coupled; they form a cognitive system, all of the components of which are causally active. [3]

---

[2] The "relevant parts of the world are *in the loop*, not dangling at the other end of a long causal chain." (Clark and Chalmers 1998, p. 9)

[3] Adams and Aizawa 2008 argue that "coupling" is the wrong way to approach extended cognition altogether. The notion of coupling conflates ways that cognition causally depends on the environment with ways that cognition constitutively depends on the environment.

As Clark says, "[T]aken as a single, integrated system, Otto-and-the-notebook exhibit enough of the central features and dynamics of a normal agent having (amongst others) the dispositional belief that MOMA is on 53rd St. to warrant treating him as such." (Clark 2005, p. 7) Clark asks rhetorically, "[I]f an inner mechanism with this functionality [passive aspects of memory] would intuitively count as cognitive, then (skin-based prejudices aside) why not an external one?" (Clark, 2005, p. 7)

The criterion for extended cognition is given by the Parity Principle:

> [I]f a process in the world works in a way that we should count as a cognitive process if it were one in the head, then we should count is as a cognitive process all the same. (Chalmers, x)

In general, tools extend cognition. A tool, "even when [only] temporarily in use, is rapidly assimilated into the brains' body maps and is treated (temporarily) just like a somewhat less sensitive part of the body." For example, the receptive visual field of a macaque using a rake for as little as 30 seconds becomes elongated as if the rake were part of the arm. (Clark 2005, p. 8) Indeed, as reported in *Neuroscience*, "[S]imply holding a stick causes a remapping of far space to near space. In effect the brain, at least for some purposes, treats the stick as though it were a part of the body."[4]   So, use of a tool, even temporarily, changes neural maps. Neural plasticity "makes it possible for new equipment to be factored deep into both our cognitive and physical problem-solving routines." (Clark 2005, 9) So, we become mutable physical and cognitive hybrids—part biological and part artifactual.

---

[4] Berti and Frassinetti 2000. Quoted by Andy Clark 2005, p. 8.

Not only is there physically extended cognition, according to EM, there is socially extended cognition as well. As many have observed, their spouses are their external memory devices. My husband is a reliable depository of facts like the name of the peninsula across the harbor from Hong Kong Island. For my access to such facts, a proponent of EM may say, my husband and I are coupled. Coupling between agents is effected by language, among other things.[5] Language "is not a mirror of our inner states but a complement to them. It serves as a tool whose role is to extend cognition in ways that on-board devices cannot." Clark emphasizes that hybridization (like Otto-and-his-notebook) is quite normal.

If minds are spread out into the environment, then presumably we who have minds are too. Indeed, we have the minds of 'natural-born cyborgs' to use Clark's term—that is, "of systems continuously renegotiating their own limits, components, data stores, and interfaces." (Clark, 2008, 42) So, it is a quick step from extended minds to extended systems. "Extended systems or extended selves" are relatively transitory couplings of biological organisms and external resources. (Clark and Chalmers, 1998, 18) We are just shifting combinations of biological and non-biological elements. In that case, one's boundaries may fall outside one's skin. "Otto *himself* is best regarded as an extended system, a coupling of biological organism and external resources." (Clark and Chalmers, 1998, 18)

As Clark put it: "We have been designed, by Mother Nature, to exploit deep neural plasticity in order to become one with our best and most reliable tools. Minds like ours were made for mergers. Tools-R-us, and always have been." (Clark 2003, 6-7)

_____
[5] Coupling can also be effected by gestures, artifacts, or shared practices.

Indeed, Clark suggests, we need to reject "the seductive idea that all these various neural and nonneural tools need a kind of privileged user. Instead it is just tools all the way down." (Clark 2003,136) Similarly, as Clark summarizes Dennett's view, we are only "a grab bag of tools and an ongoing narrative: a story we, as the ensemble of tools, spin to make sense of our actions, proclivities, and projects." (Clark 2003, 138)

Although it is never explained how a grab bag of tools can spin a narrative to make sense of its actions, the idea that we are extended systems easily accommodates machine-brain interfaces—like neural implants that allow paralyzed people to interact with the world by computer. EM can take up Kurzweil's idea that we are about to merge with our technology, and go a step farther: I not just an extended system, but I am shifting arrays of biological and nonbiological elements.

Where Did All the People Go?

But are *we*—with or without neural implants or prosthetic limbs—well-served by the idea that we are extended systems or a grab bag of tools with no privileged user?   If we are extended systems, are we still persons? It is uncontroversial that persons are rational and moral agents capable of understanding what they do while they are doing it. Are extended systems rational and moral agents capable of understanding what they are doing while they are doing it? In his Forward to *Supersizing the Mind*, Chalmers suggests that he is now coupled with an iphone. He uses the pronoun 'I' to refer to himself, and says that he has a new part: the iphone.   The iphone is a part of the extended system that is Chalmers at time t.

Clark says that "(what we ordinarily think of as) the self is a hastily cobbled together coalition of biological and non-biological elements, whose membership shifts

and alters over time and between contexts." (Clark, 2004, 177) In this case, presumably Chalmers is not a person who is a proper part of various larger cognitive systems; but rather Chalmers *is himself* a series of larger cognitive systems hastily cobbled together from biological and non-biological elements.

So, in a world of extended systems, where would persons be? Let's ask: (1) Can extended systems be rational agents? (2) Can extended systems be moral agents? (3) Can extended systems understand what they are doing while they are doing it?

If anyone is a person, I suppose that Otto is. Recall that "Otto *himself* is best regarded as an extended system." (Clark and Chalmers, 18) Let's call the extended system that Otto (putatively) is 'Extended Otto', and let's call the "biological skinbag" – Clark's expression (Clark 2003, 4)—'Unextended Otto.' So, Unextended Otto is a component of Extended Otto.

The thesis at issue is that Otto, the entity of interest, is *identical* to Extended Otto —the integrated system of Unextended Otto and his notebook. To generalize, we might say that, on the extended-mind view, a person (like Otto) is a coupling of biological organism and external resources—but with a caveat: Otto is not just *an* extended system; he is different extended systems at different times. This caveat makes for complications, not least of which is that we seem forced into endorsing either four-dimensionalism or time-bound identity. For simplicity, I shall skip such complications here.

First question: If Otto is in fact an extended system, can he be a rational agent? If we accept the Parity Principle, perhaps we should say that the extended system did some cognitive processing; but the question is, Is the *extended system* the entity that acts

rationally? Unextended Otto is rational: He wants to go to MOMA and believes that the best way to get directions is to consult his notebook. Once he has the directions, he starts walking toward 53rd St. The fact that Otto is coupled with his notebook (if he is) is irrelevant to what he is doing. He would do the same thing no matter how he got the directions. He would still walk toward 53rd St. even if he had forgotten to put the notebook back in his pocket. Moreover, if Otto was always accompanied by Inga, and got the directions by asking Inga, then, by the Parity Principle, the extended system that (putatively) is Otto would have Inga as a component. But it is not the extended system having Otto and Inga as components that is *itself* a rational agent: Otto is a rational agent, and Inga is a rational agent. They are not jointly a rational agent.

Although I have provided only two examples—Otto-and-his-notebook and Otto-and-Inga—I think that it is safe to conclude that the rational agent, the person, is Unextended Otto, not Extended Otto. The extended-system approach to what we are does not seem to account for our being rational agents.

Here's what I think is going on: An Alzheimer's victim like Otto may require external resources for cognitive processing. But *how* Otto processes the information about the location of MOMA is irrelevant to his being a rational agent. This is so, because the cognitive processing is on a subpersonal level, but Otto's being rational is on the personal level. (Dennett 1969; Baker 2007, Ch. 11) The personal level is one at which someone can do what he or she intends to do; subpersonal levels of cognitive processing are ones at which intention (typically)[6] plays no role. Rather, subpersonal levels are the loci of mechanisms that make possible what we do at the personal level.

_____

[6] But see how cyber-artist Stelarc learned to control his electronic third arm. (Clark 2003, 115 -116).

To take Otto's consulting his notebook (which is at the personal level) to be part of cognitive processing (which is at a subpersonal level) is to conflate levels.[7]

This conflation is abetted by ambiguous use of the term 'cognitive' for processes at subpersonal levels as well as for activities (e.g., consulting a notebook, asking a question) at the personal level. I use the term 'cognitive processing' exclusively for the subpersonal level. Cognitive activities at the personal level are knowingly engaged in; cognitive processing at the subpersonal level goes on without any knowledge (or effort) on the part of the person.

Second question: If Otto is an extended system, can he be a moral agent? How can an extended system be held responsible for anything? Suppose that Otto always carried a knife and felt incomplete without it; he and his knife were an extended system. One day, he stabbed someone. In this instance, Extended Otto would be the combination of Unextended Otto and the knife. But the responsible party is not Extended Otto (an entity with a knife as a part). The moral agent is Unextended Otto, who is only a proper part of Otto, the extended system. (Certainly the police would not put the knife in jail.) So, the answer to the original question—If Otto is an extended system, can he be a moral agent?—is no. The moral agent would be the proper part of the extended system. However, moral agency does not accrue to proper parts of people. (A knife-wielding hand is not a moral agent.) So, not being a moral agent, Extended Otto is not a person either.

_____

[7] Interestingly, Dennett introduced the personal/subpersonal distinction in 1969. (Dennett, 1969, 93-96). Then in 1980, he reiforced it, "I understand English; my brain doesn't." (Dennett 1980, 429). In 2000, Jennifer Hornsby argued that Dennett had abandoned commitment to an explanatory personal/subpersonal distinction when he applied intentional-stance theory to subpersonal levels. (Hornsby 2000, 15).

Third question: If Otto is an extended system, can he understand what he is doing while he is doing it?   When Unextended Otto consults his notebook, he understands that he (Unextended Otto) is consulting his notebook.   Understanding is at the personal level. But Extended Otto's consulting his notebook is construed as cognitive processing that loops out into the world, and cognitive processing is on a subpersonal level. Cognitive processing is not something that an agent understands himself to be doing merely by doing it. The entity who understands what he is doing is Unextended Otto, who knowingly consults his notebook; but according to EM, the entity who is identical to Otto is Extended Otto, whose notebook is just a vehicle in his cognitive processing.

Again, to construe consulting a notebook as cognitive processing is to conflate the subpersonal (cognitive processing) with the personal (consulting a notebook). Consulting a notebook is something that Otto intends to do and knows that he is doing; cognitive processing is something that happens with no intention or knowledge on Otto's part.

On the extended-self approach, our boundaries are fluid, and do not coincide with the boundaries of organisms. If we are extended systems, we are not agents of any sort; we are only systems of diverse and fluctuating parts. The extended-system approach is very much at home with biotechnology now and in the future. But we—persons and agents—turn out to be problematic.   However, the proponents of the extended-system approach still have a card or two up their sleeves.

Interpretationism about Persons?

A popular strategy in metaphysics is to grant that strictly speaking ordinary objects do not really exist, but that sentences seemingly about them are usefully regarded as true. (Think of Peter van Inwagen's view of tables and chairs.)   Such a strategy has

made inroads into philosophical thinking about persons. Can this strategy rescue a perhaps deflated view of persons?

I'll consider first an early strategy of Daniel Dennett's and then a more recent strategy in which Clark joins Dennett. The early strategy, exhibited in Dennett's "Conditions of Personhood," is to understand persons in terms of the most robust conditions you like, and then to eviscerate the conditions in a way that makes them satisfiable come what may technologically.

Dennett offers six conditions of personhood—rationality, intentionality, being the object of the intentional stance, being able to reciprocate and treat others from the intentional stance, having a language, being self-conscious. Fine, so far. Now comes the deflation: Fulfillment of each of these conditions depends on the theorist's taking the intentional stance toward the would-be person.

The intentional stance is the stance from which we explain and predict behavior of others in terms of beliefs, desires and intentions. Anything whose behavior can be predicted from the intentional stance—us, termites, thermostats, anything—is an intentional system. Something "is an intentional system only in relation to the strategies of someone who is trying to explain and predict its behavior." (Dennett, 1978a, 3-4) To be an intentional system just is to be an entity whose behavior is predictable by attributions of beliefs, desires and intentions. There is nothing more to being an intentional system than to be *interpretable* as an intentional system.[8]

_____

[8] Incidentally, Clark expresses agreement with Dennett's interpretationism regarding "selves." (Clark, 2005, p. 10)

Almost any imaginable biotechnological innovations would allow the resulting entities to be interpretable as intentional systems, and very likely to satisfy Dennett's six conditions of personhood. But does satisfaction of Dennett's six conditions of personhood (as Dennett construes them) make an entity a person? Not if you think that 'person' is more than a façon de parler. The vocabulary of intentional-stance theory is non-referential. (Dennett, 1969, 93-96). Dennett is insistent that there are no ontological implications about what is discerned from the intentional stance. (Dennett 1978b, 273) The intentional stance reveals nothing about reality, only about useful interpretations.

Officially, Dennett's ontological view is one of "mechanistic physicalism." (Dennett, 1978b, 273) Ontologically speaking, reality is exhausted by the physical realm, the realm of the physical sciences.

However, from the *physical* stance, there are no persons: Our being persons presupposes the intentional stance. Dennett is no reductionist: What is discerned from the intentional stance is not reducible to the the physical realm. Indeed, as Dennett says, "[W]e, *as persons*, cannot *adopt* exclusive mechanism (by eliminating the intentional stance altogether)." (Dennett 1978d, 254) This is so, because to eliminate the intentional stance is to eliminate persons.

Dennett is very clear about his interpretationism. Intentional-systems theory is just a shortcut for explanation and prediction when we don't know what is really going on (that is, what is going on in the physical realm). (Dennett, 1969, 95) He says, "[W]e never need absolutely to ascribe any intentions to anything." (Dennett 1978b, 273) So, although we may find the intentional stance useful, nothing discerned from the intentional stance has a place in a full ontological account of reality. And since persons

are discerned only from the intentional stance, persons have no place in a full ontological account of reality. So, ontologically speaking, there are no persons.

I would expect Dennett to agree fully, and to say, "Or course, *ontologically speaking*, there are no persons, but we still have the intentional stance from which we can regard entities as persons—indeed, as responsible agents." However, the intentional stance seems unable to play its assigned role, because it is pulled in two different directions. On the one hand, "...we will never need absolutely to ascribe any intentions to anything...." But on the other hand, in some cases if one failed to take an intentional stance, Dennett says, one would be "missing something perfectly objective: the *patterns* in human behavior that are describable from the intentional stance and only from that stance, and which support generalizations and predictions." (Dennett 1987, 25).

This leads to a dilemma: Does the intentional stance tell us about anything that is grounded in reality or not? Officially, the answer is no: : "[W]hen we treat each other as intentional systems, [we are] relying, unconsciously, on the fact that to a remarkably good first approximation people are rational." (Dennett 2009b, 341) There is nothing in reality (i.e., the realm of physics) that answers to rationality. But if the intentional stance does not tell us about anything that is grounded in reality, I want to ask, how can we successfully use it, in Dennett's words, "to make sense of the behaviors of so many complicated things by considering them as agents"? (Dennett, 2009b, 349) Given Dennett's official view, it is a total mystery how the intentional stance *can* reveal patterns that are perfectly objective and support generalizations and predictions.

In work later than "Conditions of Personhood," Dennett proposes a somewhat different theory, a theory of consciousness that is applicable to persons. It too relies on

the intentional stance. In *Consciousness Explained*, Dennett (rightly, in my opinion) repudiated the notion of a Cartesian self, an internal audience for an internal theater. He proposed instead that our brains are continuously editing multiple drafts. Although I am totally convinced that there is no Cartesian self, I am not at all convinced of Dennett's alternative of appealing to diffuse neural processing together with a "user-illusion".

'User-illusion' is a label for presentation of a computer screen "since a user can't tell...how the particular virtual machine he's using is implemented in hardware." (Dennett, 1991, 216) I may think that I am looking at my desktop on the screen and moving the cursor to the place where my file is kept; but of course, that's not what's really going on at the level of the machine. It's only a user-illusion. (Note that the user-illusion is visible only from the intentional stance.) Dennett takes our supposed access to items of consciousness to be a similar user-illusion. The putative illusion is to suppose that there is a subject who is having colorful experience and is dealing with meanings when what is really going on is distributed processing. (Dennett, 1991, 312)

But this is another kind of error about levels. Dennett is quite aware of the distinction between personal and subpersonal levels; indeed, he introduced the distinction in 1969 (Dennett, 1969, 93), and in 2005 reiterated, "I understand English, my brain doesn't."[9] However—and here is where I think that the error lies—on Dennett's view, the only ontologically significant level is subpersonal, the physical level of neural circuitry. The personal level, saturated with normativity, is not anchored in reality; it is only heuristic. So, Dennett is committed to what I believe is a mistake—namely, that a

_____

[9] He continues: "—nor more particularly, does the proper part of it (if such can be isolated) that operates to 'process' incoming sentences and to execute my speech act intentions." Dennett 2005. Quoted from Dennett's comment on Searle in Behavioral and Brain Sciences 3 (1980): 429. http://ase.tufts.edu/cogstud/papers/bennetthacker.pdf

complete understanding of reality would require no reference to anything at the personal

level or to anything discernible only from the intentional stance.[10]

The intentional stance is central to Dennett's whole view. Indeed, as recently as

2009, Dennett said that he considers the idea of the intentional stance to be "probably my

most important contribution" to philosophy. (Dennett 2009a, 26) He says that

intentional-systems theory shows why the use of the intentional stance is "ubiquitous and

practically ineliminable" in computer science, evolutionary biology, and animal

psychology. (Dennett 2009b, 349).

I agree that intentional vocabulary is successful, ubiquitous and practically

ineliminable, but I do not see how its success can be explained by a theory that makes no

appeal to genuine reality. The fact that we can use the intentional stance so successfully

"in ignorance of the details of the processes going on in each other's skulls (and in our

own!)" (Dennett 2009b, 341) strongly suggests the contrary: that the intentional stance

reveals a level of genuine reality. In any case, that is not Dennett's view.

According to Dennett, the intentional stance is just a convenience. Its immunity

to refutation from technology is assured by its thinness: it makes no claims about reality.

It is for this reason that the intentional stance seems promising in the face of advances in

biotechnology, come what may. But if Dennett's view is correct, then we are persons

only by (optional) interpretation. In that case, persons ontologically have no future (or

past either), because strictly speaking, there aren't any persons, or agents or subjects of

experience. Not really. On Dennett's official view, a full and complete account of reality

---

[10] And it is an error even if one interprets subpersonal levels from the intentonal
stance, as Dennett does. Dennett says that "the same principles of interpretation are used
to endow subpersonal parts with contents as are used to endow whole persons." Dennett
1994, 528.

will mention neither persons nor anything that entails that there are persons.   So, interpretationism cannot save persons.

<center>**Soft Selves**</center>

Related to the notion of 'I' as a user-illusion, but put more gently, is the notion of soft selves. Can the notion of a soft self salvage persons?   A soft self is, in Andy Clark's words, "a rough-and-tumble, control-sharing coalition of processes—some neural, some bodily, some technological—and an ongoing drive to tell a story, to paint a picture in which "I" am the central player." (Clark 2003,138) Our "true nature" is to be "distributed decentralized coalitions," "palpitating biotechnological hybrids". (Clark, 2003, 139)

However, Clark, as I suggested earlier, conflates levels. He suggests that to understand ourselves, we must understand subpersonal cognitive processing. He says,

> "[W]e suffer from a chronic tendency to misconstrue the relations between our self-conscious 'choosings' and the vast webs of non-conscious processing activity (all those whirrings and grindings of machinery, neural and perhaps non-neural, internal and perhaps external) that also structure and determine our own actions and responses. Until we form a better, more consistent image of the relationship between these factors, we cannot hope to know ourselves." ("Soft Selves," p. 10).

That is to say, we cannot hope to know ourselves until we understand the relations between our thoughts and actions, on the one hand, and the subpersonal machinery, on the other.   But a glance at history suggests that this is wrongheaded. Knowing ourselves in the sense that people care about—as Dennett noted—does not require knowing

anything about the subpersonal machinery that supports our self-knowledge. There is a venerable history—at least from St. Augustine forward—of people coming to know themselves without knowing neuroscience.

The question that I am asking about the future of persons is whether the level at which we pretheoretically suppose that we exist—the level at which we traditionally have sought self-knowledge—is real. In *The Metaphysics of Everyday Life*, I argued that the personal level *is* real and distinct from other levels. I am now suggesting that, if we were soft selves or grab bags of tools, the personal level would not be real and distinct from other levels. It would dissolve into, well, a grab bag of tools. And a grab bag of tools, I have argued, is not an agent who can understand herself—indeed, not an agent of any kind.

In any case, the appeal to soft selves presupposes the intentional stance. And as we have seen, the intentional stance is simply a scheme for interpretation, without any ontological implications. Since something is a tool only if it has a certain purpose, there are no tools—much less grab bags of tools that spin stories about themselves—except from the intentional stance. So, our "true nature" cannot be even to be a grab bag of tools. From the physical stance, tools disappear along with persons. Tools are visible only from the design and intentional stances, but reality is ontologically exhausted by the physical realm. So, thinking of persons as grab bags of tools or as soft selves provides not even a toehold for persons in reality.

Securing the Future of Persons

The alternative that I offer is a Practical-Realist approach to reality: Fundamental reality is not just "given"; it is not static, but it changes over time—with evolution of the universe and evolution of the species, with inventions that we contribute. A world with organisms is fundamentally different from a world without organisms, even if the two worlds have the same kinds of atoms and molecules distributed slightly differently. A world with iphones is fundamentally different from a world lacking electronic devices. An important ontological change was the evolution of persons—the coming into existence of subjects of experience and agents who can understand themselves and what they are doing while they are doing it. The existence of such beings is not a matter of interpretation; a being who can think of herself from the first-person *as* herself is a person, whether she or anyone else thinks so or not.

On my own view (as you no doubt know), we are essentially persons, constituted by bodies that have mechanisms that support our first-person perspectives. Typically, persons are constituted by human organisms, but there is nothing to prevent the human organisms from having nonbiological prostheses, including neural prostheses. Indeed, there is nothing to prevent a human organism from being replaced over time by a bionic body. (Animalists would have to say that with enough nonbiological replacements, you would cease to exist; but I agree with Clark and Dennett that we are not essentially skin-bound organisms.) The composition of a person's body is not crucial so long as it supports a first-person perspective.

Although I suppose that we may be constituted by bodies that are partly or even wholly nonbiological, I do not suppose that every time we pick up a pencil, we are

constituted by a body with a new part. Unlike Clark, I would distinguish prostheses that are relatively permanent from tools that are used on occasion. I consider our relation to tools used on occasion to be like outsourcing. When a company outsources its accounting, it does not take on the accounting firm as part of itself; when the US Army outsources some of its combat activities to a private firm like Blackwater, employees of Blackwater do not become part of the army. Similarly, when you use a pencil and paper for a long-division problem, the pencil and paper do not temporarily become part of your body. But neural implants are another thing. Bodies with neural implants are stable; they are not shifting coalitions. A person's body is connected and functionally integrated. So, I would take a neural implant (but not a pencil and paper) to become part of a person's body—part of what constitutes the single individual person.

Natural beings like us all start out constituted by human organisms, but over time, with enough nonorganic replacements, we could come to be constituted by different kinds of bodies. We would still exist as long as, whatever bodies constituted us, those bodies sustained our first-person perspectives.

Conclusion

Subpersonal states and entities are subject to hybridization: they may have nonbiological elements (in or beyond the body) that play essential roles in cognitive processing. But, on my view, bionic components of subpersonal states and processes are no threat to the reality of persons. A person endures as long as the subpersonal mechanisms that support a first-person perspective function correctly—whether those subpersonal mechanisms are biological or bionic or something else.

We have seen the perils of ignoring the line between the personal and the subpersonal levels.   If we keep straight the distinction between levels and treat the personal level with ontological seriousness, we can applaud the technological advances that greatly improve the lives of paralyzed and debilitated patients without losing sight of the person, the whole patient who should be ultimately the object of interest.[11]

Revised April 28, 2009

References

Adams, F. and K. Aizawa. 2008. *The bounds of cognition*. Malden, MA: Blackwell Publishing.

Baker, L.R. 2004. The ontology of artifacts. *Philosophical Explorations*, 7: 99–111.

Baker, L.R. 2007. *The Metaphysics of Everyday Life: An Essay in Practical Realism*. Cambridge: Cambridge University Press.

Baker, L.R. 2008. "The Shrinking Difference between Artifacts and Natural Objects," *American Philosophical Association Newsletter on Philosophy and Computers*, 7(2).

Berti,A. and Frassinetti, F. 2000. ''When Far Becomes Near: Re-mapping of Space by Tool Use'', *Journal of Cognitive Neuroscience* 12, pp. 415–420

Brown, Alan 2008. "Putting Thoughts into Action: Implants Tap the Thinking Brain" Alan S. Brown, Scientific American. http://www.sciam.com/article.cfm?id=putting-thoughts-into-action. Accessed Mar. 11, 2009

Chalmers, David, 2008. "Forward" in Supersizing the Mind by Andy Clark (Oxford: Oxford University Press): ix-xvi.

Clark, A. 2003. *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.

Clark, A. 2004. Author's Reply. . Review Symposium on Andy Clark's *Natural-Born Cyborgs. Metascience* 13(2): 169-181.

---

[11] Many thanks to Hilary Kornblith and to Gareth B. Matthews for commenting on a draft of this paper. Also, I am grateful to the organizers of the conference on my work at the University at Buffalo, where this paper was presented on April 24, 2009.

Clark, A. 2005. "Intrinsic content, active memory and the extended mind". *Analysis* 65.1: 1-11.

Clark, A. 2007. "Soft Selves and Ecological Control" in Distributed Cognition and the Will, Don ross, David Spurrett, Harold Kincaid, and G. Lynne Stephens, eds. (Cambridge MA: MIT).

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action and Cognitive Extension* (Oxford: Oxford University Press).

Clark, A. and D. Chalmers. 1998. "The Extended Mind," *Analysis* 58.1: 7-19.

Dahlbom, Bo. 1993. "Mind is Artificial," in *Dennett and his Critics*, Bo Dahlbom, ed. (Oxford:Basil Blackwell Ltd): 161-183.

Dennett, D.C. 1969. Content and Consciousness, 2nd ed. 1993 (London: Routledge).

Dennett, D. C 1978a. "Intentional Systems" in *Brainstorms: Philosophical Essays on Mind and Psychology* ((Montgomergy VT: Bradford Books): 3-22.

Dennett, D.C. 1978b. "Conditions of Personhood," in *Brainstorms* ((Montgomergy VT: Bradford Books): 267-285.

Dennett, D.C. 1978c. "Mechanism and Responsibility," *Brainstorms* ((Montgomergy VT: Bradford Books): 233-255267-285.

Dennett, D.C. 1978d. Where Am I*? Brainstorms* (Montgomergy VT: Bradford Books): 310-323.

Dennett, D.C. 1980. Reply to Searle. *Behavioral and Brain Sciences* 3

Dennett, D.C. 1987. "True Believers" in *The Intentional Stance* (Cambridge MA: MIT/Bradford): 13-35.

Dennett, D.C. 1991. *Consciousness Explained* (Boston: Little Brown).

Dennett, D.C. 1994. "Get Real," *Philosophical Topics* 22: 505-568.

Dennett, D.C. 1998. "Self-Portrait" in *Brainchildren: Essays on Designing Minds* (Cambridge MA: MIT/Bradford)

Dennett, D.C. 2005. "Philosophy as Naive Anthropology: Comment on Bennett and Hacker, final draft, October 18, 2005. http://ase.tufts.edu/cogstud/papers/bennetthacker.pdf. accessed April 18, 2009..

Dennett, D.C. 2009a. "Daniel Dennett," *Mind and Consciousness: 5 Questions*, Patrick Grim, ed. (Automatic Press VIP): 25-30.

Dennett, D.C. 2009b, "Intentional Systems Theory" in *The Oxford Handbook of Philosophy of Mind*, Brian McLaughlin, Ansgar Beckermann, Sven Walter, eds. (Oxford: Clarendon): 339-350.

Hornsby, J. 2000. "Personal and Subpersonal: A Defense of Dennett's Early Distinction," *Philosophical Explorations* 3: 6-24.

Hurley, S. L. 1998. *Consciousness in Action*. Cambridge MA: Harvard Unviersity Press.

Kurzweil, Ray 2005. *The Singularity is Near*. New York: Penguin Books.

Preston, Beth. 1998. "Cognition and Tool Use," *Mind & Language* 13: 513-546.

Singer, Emily 2008. "Growing Neural Implants: New Approaches Could More Seamlessly Integrate Medical Devices into the Body," in *Technology Review* (published by MIT) http://www.technologyreview.com/Biotech/21087/?a=f

Stix, Gary, 2008. "Jacking into the Brain—Is the Brain the Ultimate Computer Interface?" by Gary Stix, Nov. 2008 Sci. Am. http://www.sciam.com/article.cfm?id=jacking-into-the-brain&page=2 accessed Mar. 11, 2009.

Unattributed 2008. "Researchers Develop Neural Implant the Learns With the Brain," *Physorg.com*. http://www.physorg.com/news133535377.html accessed March 7, 2009.