© World Scientific Publishing Company

DOI: 10.1142/S1793843012400021



A FRAMEWORK FOR APPROACHES TO TRANSFER OF A MIND'S SUBSTRATE

SIM BAMFORD

Complex Systems Modeling Group Istituto Superiore di Sanità, Rome, Italy simbamford@amail.com

I outline some recent developments in the field of neural prosthesis concerning functional replacement of brain parts. Noting that functional replacement of brain parts could conceivably lead to a form of "mind-substrate transfer" (defined herein), I briefly review other proposed approaches to mind-substrate transfer then I propose a framework in which to place these approaches, classifying them along two axes: top-down versus bottom-up, and on-line versus off-line; I outline a further hypothetical approach suggested by this framework. I argue that underlying technological questions about mind-substrate transfer, there is a fundamental question which concerns our beliefs about continuity of identity.

Keywords: Neural prosthesis; neuroprosthesis; prosthetic; brain-machine interface; robotics; mind uploading; whole brain emulation; mind-substrate transfer; substrate-independent mind; metempsychosis; identity.

1. Introduction

The terms "whole brain emulation", "mind uploading" and "substrate-independent minds" have been used informally in recent years to describe a set of related ideas regarding hypothetical possibilities for transferring or emulating the functioning of a human's brain or "mind" on a synthetic substrate. It is my aim to propose a common framework in which these ideas can be discussed. It should be clear from the outset that these ideas require a great deal of speculation and rest on some difficult-to-define concepts. In order to proceed I will refer to this set of ideas as "Mind-Substrate Transfer" (MST), and I defer a definition of this until Sec. 4.

Some of these perceived possibilities are related to the development of neural prosthetic technology. In Sec. 2, I briefly discuss some recent advances in this technology and speculate about the possibility of using (neural) prostheses to achieve MST. In Sec. 3, I introduce other approaches to MST which have been proposed. In Sec. 4, I try to define MST and I suggest a fundamental question. In Sec. 5, I propose a framework within which it may be useful to group these ideas. In Sec. 6, I propose as a thought experiment an alternative approach which is suggested by this framework. In Sec. 7, I compare and contrast the various proposals.

2. Neural Prosthesis

Prosthesis, the replacement of body parts with artificial substitutes, has a long history. Prostheses have advanced from inert to movable, actively powered and controllable devices. In the previous century came the first attempts to replace parts of the nervous system, or support their functioning, e.g., artificial cardiac pacemakers [Webster, 1995], cochlear implants [House, 1976], and deep-brain stimulators [Kumar et al., 1998]. Current research in this area is widespread, concerning e.g., recording from the motor cortex for control [Taylor et al., 2002], and restoration of vision, whether targeting the retina [Hetling and Baig-Silva, 2004] or visual cortex [Normann et al., 2009]. Existing prosthetic interventions typically act as an input or output of the nervous system, e.g., inputting signals to replace lost sensory functions, or taking outputs from the nervous system and using them to control (prosthetic) limbs. In recent years, some interventions have acted as both input and output, i.e., closed-loop interaction, for example artificial bridges between two cortical regions [Jackson et al., 2006; Marzullo et al., 2010] and a replacement circuit for a spinal central pattern generator [Vogelstein et al., 2008].

In 2011 some significant advances were published regarding closed-loop interaction, in which a signal is taken from the brain, some function of the brain is performed by a synthetic substitute and the output is returned to the brain, thus bypassing and replicating some function internal to the brain. Berger et al. [2011] presented a system that recorded from hippocampal area CA3, performed a preparametrized transformation on the resulting spike stream in real time and applied the result as stimulation to downstream area CA1, allowing a behaving rat to perform a memory-related task, overcoming chemical inactivation introduced between the recording and stimulation sites which would otherwise have prevented the rat from performing the task. Similarly, Prueckl et al. [2011] presented a system that recorded from pontine nucleus and inferior olive, input these recordings in real time to a model of classical conditioning in the cerebellum and applied the output as stimulation to the facial nucleus, bypassing a micro-circuit of the cerebellum and allowing a rat to obtain a conditioned response to a stimulus requiring timing accurate to tens of milliseconds, where in this case it was general anaesthesia rather than localized chemical inactivation that would otherwise have prevented the rat from learning the response. Both projects concern memory formation: In the former, a pathway was restored that enabled the formation of short-term memories, whereas in the latter project the synthetic part of the system itself became the substrate of a learnt response. Both projects have the miniaturization of the synthetic parts of the system for implantability as work in progress. I present myself as a collaborator in the latter project to make it clear that such expertise as I have is in the domain of neural and neuromorphic engineering; the rest of this article, however, is concerned with philosophy, and the reader should take what follows with a pinch of salt.

The ability to replace the functionality of a topologically internal part of the brain, albeit within the limited experimental settings of these projects, raises certain questions: How far can the replacement of brain parts be extended? Would it be possible to replace (progressively or otherwise) all parts of the nervous system (and body), and if so, would this result in a viable (cybernetic or synthetic) organism? To what extent could such an organism resemble the original? Related questions have long been raised in the domain of science fiction. I start by answering these questions from the point of view of the technology currently available; then I set aside that answer and speculate regardlessly about such possibilities.

There are good reasons to think that the current batch of technology could not be extended to this aim. First, the specificity of electrodes is limited; the precision with which the electrical behavior of the nervous system must be recorded in order to extract sufficient information for given tasks to be performed is, in general, an open question; note that Berger et al. [2011] extracted individual spikes from their recordings, sampling only a small proportion of the neurons in the area targeted by electrodes, whereas Prueckl et al. [2011] worked with multi-unit signals, i.e., aggregates of the signals from many neurons in an area, rejecting any information which might consist in individual spike timings. Electrodes damage tissue as they are inserted. Considering then the possibility of the separate replacement of many brain areas, the cumulative damage that would be caused by inserting enough electrodes to sample very sparsely from a majority of the areas of the brain would be substantial. The stability of electrode connections as well as damage they may cause in the long term are long-standing problems [Oh et al., 2002]. Alternative interfacing technologies such as optogenetic stimulation [Zhang et al., 2009] may help to address issues of specificity but will not resolve the problem of how to sample extensively throughout the 3D volume of the brain. If the approach of external or minimally invasive interfacing is followed instead, this can be effective for some tasks [Pistohl et al., 2008, but such an approach is unlikely to support a high enough bandwidth of information for large-scale replacement of many simultaneous functions. The volume, power requirements, and above all the very design of a synthetic substrate which could emulate the functioning of the brain are open questions towards which the nascent field of neuromorphic engineering has taken only small steps so far [Indiveri et al., 2009. Note that the replacement parts in the two aforementioned studies provide equivalent functionality only within a limited experimental setting; it is highly unlikely that no behavioral changes would result from the respective interventions in freely behaving animals. Note also that the choice of brain area in both studies was strategic: the hippocampus and cerebellum are two relatively wellcharacterized areas with clear directional connectivity — the connectivity of the cortex is much more complex. Perhaps the largest problem is that, whilst a synthetic part may have the ability to adapt and take part in the formation of new memories and learnt abilities, it is unclear how existing memories which have been acquired by specific individuals could be transferred to synthetic parts. The study of Berger et al. [2011] made a step in this direction by parametrizing the transformation performed by their synthetic component according to the normal input—output relationship demonstrated by the connection being replaced (for the specific experimental

protocol only). There was therefore a period in which the synthetic connection could be operated alongside the existing biological connection and the performance of the two compared. No attempt to apply parametrizations from one animal to another was reported, and so it is not yet clear if any individual-specific information was captured by the synthetic replacement (in Prueckl et al. [2011], parametrization of the synthetic model was with respect to experimental norms).

I will now speculate on the some related issues. Let us assume that the aim was to replace the whole of the central nervous system, but not the peripheral nervous system or the rest of the body. If this was performed in a single intervention, then the total number of connections that would have to be severed and remade (or otherwise functionally replaced) would perhaps be minimized; the surgical replacement of an entire nervous system, however, even supposing that a functional replacement were available, would be such a complex procedure as to stretch the bounds of conceivablity. At the other end of the spectrum it is possible at least to imagine that the nervous system were replaced progressively, small component by small component, where the components might be identifiable brain regions, or, less credibly, smaller components, even down to individual nerve cells (readers may spot a similarity with Moravec's response to the "Chinese room" argument [Moravec, 1988]). Such a procedure would increase greatly the number of surgical interventions required, with the risk and other inconveniences that each entails, with the number of connections that would be required between biological and synthetic substrate peaking at a much higher level at some point during this process. This is because interconnectivity within the brain is far higher than connections at the periphery of the central nervous system. Such a procedure may provide greater opportunities to assess and correct differences in functioning as the result of each intervention, perhaps along the lines suggested by the work of Berger et al. [2011]. The scenario of replacing one cell at a time, might actually allow the most opportunities to refine performance, were it not for the manifest technological infeasibility.

In summary, the prosthetic replacement of a complex nervous system remains infeasible with known technology. However, the recent landmark results referred to above renew questions that such a prospect raises, and give some additional insight into the problems that would have to be overcome.

3. Approaches Currently Under Consideration

I now turn my attention to other approaches to MST which have been proposed in recent years. Two main approaches which I wish to discuss are: (1) reconstruction from a scan, and (2) reconstruction from behavior, explained as follows (my reporting of these ideas should not be taken as indicating my belief in their possibility or desirability):

(1) Reconstruction from a scan: the brain could be scanned to yield information on structure at high resolution (e.g., capturing the size and approximate composition of each synaptic density). The functioning of the system could then be derived using generic knowledge of neural and synaptic function, and then simulated, where the simulation may be a neural-level simulation or may be at a more abstract functional level (the option of simulation at around the level of detail of individual neurons and synapses also goes under the name of "whole brain emulation" [Sandberg and Bostrom, 2008]). This approach may involve plastination or cryonic suspension, perhaps through a controlled euthanasia, with a view to detailed scanning at a later point; it would also include the non-destructive scanning of a live brain at a sufficient level of resolution were such a thing to become possible. Proponents of the reconstruction-from-a-scan approach include Hayworth [2010].

(2) Reconstruction from behavior, in which information about the behavior of an individual is collected and later used to parametrize a generic substrate in order to in some sense reconstruct the individual. The information collected may be left-over third-party information, e.g., video footage, or it may be information from live data capture that the individual uses. Proponents of the reconstruction-from-behavior approach include Bainbridge [2009] and Rothblatt [2007].

I defer criticism of these ideas to Sec. 7.

4. Defining MST

To recap, there are three approaches that are so far under consideration in this article: reconstruction-from-a-scan and reconstruction-from-behavior from Sec. 3, and gradual-replacement-of-system-parts from Sec. 2. I have delayed defining MST until now in order to do so with the knowledge of these approaches in mind. To discuss MST, traditional questions in AI, such as whether a machine could become conscious or experience human-like qualia, are (prematurely) set aside. The question instead becomes, assuming that a human-like synthetic substrate were available, whether a specific human individual could in some sense come to inhabit that substrate? One can draw a comparison with many expressions of metempsychosis in traditional religions, but instead of the "immortal soul", it is the "mind" that should be transferred, where the definition of "mind" is not without its difficulties, and will not, in fact, be attempted here. Hayworth [2010, p. 7] wrote:

"The debate over mind uploading revolves around a central question, 'What do you consider to be you?' Mind uploading is useless if this personal definition of 'you' is not successfully transferred."

Indeed it is difficult to assess the aforementioned approaches without addressing the question of what this "personal definition of you" consists of. The assumption of Hayworth [2010] is that the information coded in the brain at some microscopic level would be sufficient to reconstruct an individual's unique memories and learnt responses and guide the conscious experiences of the substrate towards conscious

experiences resembling those of the individual: Rothblatt [2007] instead focuses on externally observable behavior as the carrier of information with which the individual could be reconstructed. In both cases, there would be an original human and a resultant human-like synthetic organism (or "synthetic human") physically and temporally separated from each other, yet the information transfer from one to the other is intended to be sufficient that the synthetic human identified with, or were identified with, the original. What is at stake is a belief that the original and synthetic human shared a common identity, where this belief might be held by the synthetic human resulting from the procedure, the original human prior to the procedure, by friends, or by society at large. In this context I consider only beliefs about identity, because I would argue that identity is itself a concept which has no objective reality. It is convenient for us to consider our identity continuous from one day to the next. although during a lifetime there are drastic changes in our appearance, behavior, memories and so on. The assumption of continuous identity attached to a (spatiotemporally continuous) human body is also encoded into and necessary for our society and legal systems. This assumption becomes strained when individuals undergo memory loss or certain forms of psychosis (the case of Phineas Gage [Damasio et al., 1994] is one of many classic examples), and it cannot be unproblematically extended into a hypothetical future in which MST were possible and in which, for example, multiple copies of an individual could be created.

I submit that if MST becomes successfully defined as a field, its fundamental question should be "how can a belief in continuity of identity be supported for a synthetic individual based on a particular human individual?" The only real way to compare and contrast the above approaches, apart from technological feasibility, which is currently stretched in all of them, is to consider what aspects they have that make the adoption of such beliefs more or less likely. If a hypothetical individual resulting from a procedure claimed the identity of the pre-procedure individual, we should ask what it were about the procedure that made this claim more credible than the claim of someone alive today to be a reincarnation of Florence Nightingale? Importantly, although the merits of technological procedures could help to establish the credibility of such claims by demonstrating a certain quality or quantity of information transfer, we should recognize that it would not be the procedure itself that ultimately allowed an individual to transfer substrate, but rather personal and societal acceptance of such transfers.

5. Proposed Framework for Mind Transfer Approaches

In order to consider and compare the aforementioned approaches, I propose a framework in which they are categorized along two axes.

5.1. "On-line" versus "Off-line"

What both reconstruction approaches have in common is that they are both "off-line" approaches; firstly data about the individual is gathered, then the data is used to

reconstruct the individual (perhaps after the death of the individual, though not necessarily in all scenarios). In the gradual-replacement approach by contrast, the synthetic substrate is assembled alongside the pre-existing biological substrate. The biological substrate may be gradually disassembled but they operate in parallel for a time to implement the same individual. Let's call this an "on-line" approach.

5.2. "Bottom-up" versus "Top-down"

What reconstruction-from-a-scan and gradual-replacement-of-system-parts have in common is that they are both "bottom-up": In the case of the scanning approach, the focus is on information about neurons and synapses, or perhaps some other level of information deemed to be the lowest level necessary for reliable simulation or extraction of function. In the gradual-replacement scenario, small parts of the system are replaced and the co-functioning of these individual parts would be expected to sum up to the behavior of the system as a whole. As noted in Sec. 2 the replacement parts may be much larger and many fewer than individual nerve cells, for example a functional replacement for all the hair cells of the cochlear [House, 1976], for all the ganglion cells of the retina [Hetling and Baig-Silva, 2004], or for the CA3-CA1 synapses of the hippocampus [Berger et al., 2011] but this would be because of technological limitations (without wishing to imply that there were a clear technological pathway in any case to achieving replacement of the full system). The reconstruction-from-behavior approach by contrast is "top-down", in that the focus is on behaviors, which arise from the brain, body and environment of the individual as a whole, and from which the lower level details of an implementation of the system could be derived.

I have therefore defined the three approaches along two axes: on-line versus offline, and top-down versus bottom-up.

6. Alternative Approach

Accepting this framework, it becomes apparent that there is fourth category to which no approach is assigned. What would an on-line, top-down approach to MST look like? I will now propose an approach that would fit these criteria, without making any claims as to its feasibility or desirability and also without suggesting that it would be the only approach that could fit into this category.

I follow the aforementioned off-line approaches in assuming the future availability of a generic human-like substrate. To add some more specificity, this hypothetical substrate would be robotic rather than virtual, would have human form, the same degrees of freedom and capabilities of movement as the human body, the same sensory capabilities as a human, and with a synthetic control system modeled closely on the human nervous system, i.e., brain-like functional architecture, with the capacity for volition and independence as a human, but (if possible) without strongly imprinted goals and knowledge, so that it started in an essentially child-like state.

I will assume that such a substrate would generate, or at least be capable of developing, human-like states of subjective consciousness. Following my previous assertions, the question to ask is how to engineer a transfer of information from a particular human to such a substrate so as to best support a belief that the two shared a common identity?

If the robot spent most of its time in the presence of a particular human we might expect it to learn something of the knowledge and behavior of that human, in the same way that a child picks up some behaviors from its mother. The robot would have a different point of view from the human because of inhabiting a different body in a different physical location. It would also have non-symmetrical interactions with the human. These two factors would limit the extent to which the robot and the human would identify with each other — they would consider themselves to be separate individuals and at best the human might feel that the robot had become a helper or a friend, though with no guarantee that the robot would even co-operate towards the same goals as the human.

How could the sense of identity between the two individuals be increased? Here are two possibilities — not an exhaustive list. Firstly some coupling could be engineered between the reward systems of the robot and the human. If the human and the robot experienced the same sense of reward for actions of either body, they might feel constrained to work towards the same goals. Secondly, the body of the robot may be constrained to overlap that of the human. This may be because the robot were an exoskeleton in which the human sat or it may be more subtle, with the robot consisting of a control system connected to sensors and actuators implanted within the human, perhaps even the human's own sensors and actuators. Then the robot and human would be constrained to have the same physical location and sensory point of view and they would have to make the same actions. This would give the two control systems a good rationale for identifying as the same individual. Further strategic (electrical) coupling between the human and robotic brains might be expected to further dissolve the barriers between the human and robotic "minds". These possibilities of course leave many questions unanswered, for example about how control would be shared between the control systems, and how a process of substrate transfer might then conclude. However the general idea should be clear by closely coupling the behavior of a human and a synthetic substrate the two might learn to identify as the same individual.

This approach is on-line in that the biological and synthetic substrates would operate in parallel for a period of time. It is top-down in that it involves a coupling between the high-level behavior of the biological and synthetic substrates, inhabiting the same body in the same environment, in order to identify them as the same individual (although coupling lower-level and internal behavior might increase the success of the procedure). I have proposed this primarily as a thought experiment; I am not aware of anyone seriously suggesting this kind of coupling, especially as a form of MST. Having said this, certain current work on human—robotic interaction

could be seen as antecedent to such an approach. For example, Hiroshi Ishiguro [Miyake et al., 2011] specializes in creating robotic look-alikes of particular people. Then, Adorno et al. [2011] have recently demonstrated the ability of a robot to perform a task partly by taking control of a human's arm.

7. Comparison of Approaches

I now consider some of the general features of the approaches discussed. On-line approaches may have the advantage over off-line approaches that, depending on the details of the procedure, it might be difficult to identify any single point in the procedure at which the subject of the procedure passed from "being biological" to "being synthetic", or a point at which the subject could be considered dead or defunct; this could increase the likelihood of engendering a belief in the continuity of the identity of the individual from before the procedure to after it. On the other hand, by being able to put off the development of most of the necessary technology for an indefinite period, off-line approaches are, paradoxically, closer to being potentially relevant for people alive today. There is also a sense in which allowing a period of time to pass could be advantageous, in that people normally change through time, becoming less similar to their former selves the more time passes. If we meet someone we once knew after a long time has passed, we accept changes in their appearance, losses in the memories we previously shared with them, and changes in their goals and interests that would seem pathological if they were to happen overnight. If, then, a procedure were required to deliver an individual who was at least as similar to the pre-procedure individual as they would likely have been if they had simply continued to live in the intervening period, then the requirement for similarity (in terms of memories, behaviors, appearance, etc.) would be reduced as more time were allowed to pass, effectively easing the constraints on the procedure.

Physical embodiment is fundamental for on-line approaches whereas for off-line approaches it is sometimes considered optional or of secondary importance. In this respect it is worth considering recent arguments on the importance of embodiment for cognition [Clark, 2008; Noe, 2009]. Top-down approaches might be less physically invasive than bottom-up approaches, again depending on the details of the procedure.

With bottom-up approaches, one open question is whether the method of data capture would actually capture all information relevant to the reconstruction of an individual (where the relevance of information would be judged with respect to its importance in engendering a belief in the continuity of the identity of the pre-procedure individual). More fundamentally, there is the question of whether all the components of the system (whether they were simulated neurons and synapses or other functional abstractions) would result in human-like behavior or indeed viable behavior of any kind; low-level functional abstractions may ignore information essential to overall system integrity. Top-down approaches would avoid this problem only if the assumption of the availability of a human-like substrate were allowable.

Given this assumption, one of many questions over the reconstruction-from-behavior approach would be how parametrization of something as behaviorally complex as a human could be achieved; it could turn out that the problem of mapping behavioral observations onto the free parameters of a sufficiently complex human-like substrate were so under-constrained as to render the result worthless (where worth would be judged according to the degree of similarity by some measure that could be achieved between the original and the copy). Alternatively, since behavioral observations from an extensive time period and diverse environmental conditions would be expected to yield a reconstruction at a particular point in time, the problem could instead be unsatisfiable, i.e., it could be that no parametrization would be capable of demonstrating, within a trial environmental context, behavior which was consistent with all the behaviors of an individual during the sampling period. Even if the problem were appropriately constrained it could be computationally intractable.

One criticism of all approaches is that none could be considered untraumatic to an individual; reconstruction-from-a-scan, for example, typically envisions the reincarnation of an individual in a synthetic substrate, possibly in a virtual environment, following their natural death, possibly after an extended defunct period, perhaps with altered legal status and certainly with different life prospects. Such traumatic changes could be expected to induce major behavioral changes in an individual, perhaps reducing any sense of identity that had been engineered between the pre- and post-procedure individuals. Such issues would need to be carefully considered in ethical debates on such procedures.

8. Conclusion

I have outlined recent advances in the field of neural prosthesis and discussed their possible relevance in terms of an approach to MST by gradual replacement of a system's parts. I have then brought into consideration two other approaches in order to think about MST as a field. I have proposed a fundamental question for this field: "how can a belief in continuity of identity be supported for a synthetic individual based on a particular human individual?". I have proposed a framework for different approaches to MST, classifying them along two axes: top-down versus bottom-up, and on-line versus off-line. I have then outlined a further hypothetical approach suggested by this framework, based on close coupling of the behavior of a human and a robot. Finally I have highlighted various comparisons between and problems with different classes of approach.

References

Adorno, B. V., Bó, A. P. L. and Fraisse, P. [2011] "Interactive manipulation between a human and a humanoid: When robots control human arm motion," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Bainbridge, W. [2009] "Religion for a galactic civilization 2.0," http://ieet.org/index.php/IEET/more/bainbridge20090820/.

- Berger, T. W., Hampson, R. E., Song, D., Goonawardena, A., Marmarelis, V. Z. and Deadwyler, S. A. [2011] "A cortical neural prosthesis for restoring and enhancing memory," *Journal of Neural Engineering* 8.
- Clark, A. [2008] Supersizing the Mind: Embodiment, Action, and Cognitive Extension (Oxford University Press).
- Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M. and Damasio, A. R. [1994] "The return of Phineas Gage: Clues about the brain from the skull of a famous patient," Science 264, 1102-1105.
- Hayworth, K. [2010] "Killed by bad philosophy: Why brain preservation followed by mind uploading is a cure for death," accessible at: http://www.brainpreservation.org/documents/killed by bad philosophy.pdf.
- Hetling, J. R. and Baig-Silva, M. S. [2004] "Neural prostheses for vision: Designing a functional interface with retinal neurons," *Neurological Research* 26, 21–34.
- House, W. F. [1976] "Cochlear implants," Ann. Otol. Rhinol. Laryngol. 85 suppl 27(3Pt2); 1–93.
- Indiveri, G., Chicca, E. and Douglas, R. J. [2009] "Artificial cognitive systems: From VLSI networks of spiking neurons to neuromorphic cognition," Cognitive Computation 1, 119–127.
- Jackson, A., Mavoori, J. and Fetz, E. E. [2006] "Long-term motor cortex plasticity induced by an electronic neural implant," Nature 444, 56-60.
- Kumar, R., Lozano, A. M., Kim, Y. J., Hutchison, W. D., Sime, E., Halket, E. and Lang, A. E. [1998] "Double-blind evaluation of subthalamic nucleus deep brain stimulation in advanced Parkinson's disease," Neurology 51, 850-855.
- Marzullo, T. C., Lehmkuhle, M. J., Gage, G. J. and Kipke, D. R. [2010] "Development of closed-loop neural interface technology in a rat model: Combining motor cortex operant conditioning with visual cortex microstimulation," *IEEE Transaction on Neural Systems* and Rehabilitation Engineering 18, 117–126.
- Miyake, N., Ishiguro, H., Dautenhahn, K. and Nomura, T. [2011] "Robots with children: Practices for human-robot symbiosis," in *HRI '11 Proceedings of the 6th International Conference on Human-Robot Interaction*.
- Moravec, H. [1988] Mind Children (Harvard University Press).
- Noe, A. [2009] Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness (Hill and Wang).
- Normann, R. A., Greger, B. A., House, P., Romero, S. F., Pelayo, F. and Fernandez, E. [2009] "Toward the development of a cortically based visual neuroprosthesis," *Journal of Neural Engineering* 6, 1–8.
- Oh, M. Y., Abosch, A., Kim, S. H., Lang, A. E. and Lozano, A. M. [2002] "Long-term hard-ware-related complications of deep brain stimulation," Neurosurgery 50, 1268-1276.
- Pistohl, T., Ball, T., Schulze-Bonhage, A., Aertsen, A. and Mehring, C. [2008] "Prediction of arm movement trajectories from ECoG-recordings in humans," *Journal of Neuroscience Methods* 167, 105–114.
- Prueckl, R., Taub, A. H., Hogri, R., Magal, A., Herreros, I., Bamford, S., Ofek Almog, R., Shacham, Y., Verschure, P. F. M. J., Mintz, M., Scharinger, J., Silmon, A. and Guger, C. [2011] "Behavioral rehabilitation of the eye closure reflex in senescent rats using a real-time biosignal acquisition system," in *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*.
- Rothblatt, M. [2007] "Cyberev project," http://www.cyberev.org.
- Sandberg, A. and Bostrom, N. [2008] "Whole brain emulation: A roadmap, technical report n.2008-3," available at: www.fhi.ox.ac.uk/reports/2008-3.pdf. Technical report, Future of Humanity Institute, Oxford University.

- Taylor, D. M., Tillery, S. I. H. and Schwartz, A. B. [2002] "Direct cortical control of 3D neuroprosthetic devices, Science 296, 1829–1932.
- Vogelstein, R. J., Tenore, F. V. G., Guevremont, L., Etienne-Cummings, R. and Mushahwar, V. K. [2008] "A silicon central pattern generator controls locomotion in vivo," Biomedical Circuits and Systems, IEEE Transactions on 2(3), 212–222.
- Webster, J. G. [1995] Design of Cardiac Pacemakers, (IEEE Press).
- Zhang, J., Laiwalla, F., Kim, J. A., Urabe, H., Van Wagenen, R., Song, Y. K., Connors, B. W. Zhang, F., Deisseroth, K. and Nurmikko, A. V. [2009] "Integrated device for optical stimulation and spatiotemporal electrical recording of neural activity in light-sensitized brain tissue," Journal of Neural Engineering 6, 1–13.