

Mind Uploading: A Philosophical Counter-Analysis

Massimo Pigliucci

1 Introduction

The possibility of “mind uploading” and related ideas about the nature of consciousness have been discussed relatively little in the technical philosophical literature. One prominent exception is the analysis by David Chalmers (see Chapter 6 above), first presented in the *Journal of Consciousness Studies* (2010). The present chapter takes Chalmers’ arguments to be one of the best defenses of the concept of mind uploading and uses them as the starting point for a philosophical counter-analysis.

My criticisms are based in part on the contention that supporters of uploading do not take seriously enough the apparent fact that consciousness is – to the best of our knowledge – a biological phenomenon that is unlikely to be entirely (or even largely) substrate-independent, although this does not imply that the current biological substrates are the only ones capable of generating consciousness. Indeed, substrate independence of the type envisioned by Chalmers implies a form of dualism that should be unacceptable in modern philosophy of mind. Moreover, I argue that the Computational Theory of Mind underlying the concept of mind uploading is incomplete and thus only partially applicable to human-like intelligence/consciousness. Finally, I maintain that mind uploading – even if somehow technically possible – would be at best a form of mental cloning, most certainly *not* a way to preserve one’s own consciousness.

2 What I am going to attempt here

There has been quite a brouhaha about mind uploading and the related concept of a (soon, but not quite yet) forthcoming “Singularity” episode in the history of artificial and human intelligence. Indeed, this book is one example of what the fuss is all about, with a preponderance of contributions from Singularity and mind-uploading (MU) theorists, several of whom are operating at least partially outside classical academic institutions, for instance at the Singularity Institute (now the Machine Intelligence Research Institute), the Singularity University, and so forth.

This chapter will set aside the question of whether a Singularity will occur (in the near future or ever), to focus on the closely related issue of MU, specifically as presented by one of its most articulate proponents, David Chalmers (2010; Chapter 6 in this collection). MU and the Singularity share a common basis in that they both rely on the success of a strong type of artificial intelligence research program, which is in turn based on some version of the Computational Theory of Mind (CTM). I will proceed in the following fashion: first, I will briefly recall Chalmers’ main arguments; second, I will argue that the ideas of MU and CTM do not take seriously enough the fact that consciousness is a biological phenomenon, as pointed out in different contexts by both John Searle and Jerry Fodor, among others; third, I will stipulate *for the sake of argument* that MU is somehow possible, and show that it would, at best, amount to a sort of mind cloning, but most certainly not to a preservation of anyone’s consciousness in a medium different from one’s own brain.

3 What Chalmers says, redux

Chalmers first presented his broad ideas about the Singularity and MU in a paper published in the *Journal of Consciousness Studies* (Chalmers 2010), and some substantive criticisms of the paper have already appeared in print (e.g., Prinz 2012). Limiting ourselves to MU, as Chalmers does in his abridged discussion of the issues in this volume, we can summarize his arguments as follows.

Chalmers begins with a taxonomy of possible methods of MU, including destructive, gradual, and non-destructive uploading, but quickly and correctly zeros in on the idea that there are two components to a philosophical analysis of MU: first, the issue of the nature of consciousness; second, the question of personal identity. For MU to be possible, it has to be true that consciousness is a particular type of computational phenomenon; and for

MU to be worth our while we also need to agree on a particular view of personal identity. Section 4 of this chapter will deal with the first issue, leaving the second for section 5.

To my way of seeing the problem of consciousness, it is astounding that Chalmers begins with an admission that ought to halt him right in his tracks:¹

The issue here is complicated by the fact that our understanding of consciousness is so poor. No one knows just why or how brain processes give rise to consciousness There is nothing even approaching an orthodox theory of why there is consciousness in the first place. Correspondingly, there is nothing even approaching an orthodox theory of what sorts of systems can be conscious and what systems cannot be.

But that apparently does not stop Chalmers and other supporters of MU from proceeding *as if* we had a decent theory of consciousness, and by that I mean a decent *neurobiological* theory (as opposed to a general philosophical account). Odd, to say the least.

Chalmers then portrays an antagonism between biological and functionalist views of consciousness, a distinction that is crucial to all his further arguments:

Biological theorists of consciousness hold that consciousness is essentially biological and that no nonbiological system can be conscious. Functionalist theorists of consciousness hold that what matters to consciousness is not biological makeup but causal structure and causal role, so that a nonbiological system can be conscious as long as it is organized correctly.

I will show that Chalmers here does not provide his readers with an accurate view of what biological theorists actually say, and that moreover his view of functionalism appears to work only because it is founded on an equivocation about what functionalism *means*.

The next step in the argument is to present readers with a thought experiment (since actual experiments on MU cannot be done): What would happen if we gradually replaced the components of a biologically conscious system (i.e., neurons, etc.) with *functionally equivalent* (more on this later) non-biological components? Chalmers tells us that there are three possibilities here, regarding consciousness:

It might suddenly disappear, with a transition from a fully complex conscious state to no consciousness when a single component is replaced. It might gradually fade out over more than one replacements, with the complexity of the system's conscious experience reducing via intermediate steps. Or it might stay present throughout.

He then (too) quickly concludes that both sudden disappearance and gradual fading “seem implausible,” thereby accepting the scenario most favorable to MU: consciousness is maintained throughout the replacement procedure.² The upshot is that, according to Chalmers, consciousness is an organizational invariant (i.e., it is substrate-independent), and MU is therefore possible.

He then moves to the second issue, that of personal identity. Again he contrasts an “optimistic” view of MU, according to which mind uploading is a form of survival of the individual, to a “pessimistic” view, according to which MU is a technologically sophisticated form of suicide. The discussion in turn hinges on what criteria we use for personal identity, a notoriously contentious subfield within metaphysics (Olson 2010). Chalmers presents his readers with three choices: biological theories, psychological theories, and “closest-continuer” theories (see Chapter 6 above for a brief explanation), commenting that someone holding a biological view of personal identity is likely to be a pessimist about MU, those adopting a psychological view are going to be optimists, and those espousing a closest-continuer view will hedge their bets depending on whether the upload is destructive or not.

Although he proceeds by examining arguments about different forms of uploading, Chalmers admits that he does not have a settled view of personal identity, and that he would “hesitate” before undergoing destructive uploading. The last part of Chalmers’ chapter goes into science-fictional scenarios involving mind reconstruction from cryonic preservation or, even more fancifully, from a combination of leftover genetic material and data records. I will not entertain those any further, largely on the ground that I will show that Chalmers gets into serious trouble way ahead of reaching those far points in the discussion.

4 Consciousness, computation, and mind uploading

The fundamental premise of Chalmers’ (and others’) arguments about MU is some strong version of the Computational Theory of Mind, which is defined by Horst (2009) as a “particular philosophical view that holds that the mind literally is a digital computer ... and that thought literally is a kind of computation.” This is of course not the place for an in-depth critique of the CTM, but I’d like to point out a number of *prima facie* reasons why it simply cannot do the sort of job needed by supporters of MU.

To begin with, there is much misunderstanding of what “computation” stands for here. Jerry Fodor, one of the originators (Fodor 1975) – together with Hilary Putnam (1960) – of the CTM, has expressed bewilderment at how it has generated a totalizing view in philosophy of mind that simply

cannot be right: distinguishing between “modular” and “global” mental processes, and arguing that the former, but not the latter (which include consciousness) are computational in any strong sense of the term, he has commented that it “hadn’t occurred to me that anyone could think that it’s a very large part of the truth; still less that it’s within miles of being the whole story about how the mind works” (Fodor 2000, a direct response to Pinker 1997). If Fodor is right, then the CTM cannot be a complete theory of mind, because there are a large number of mental processes that are not computational in nature.³ In turn, this would mean that MU is not possible, since it is premised on the idea that minds are essentially (and completely) computational in nature.

Indeed, Copeland (2002) traces much trouble with the concept of computation in philosophy of mind to a widespread and persistent misunderstanding of the famous Church–Turing thesis about universal computability. Turing’s version of the thesis says that logical computing machines, which eventually became known as Turing machines, can do anything that can be described as a rule of thumb or purely mechanical (“algorithmic”); the Church version says that a function of positive integers is effectively calculable only if recursive (see Copeland 2002 and references therein), which turned out to be equivalent to Turing’s statement.

None of the above implies the sort of much stronger declarations that have been made by computationally inclined philosophers of mind. Specifically, the following thesis is *not* established by Church–Turing: Whatever can be calculated by a machine (working on finite data in accordance with a finite program of instructions) is Turing-machine-computable (Copeland 2002). Moreover, Turing was explicitly interested in what *cannot* be computed (i.e., in the limits of computability), and we now know of a number of problems that fall into this category (other than the classical halting problem). And yet mind computationalists often talk (see Copeland 2002 for a number of textual examples) as if Church–Turing has essentially established the CTM, and therefore indirectly also the possibility of mind uploading.

An additional problem for the CTM is that it is often not clear whether its supporters are arguing that a computer can *simulate* a human mind or that it can function as a human-type mind. The difference is crucial. Searle (1980) drew the analogy with other biological processes, such as, say, photosynthesis. We can most certainly simulate what is going on during photosynthesis, down to the quantum level, as it turns out. But there is a crucial thing we don’t get out of simulated photosynthesis: sugar, the only outcome that matters to real plants.

Searle did not claim that this somehow shows that consciousness is impossible outside of biological systems (contra what is implied by Chalmers in his summary of biologically informed positions in philosophy of mind). Rather,

the question is empirical, and cannot, therefore, be settled by any thought experiment, regardless of how ingenious it is. It can be rephrased thus: How strong are material constraints on the production of the phenomenon of consciousness? And the answer is: We don't know, but they are unlikely to be weak.

Another way to put the issue is this: Is minding (the conscious thinking activity of the brain) more like photosynthesis or more like the sort of abstract symbol shuffling that characterizes the operations of an electronic computer? At the very least, this is an open question that cannot simply be brushed aside by hard-core computationalists. And let us remember that – biologically speaking – any activity of the brain does take place by way of, and results in, physical products (neurotransmitters, electrical impulses, chemical interactions, and so forth). To claim that these are only incidental – as opposed to constitutive – of the ability to be self-conscious is to veer deeply into a form of dualism that ought to make contemporary philosophers at least a little bit uncomfortable.

Other than the problem that there usually is a difference between X and the simulation of X, there is an additional reason why I think biological naturalism is a better way to think about consciousness than hard-core computationalism. Consider life itself: there is no question – I hope – that being alive is a qualitatively different state of matter from not being alive, just as having consciousness is a qualitatively different state from not having it. (This, of course, has no pseudo-mystical or vitalistic implications at all, and does not negate the basic fact that both rocks and biological organisms are made of atoms.) Now, although we know quite a bit about the chemistry of life on earth, it unfortunately is so far the only example we have in the entire universe. Astrobiologists have therefore been somewhat free to speculate about possible alternative chemistries capable of producing life forms. Setting aside the very thorny question of what we mean by “life” (which, of course, is analogous to the thorniness of defining consciousness), such speculations have focused on silicon as pretty much the only other *potential* game in town (and by town, I mean the universe), aside from carbon (e.g., Schulze-Makuch and Irwin 2006). That is because we know enough about the other elements of the periodic table that we are reasonably sure that their chemistry cannot lend itself to anything like the functionality necessary for the complex metabolism and reproduction typical of living beings. Indeed, there are good reasons to doubt the viability even of silicon-based life forms, given the much more restricted chemical flexibility of silicon when compared to carbon.

Which brings us to the question of whether and in what sense it may be possible to extract or “upload” human consciousness to a mechanical device made of something other than carbon-based neurons. There are two further

issues here. First, what do we *mean* when we are asking that question? Second, assuming that what we mean is something coherent, is it technically possible? The latter question is, again, empirical in nature, and I think the best attitude a philosopher can take towards it is to wait and see whether science and technology will be able to provide us with an answer. But an exploration of the first question lays bare a number of troubling equivocations in Chalmers' (and others') position concerning the CTM and MU.

The most obvious equivocation concerns whether we are talking – as Chalmers does through most of his chapter – about replacing carbon-based components in a human brain with functionally equivalent components made of something else (most obviously, silicon), vs. whether we are considering more exotic possibilities, such as somehow transferring human consciousness inside a computer (as he hints in his 2010 paper). The first scenario requires “only” a convincing (empirical) demonstration that, say, silicon-made neurons can function just as well as carbon-based ones, which is, again, an exclusively empirical question. They might or might not, we do not know. What we do know is that not just any chemical will do, for the simple reason that neurons need to be able to do certain things (grow, produce synapses, release and respond to chemical signals) that cannot be done if we alter the brain's chemistry too radically. The second scenario, instead, brings us right back to a curious form of dualism, since it essentially assumes that consciousness is substrate-independent. I find this position downright bizarre, and not at all disanalogous to claiming that photosynthesis, or life itself, is likely to be substrate-independent. Here I follow Searle's (2008) “biological naturalism” position and demand that my colleagues take biology a bit more seriously, since after all consciousness – so far as we can tell – is a biological phenomenon. Needless to say, the second scenario also requires a very strong version of the CTM, which we can reject for the various reasons already mentioned (Fodor's “not everything in the mind is computable” objection; misunderstanding of Church–Turing; etc.).

There is another, more subtle, type of ambiguity at work here. Chalmers (like others) makes a lot out of his “functional” (as opposed to biological) approach to consciousness. As he puts it: “Functionalist theorists of consciousness hold that what matters to consciousness is not biological makeup but causal structure and causal role, so that a non-biological system can be conscious as long as it is organized correctly.” But what does Chalmers mean by “causal structure and causal role”? This phrase may very well hide the necessity that neurons be made of certain particular materials, otherwise they won't work – surely, the physico-chemistry of a system has *something* to do with its causal structure and role, no? If, for instance, we replace all the carbon in a brain with, say, krypton (a so-called “noble” gas) the causal functionality of the system will be irreparably disrupted. That's because

functionality isn't just a result of the proper arrangement of the parts of a system, but also of the types of materials (and their properties) that make up those parts. But if we expand our concept of functionality in philosophy of mind to take this objection into account, it is no longer clear what is the difference between a functionalist and a biological perspective. Contra Chalmers (as cited above), people like Searle (and myself) who hold to a biological approach to consciousness are not claiming that only currently known biological systems are capable of conscious states. Indeed, already in his classic Chinese room paper Searle stated this explicitly:

Part of the point of the present argument is that only something that had those *causal powers* could have that intentionality. Perhaps *other physical and chemical processes could produce exactly these effects*; perhaps, for example, Martians also have intentionality but their brains are made of different stuff. That is an empirical question, rather like the question whether photosynthesis can be done by something with a chemistry different from that of chlorophyll. (Searle 1980; my emphasis)

So functionalists cannot have it both ways: either they are saying that the substrate truly does not matter at all, in which case they are endorsing some sort of dualism based on an untenably strong version of CTM and a misunderstanding of Church–Turing; or they are admitting that substrate does matter, in which case they are really falling back onto some sort of (perhaps expanded) biological view of consciousness. And recall that, according to Chalmers' own reckoning, a biologically inclined philosopher will *ipso facto* be a pessimist about mind uploading.

5 Mind uploading, personal identity, and Kirk's death by transporter

We now turn to the second of Chalmers' fundamental questions about MU, that of personal identity. As I mentioned above, he provides his readers with three families of alternatives: biological views, psychological views, and what he terms, following Robert Nozick, "closest-continuer" views, stating – correctly, I think – that adopting the first type leads one to think of MU as a form of suicide, preferring the second one inclines people toward seeing MU as a type of preservation of personal identity, and going for the third option ends up in a type of bet-hedging that depends on the form of uploading (destructive or not). While Chalmers is correct in *logically* separating from the question of personal identity the question whether consciousness is a type of computational phenomenon, the two are obviously deeply

related when it comes to mind uploading. It would be odd, for instance, to endorse a biological view of consciousness, leading to skepticism about the CTM, and yet somehow switch to a psychological view of personal identity and gingerly walk into an uploading machine.

Quite honestly, I have always found the standard way of framing the issue of personal identity in metaphysics somewhat odd: it seems that for unfathomable reasons many philosophers think of the types of alternatives laid out by Chalmers as mutually exclusive, while a more commonsensical approach would suggest that they are complementary. It will come as no surprise that it is Searle, again, who takes this latter approach, even in introductory treatments of the subject matter (Searle 2005). He suggests that spatio-temporal continuity of the body is certainly a major criterion we all use to decide about personal identity (so, for instance, our monozygotic twins are not us, even when they are for all effective purposes genetically and physically identical to us, and even if they have the same thoughts – including memories – and general personality as we do). But spatio-temporal continuity is not sufficient to account for the first-person *experience* of personal identity; we *also* (as opposed to *instead*) need memory and personality. And Searle is quick to recognize that there are partial exceptions to all of these criteria (e.g., people whose memory is lost or personality profoundly altered as a consequence of accident or disease). Indeed, these exceptions help us to focus on what we count as personal identity *under normal circumstances*.

And there lies the rub for all discussions about personal identity: we make a certain sense of the concept because human beings are a particular type of biological species (biology again!). If we were a very different type of being, we would have a very different concept of personal identity, or perhaps even no such concept at all. There is, in other words, no *metaphysical fact* of the matter about personal identity – which perhaps helps explain both Chalmers' own lack of commitment on the subject and the fact that neo-scholastic philosophers⁴ simply cannot agree on it.

As Chalmers himself observes, the issue here is analogous to the famous “transporter” problem inspired by the futuristic technology of the *Star Trek* series, and the question is often framed informally in philosophical circles as to whether Kirk (the original captain of the *Enterprise*) dies as soon as he steps in the transporter. Of course, for the dramatic purposes of the series, the answer is no: Kirk is scanned by the transporter's computer, information about the position of every one of his molecules is transmitted to the destination, and Kirk is “reassembled” (really, assembled *de novo*) on the other side (Robinson and Riley 2010) – which, incidentally, makes teletransportation a destructive type of uploading.⁵

So, in the spirit of anti-neo-scholastic philosophy, let us use a naturalistic, commonsensical approach for a change, and briefly examine the situation

from the point of view of someone about to step into a transporter, or to allow what he thinks is going to be an instance of MU. To begin with, we can do away with any bet hedging: if your decision about pushing the button depends on whether the procedure is going to be destructive or not, I submit that you really do not believe that what's about to be created is a continuation of you. At best you think there will be a copy of you in existence at the end of the process – either *alongside* you as the original (nondestructive uploading), or not (in which case you'll be dead, not transported or uploaded). And indeed, it seems to me that this is pretty much all there is to say on the matter: if it is possible to do the transporting or uploading in a non-destructive manner, *obviously* we are talking about duplication, not preservation of identity. But if the only difference between the two cases is that in one the original is destroyed, then how on earth can we avoid the conclusion that when it comes to destructive uploading we just committed suicide (or murder, as the case may be)? After all, *ex hypothesi* there is no substantive difference between destructive and non-destructive uploading in terms of end results. This means that even by allowing the possibility of MU for the sake of argument, we end up rejecting it by way of what ought to be straightforward (and certainly commonsensical, judging from Chalmers' own hesitation) considerations about personal identity. I realize, of course, that to some philosophers this may seem far too simple a solution to what they regard as an intricate metaphysical problem. But sometimes even philosophers agree that problems need to be dis-solved, not solved: "Since everything lies open to view there is nothing to explain" (Wittgenstein 1953: 126).

6 Conclusion

I have argued that the most reasonable answers to Chalmers' two questions concerning mind uploading are pessimistic: (1) No, the mind "doesn't work that way," as Fodor quipped in the title of his book. That is, we have serious reasons to doubt a straightforward computational account of consciousness, and very good *prima facie* reasons to endorse a moderate version of biological naturalism. (2) No, personal identity (as understood by normal human beings) would not be maintained after a process of MU, which would result either in a duplication of the original individual (if non-destructive) or in the death of the original (if destructive).

The first conclusion, it should go without saying, does *not* imply that we will never be able to build "conscious machines." After all, human beings *are* conscious (biological) machines, if by machine one means something built in a manner that doesn't invoke mystical or supernatural processes. And,

again, it should not even be interpreted as saying that the only way to get consciousness is via carbon-based life forms. It only says that consciousness is unlikely to be substrate-independent, and that there plausibly are very strong constraints on the sort of substrate that is suitable to the process. The proof, of course, is in the pudding: the moment someone is capable of producing human-type thinking and consciousness in a computer, I will stand corrected (though there is the pesky issue of how would we know, a particular version of the problem of other minds: Hyslop 2009; of one thing I'm sure: a simple Turing test is not going to do it).

Even if and when we are able to produce a computerized version of human consciousness, however, I will still strongly advise against pushing the “upload” button: in the light of a naturalistic and commonsensical approach to personal identity, you would simply be committing a very technologically sophisticated (and likely very, very expensive) form of suicide.

Acknowledgments

I would like to thank the editors of this volume for inviting me to contribute as a skeptic to their collection; that's the way intellectual discourse advances. I also wish to thank Leonard Finkleman for thoughtful and helpful comments on a previous draft of this manuscript. Finally, sincere thanks to many of the readers of my blog, *Rationally Speaking*, who have challenged and continue to challenge me in clever and instructive ways whenever I write about these (or, really, any other) topics.

Notes

- 1 All quotations from Chalmers are from the version of his arguments presented in the chapter in this volume.
- 2 Chalmers here rehashes his “dancing qualia” argument, presented in Chalmers (1995) and criticized by various authors, e.g., van Heuveln et al. (1998).
- 3 Unless one expands the definition of computation to encompass so-called pan-computationalism (Piccinini 2010 and citations therein), the idea that everything computes; at that point, however, talk of computing becomes vacuous and at any rate not particularly helpful to understanding the human mind.
- 4 I am using the term in the admittedly pejorative way articulated by Ladyman and Ross (2007) and exemplified in Chalmers et al. (2009), and in opposition to naturalistic or “scientific” philosophy: Ross et al. (2013).
- 5 Of course, also for dramatic purposes, Kirk is always accompanied by one or more “red shirts” from the *Enterprise* security team, who are almost guaranteed to become casualties before the end of the episode. For a complete analysis backing up this particular point, see Barsalou (2013).

References

- Barsalou, M. 2013. Keep your shirt on: A Bayesian exploration. *Significance Magazine*, <http://www.significancemagazine.org/details/webexclusive/4381371/Keep-your-redshirt-on-a-Bayesian-exploration.html> (accessed October 5, 2013).
- Chalmers, D.J. 1995. Absent qualia, fading qualia, dancing qualia. In T. Metzinger, ed., *Conscious Experience*. Exeter: Imprint Academic, pp. 309–328.
- Chalmers, D.J. 2010. The Singularity: A philosophical analysis. *The Journal of Consciousness Studies* 17: 7–65.
- Chalmers, D.J., Manley, D. and Wassermann, R. 2009. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.
- Copeland, B.J. 2002. The Church–Turing thesis. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/church-turing/> (accessed October 5, 2013).
- Fodor, J. 1975. *The Language of Thought*. New York: Thomas Crowell.
- Fodor, J. 2000. *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: MIT Press.
- Horst, S. 2009. The computational theory of mind. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/computational-mind/> (accessed October 5, 2013).
- Hyslop, A. 2009. Other minds. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/other-minds/> (accessed October 5, 2013).
- Ladyman, J., and Ross, D. 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Olson, E.T. 2010. Personal identity. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/identity-personal/> (accessed October 5, 2013).
- Piccinini, G. 2010. Computation in physical systems. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/computation-physicalsystems/> (accessed October 5, 2013).
- Pinker, S. 1997. *How the Mind Works*. New York: W.W. Norton.
- Prinz, J. 2012. Singularity and inevitable doom. *Journal of Consciousness Studies* 19: 77–86.
- Putnam, H. 1960. Minds and machines. In S. Hook, ed., *Dimensions of Mind*. New York: New York University Press, pp. 148–180.
- Robinson, B., and Riley, M. 2010. *Star Trek: U.S.S. Enterprise Haynes Manual*. New York: Pocket Books.
- Ross, D., Ladyman, J., and Kincaid, H., eds. 2013. *Scientific Metaphysics*. Oxford: Oxford University Press.
- Schulze-Makuch, D., and Irwin, L.N. 2006. The prospect of alien life in exotic forms on other worlds. *Naturwissenschaften* 93: 155–172.
- Searle, J. 1980. Minds, brains and programs. *Behavioral and Brain Sciences* 3: 417–457.
- Searle, J. 2005. *Mind: A Brief Introduction*. New York: Oxford University Press.
- Searle, J. 2008. Biological naturalism. In M. Velmans and S. Schneider, eds., *The Blackwell Companion to Consciousness*. Malden, MA: Blackwell, pp. 325–334.
- van Heuveln, B., Dietrich, E., and Oshima, M. 1998. Let's dance! The equivocation in Chalmers' dancing qualia argument. *Minds and Machines* 8: 237–249.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.