Articles

The Effects of Closed-Loop Medical Devices on the Autonomy and Accountability of Persons and Systems

PHILIPP KELLMEYER, THOMAS COCHRANE, OLIVER MÜLLER,
CHRISTINE MITCHELL, TONIO BALL, JOSEPH J. FINS, and
NIKOLA BILLER-ANDORNO

Abstract: Closed-loop medical devices such as brain-computer interfaces are an emerging and rapidly advancing neurotechnology. The target patients for brain-computer interfaces (BCIs) are often severely paralyzed, and thus particularly vulnerable in terms of personal autonomy, decisionmaking capacity, and agency. Here we analyze the effects of closed-loop medical devices on the autonomy and accountability of both persons (as patients or research participants) and neurotechnological closed-loop medical systems. We show that although BCIs can strengthen patient autonomy by preserving or restoring communicative abilities and/or motor control, closed-loop devices may also create challenges for moral and legal accountability. We advocate the development of a comprehensive ethical and legal framework to address the challenges of emerging closed-loop neurotechnologies like BCIs and stress the centrality of informed consent and refusal as a means to foster accountability. We propose the creation of an international neuroethics task force with members from medical neuroscience, neuroengineering, computer science, medical law, and medical ethics, as well as representatives of patient advocacy groups and the public.

Keywords: autonomy; accountability; neurotechnology; brain-computer interfaces; neuroethics

Introduction

Many emerging neurotechnological medical devices operate on the principle of closed-loop interaction. In one type of closed-loop interaction, an amplifier/computer system records a subject's brain activity in order to create an output that can be monitored by the subject. The subject or the device can then modulate subsequent brain activity, thus closing the loop. In a closed-loop brain-computer interface (BCI), for example, electrodes record neural activity from the scalp or from a neural implant, and a computer algorithm translates the neural activity to control a spelling program or a robotic arm; the subject can then monitor the output and manipulate the program or the robotic prosthesis.[1] In another type of closed-loop interaction, an amplifier/computer system records a subject's brain activity in order to create an output that controls brain stimulation, thus again closing the loop. Systems for treating pharmaco-resistant epilepsy are one prominent example of this type of closed-loop device.[2]

Most BCI systems today record neural activity extracranially with electroencephalography (EEG) or indirectly by measuring hemodynamic changes with near-infrared spectroscopy (NIRS) or functional magnetic resonance imaging (fMRI). In recent years, however, the development of invasive BCIs based on intracranial recording of neural activity with electrocorticography (ECoG) or penetrating needle multielectrode grids (such as the "Utah array") has advanced considerably.

623

Because of technical improvements in signal quality and information transfer rate, this technology will have a variety of clinical applications in coming years.[3,4,5,6]

Based on the available functions that BCI systems provide today—communication through spelling systems and the operation of robotic prostheses—there are different target patient populations for clinical research. Patients with neurodegenerative diseases such as amyotrophic lateral sclerosis (ALS) may benefit from a BCI system. Neurodegeneration in ALS in later stages results in partial or complete loss of all voluntary motor control—the locked-in state, which limits or prevents communication by natural means. A BCI device may thus enable or preserve communication in these patients.[7] For stroke victims with chronic aphasia, BCI could potentially improve communication, and for severely paralyzed stroke patients with largely intact cognitive abilities, it could restore motor function.[8] In patients in a minimally conscious state (MCS), BCI devices may also enable communication with the outside world.[9,10]

In what follows, we examine the effects of these novel closed-loop neuro-technological devices on the autonomy and accountability of persons and systems. Drawing on these observations, we propose the development of an international consensus process on the ethical and legal ramifications of this technology.

The Autonomy of Persons and Systems

Personal Autonomy

When we discuss *personal autonomy* in medical ethics[11] we implicitly invoke the concept of moral agents and usually assume that only persons count as agents. However, we intend to compare personal autonomy with the autonomy of computerized systems, and we must therefore specify precisely what type of agents can and should be described as autonomous. For our purposes, in order to be considered autonomous, an agent must (1) interact with objects or other agents, (2) possess reliable heuristic and decisionmaking capacity, (3) be the de facto originator of particular actions, and (4) act in accord with her (its) beliefs.

Although these criteria are usually considered *necessary* for autonomy, there are some individuals to whom we would ascribe some degree of autonomy or agency despite their being unable to be the de facto originator of particular actions. At the level of informed consent and refusal, their responses may be reactive at a lower level of assent or dissent, in large part because they can neither initiate nor fully explain their actions. This is the case for patients who are in a minimally conscious state who have been assisted by other neuroprosthetic devices such as deep brain stimulation (DBS), neuroimaging, or an EEG-based BCI system.

In order to accommodate subjects who are unable to originate actions, we propose the following as a working definition of "personal autonomy": personal autonomy arises from the subject's experience of congruence of motive and action, which gives rise to the feeling of individual agency. Ideally, these actions should be in agreement with more invariant or long-term aspects of the agent's character, dispositions, and moral values. Of note, this level of functionality requires the retention of sufficient memory and personal identity.

Effects of Closed-Loop Medical Devices

System Autonomy

BCIs that operate automatically in response to the environment that they monitor can be considered to have a limited sort of autonomy. We call this *system autonomy*. Establishing criteria for what constitutes truly autonomous system behavior, however, is not straightforward. A BCI could satisfy most of our previously discussed criteria for agency, without us being very tempted to consider it an agent. It is easy to imagine a machine that could satisfy the first three: interaction with objects or agents, possession of heuristics and decisionmaking capacity, and capacity to originate actions. The fourth criterion, acting in accord with beliefs, is more problematic. One could perhaps plausibly describe a system as having beliefs to the extent that it is *designed* to have beliefs. For example, a thermostat designed to turn on the air conditioning when the temperature rises above a certain point might be said to "believe" that it is too hot. One might respond that such a simple and fully predictable behavior seems to stretch the definition of "belief" too far, and we agree—but what about systems behaviors that are much more complex, involving machine learning and adaptation, and that are not fully predictable? In this setting, it becomes much more plausible to describe systems as having beliefs—and to the extent that this is true, we must ascribe some (limited) degree of agency to such a system. One important feature of autonomous system behavior, then, is that it is not fully predictable by the system engineer and not fully controllable by human agents interacting with the system.

Distinctions between Personal and System Autonomy

In order to clarify the distinction between system autonomy and personal autonomy when it comes to BCIs, we first examine the principle of closed-loop interactions and then discuss possible criteria for autonomous system behavior. As an example, we use closed-loop devices for epileptic seizure control and for paralyzed patients, but our discussion also applies to other closed-loop medical devices.

   In closed-loop systems for epileptic seizure control, implanted electrodes continuously monitor neural activity and may deliver electric stimulation to interrupt or prevent the seizure.[12,13] Autonomous system behavior in this case would entail a measure of decisional authority built into the algorithm of the device. On the one hand, an engineer could program the algorithms in such a way as to fully predetermine the actions taken by the system when it detects epileptic electrocorticographic patterns. Such a programmed system would have no autonomy in terms of systematic decisionmaking or agency. On the other hand, if the system uses an adaptive machine learning algorithm, which can generate its own inferences from the neural data and stimulate ad lib, the system may have some degree of autonomy in the sense that it "chose" its response without being explicitly directed to do so by the engineer. Of course, one could argue that the decision on how much decisional capacity or agency to build into the system is ultimately made by the programmer, and this *meta-autonomy* of design thus leaves the burden of accountability with the engineer, regardless of the degree of autonomy of the system.[14] This raises a deeper set of questions about whether such a system can ultimately be fully autonomous, but this will have no implications for our thesis here.

How complex must a BCI system be before we begin to ascribe partial or full autonomy to it? Few if any would ascribe autonomy to a BCI that operated like a thermostat—delivering a stimulus whenever cortical activity exceeds a certain threshold. But BCIs are being developed with adaptive algorithms based on machine learning, which makes the BCI's decisionmaking much more sophisticated and complex. If the engineer cannot anticipate or predict the system's behavior with absolute precision, then we submit that such a system should be ascribed at least some degree of autonomy. Engineers sometimes describe this difficulty of predicting system behavior as the *decisional space* being occupied by the system. When the full range of system behavior is not fully predictable anymore, and the system occupies some of the decisional space, the system represents at least a weak instance of artificial intelligence—it is an "intelligent system" and can be considered an intelligent entity.[15] Although the presence of intelligence does not *necessarily* imply the presence of autonomy, we submit that this type of adaptive decisionmaking capacity implies both intelligence and autonomy—it certainly seems to satisfy our criteria for moral agency.

Because we argue that such learning systems should be considered to have at least a version of autonomy, we must ask ourselves whether the autonomy of the human part of the BCI system should be extended to include the algorithmic part. One could consider such a human-machine decisionmaking system as either a single autonomous entity or two separate entities, each with perhaps only some privileges and duties of autonomy. Regardless of whether there are one or two entities, as the autonomy of the electronic component increases, the autonomy of the human subject may diminish—a subject of concern, in our opinion. On the other hand, we later discuss how *shared autonomy* between a paralyzed patient and a BCI system may also strengthen the autonomy of the patient by enabling decisionmaking.

Of course, system autonomy is only comparable to personal autonomy on the level of agency. Thus, we can describe different forms of "shared agency" in human-machine interaction, but on a level of autonomy as a capacity to reflect actions regarding moral implications or regarding the accordance of an action to a person's goals or way of life, we usually use the concept of autonomy in a broader sense.

We next examine how keeping the human subject in the loop may make a difference in terms of preserving the subject's autonomy and agency.

Keeping the Subject in or out of the Loop?

In the example of closed-loop medical devices for epileptic seizure control, we can either give the subject some feedback and control over the system—keeping her in the loop—or design the system to learn, adapt, and act without any input from the subject—keeping her out of the loop. One could leave the subject out of the loop by having the system monitor for seizure risk and respond according to its algorithms, but without notifying the patient. This has the advantage of convenience—the machine operates independently to reduce the risks of seizure—but the disadvantage of leaving the patient out of decisions regarding whether to act in response to an increase in seizure risk. Alternatively one could keep the subject in the loop, for example, by showing her a risk-indicating "traffic light" signaling no risk (green) to high risk (red) of an impending seizure.

Effects of Closed-Loop Medical Devices

If the subject is in the loop, she retains some autonomy over decisionmaking. For instance, if the traffic light shows yellow or red, the patient might interrupt activities like riding a bike or standing on a ladder. This has the advantage of giving the patient as much knowledge about the state of her system as possible, thus enhancing her autonomy with respect to her seizures, but has the disadvantage of requiring the patient to act in a way that maximizes safety. Criteria for distinguishing when human individuals should be able to override medical devices and when medical devices should intervene to override human autonomy remain to be developed.

We do not know yet which type of system is preferred by subjects, because we do not have sufficient empirical data on acceptance of this type of predictive and interventionist neurotechnology. However, the acceptance of electronic bodily devices for measuring, recording, and analyzing biometric data to modify behavior suggests a willingness of some individuals to share responsibility for health outcomes with their medical devices. Structured interviews with epilepsy patients show that the acceptance of a closed-loop system for seizure detection and intervention depends on whether the system causes palpable disruptions of daily routines.[16] A subject whose routines were disrupted by warning signals about seizure risk, therefore, might prefer to be kept out of the loop, ceding (or perhaps delegating) some of her autonomy to the BCI.

For BCIs, the question of whether and to what degree we keep the subject in the loop may have a profound impact on her experience of autonomy and personal identity too. For example, if the BCI uses the neural data to operate a robotic arm or an exoskeleton, the effects on embodied perception, body schemata, and the feeling of ownership are difficult to predict.[17,18] This leads to interesting anthropological and neurophilosophical questions: How does the incorporation of closed-loop neurotechnological devices, such as an intracranial BCI system, alter a subject's perception and sense of self and personal identity? How does the sense of agency and ownership change in times of neurotechnology? One could consider the device part of one's own bodily experience—or a person could say: "Somehow it's me, but I don't feel that I am the author of the action in the full sense." Another question concerns the boundary between a system merely *assisting* versus *enhancing* (or *augmenting*) a subject's behavioral range or power—neuroenhancement. These issues are beyond the scope of this article, but we go on to propose a structured method of considering them.

Strengthening or Undermining Personal Autonomy through System Autonomy?

Establishing a BCI system for spelling may enable a severely paralyzed patient to communicate, thus permitting participation in decisions and thereby preserving personal autonomy. Take the case of a person with ALS in a completely locked-in state with no remaining oculomotor movement. For such a patient, communicating via a BCI spelling system could restore some measure of autonomy through participation in decisions regarding her own well-being and care. In such a scenario, this would still be the case if some of the decisionmaking as to which letters the system selects relied on a machine learning algorithm. Partially relegating decisionmaking capacity, like spelling choices, to the algorithmic part of the BCI system may be necessary from an engineering viewpoint, to improve the decoding accuracy and spelling efficiency of the BCI. As in the case of the epilepsy patient

in which a visual feedback signaling keeps her in the decisionmaking loop by indicating the seizure risk, giving a paralyzed patient who is using a BCI final control over authorizing the spelling result could mitigate the potential for unintended system behavior.[19]

Importantly, autonomy is not a dichotomous state in our opinion but, rather, a quality that varies in proportion to the behavioral adaptiveness of the system.[20] Researchers and ethicists discuss these considerations under the rubrics "roboethics" and "responsible robotics," looking at the benefits, risks, and limits of semi- or fully autonomous algorithmic control systems for intelligent machines.[21,22,23,24] At present, in our view, the degree of system autonomy of the available BCI solutions strengthens rather than undermines personal autonomy by enabling otherwise severely impaired patients to regain control of communication and movement to some degree.

As for the impact of this technology on the quality of life of patients, few in-depth case studies have explored the social impact of long-term use of BCI systems for paralyzed patients. One study, for instance, has shown that they may improve quality of life, for example, by allowing patients to participate in creative activities like painting;[25] this highlights the importance of developing neurotechnological devices, like BCIs, from a user-centered perspective that takes the needs of patients and their families as well as caregivers into account.[26,27]

As for the limitations of our current knowledge and possible negative effects of increasing system autonomy, the following issues warrant particular attention from our point of view. Most clinical trials on implantable closed-loop devices for communication, mobility, or seizure control have thus far involved only very few subjects. Limited empirical data are available on the attitudes of patients and the effects on the experience of autonomy for patients who used BCIs and/or received closed-loop implants.[28,29] Future clinical trials with neural implants should therefore include gathering of such qualitative and quantitative data with structured, in-depth interviews as well as focus group discussions. Collecting these empirical data is also important for examining the possible negative effects of closed-loop systems on behavior, dispositions, and mental health. The need for such systematic caution is warranted, in part, by the history of adverse neurobehavioral effects, such as increased suicidal ideation, pathological gambling, or changes in personal identity in patients who received DBS.[30,31]

Furthermore, restoring a communication channel for exercising autonomous decisionmaking in paralyzed patients with a BCI may also lead to a "burden of normality."[32] This term describes an uncommon but important paradoxical negative effect in, for example, patients with Parkinson's disease (PD) who have a DBS implant.[33] These patients occasionally suffer from disruption of social relationships following the implantation of the DBS device. The slowly progressing character of PD leads to a fortification of asymmetrical social power relations over time in which the patient becomes the disempowered recipient of care and the partner becomes the empowered caregiver and, often, the patient's legal guardian. When DBS reverses motor disability (sometimes suddenly and dramatically), occasionally a drastic relational adjustment (even role reversal) occurs. This, in turn, can lead to marital problems, even resulting in divorce in some documented cases.[34,35] This analogy from PD may not fully map onto the experience of a severely paralyzed patient, because no current BCI system can undo the profound disability of a locked-in ALS patient. Nevertheless, the restoration of a previously lost ability—in

this case communication and the ability to assert one's autonomy—may exert similar stress on a relationship.

Taking these positive and negative effects into account, it is important to develop guidelines and an ethical framework to assess potential future closed-loop systems that may disrupt personal autonomy, particularly with respect to decisionmaking and agency. We next take up the question of the moral and legal accountability of autonomous system behavior.

The Accountability of Persons and Systems

The intrinsic value of personal autonomy, including the feeling of agency and the freedom of decisionmaking, also entails a number of moral and legal obligations and responsibilities. Put another way, we must consider how to approach moral and legal *accountability* when human subjects and (semi)autonomous systems work in concert.

Taking macrotechnological trends as an indicator, we seem to be increasingly willing to relegate decisionmaking to autonomous systems—think self-driving cars, computer-assisted flying, and diagnostic software. This is especially the case if these systems reliably perform better than human agents do at certain tasks. Imagine the number of road traffic accidents that could be avoided if we had a high-performance autonomous vehicle with an error rate well below that of human drivers.

Nevertheless, in all these cases of autonomously behaving complex systems, it remains unclear who is to be held morally and legally accountable in the case of unintended catastrophic system failure. Take the case of a grave error in a system's algorithm that results in unanticipated and undesired outcomes like injury or death. Who (or what) is to be held accountable in such a case? What if a medical device fails to predict and interrupt an epileptic seizure, which results in the subject being in an unsafe environment, leading to her injury and/or the injuries of others? Who will be taken to court—the subject, the programmer, or the device company? Or the regulatory body that authorized the device? Is responsibility shared so diffusely in the context of autonomously behaving complex systems that it becomes impossible to hold any one agent accountable? These serious questions suggest the need to retain human oversight at key decision points in order to maintain moral oversight of consequential actions and decisions, lest they become amoral.

A current example of this problem of accountability is the implantable cardioverter/defibrillator (ICD), which monitors cardiac rhythm and may administer electric shocks in case of a dangerous arrhythmia.[36] In observational studies, up to 13 percent of patients wearing ICDs received unnecessary electric shocks based on false positive classification by the algorithm.[37] In this case, the patient wearing the device is kept out of the loop and has no control over whether or not an electric shock should be administered; the patient thus cannot be held accountable for any adverse consequences of a false shock. For these devices, the question of liability and accountability is largely unresolved.[38,39]

For the epilepsy patient with an intracranial device for seizure control, the question of whether the subject is kept in or out of the loop also becomes relevant for our assessment of accountability. This will also hinge on basic constructs related to informed and ongoing consent to the intervention's use. If a subject is unaware of

a moment-to-moment increase in her seizure risk, one cannot reasonably hold her accountable for the consequences of a resultant seizure. If an unintended system failure occurs in that scenario—for example, if the device misses an evolving seizure, which puts the subject in danger or causes damage to a third party as a result—we are again faced with the question of who should be held accountable. The out-of-the-loop subject has some responsibility for consenting to the consequences of the implanted system but is not accountable for the consequences of any particular seizure. If, on the other hand, the subject remains in the loop (e.g., via a visual feedback system), the subject's failure to modify her behavior in accordance with the indicated level of risk may indeed result in moral and legal accountability. The in-the-loop subject, therefore, has an ongoing interaction with the technology in a way that the out-of-the-loop subject does not. The in-the-loop subject may have increased autonomy when compared to the out-of-the-loop subject, but she also has increased responsibility. Furthermore, this interaction may also strengthen ongoing consent for such a bodily intervention. It could be argued that if a patient who could give consent is not provided information necessary to exercise that right, the moral warrant for continued participation is lost, because it now becomes an unconsented-to intervention, and therefore there is a breach of both information exchange and voluntariness, both key elements of the Nuremberg Code.[40]

When it comes to BCIs for restoring mobility or communication using current technology, the potential benefits outweigh the risks of unintended system behavior, because patients stand to gain major benefits, and current technologies cannot pose much risk if they malfunction. For example, if a patient in the locked-in state has a BCI that malfunctions, the likely results (like spilling a drink or spelling errors) pose relatively little risk to her and virtually no risk to others at present. This may change in the future, however, when BCI technology may become more sophisticated (e.g., BCI control of large-scale and highly mobile robotic devices) and if its applications move from the population of severely impaired patients to a wider variety of users.

Next, we make some suggestions regarding an international ethical and legal framework for regulating the development and application of this technology.

Regulating Closed-Loop Neurotechnological Devices: A Participatory Approach

Whether the design principles for closed-loop neurotechnological devices should only *guide* the decisionmaking policy of researchers, policymakers, and research ethics committees or whether they should have a binding legal character is certainly debatable. This leads us to the important question of how to organize and channel the political process for regulating these emerging neurotechnologies. We propose a political decisionmaking process that is guided by working models of Jürgen Habermas' theory of "deliberative democracy."[41,42] In this bottom-up model of political legitimatization, extensive public discourse, dialogue, and deliberation involve all stakeholders in a regulatory or legislative issue. This procedure avoids top-down "expertocratic" imposition of regulation that is developed without adequately consulting the needs of the most vulnerable stakeholders in a particular issue—in this case the severely neurologically impaired patients and their families.[43] Similar initiatives—for example, the recent International Summit

on Human Gene Editing in 2015—have been formed for discussing the impact of genetic engineering.[44] In the case of emerging neurotechnology, particularly invasive closed-loop neural implants, this deliberative process would include representatives of patient advocacy groups, the public, medical professionals, neuroengineers, computational scientists, neuroethicists, and legal professionals. One possibility to organize and channel such a deliberative process would be an independent task force that regularly brings together these stakeholders locally and in international meetings to work out an ethical and legal framework for regulating current and future neurotechnology. This task force should initiate specially tailored focus groups, including the aforementioned parties, in order to identify the challenges and the crucial ethical and societal aspects. This process could be complemented by efforts in science communication and public outreach by researchers in the field in order to involve the public in learning about and discussing current research.[45]

Conclusions

We have seen that closed-loop neurotechnological devices may preserve, restore, or strengthen personal autonomy, particularly in the case of BCIs for severely paralyzed neurological patients. We have identified an accountability gap in scenarios in which decisionmaking capacity and agency are relegated to an autonomous system. There may be value in keeping a subject in the loop for preserving and strengthening the autonomy of decisionmaking and agency, though this involves a tradeoff in responsibility and accountability—the more in the loop a subject is, the more responsible she is for the outcomes. We describe a need to consider the possibility of shared autonomy between subjects and systems and to discuss whether traditional mandates for informed consent (or refusal) and agency are adequate ethically and legally in such scenarios. We have also identified a paucity of systematic studies on the effects of neurotechnological devices on subjects' experience and feelings of personal autonomy. As a subject's first-person experience lies at the heart of autonomy, these empirical data may increase our understanding of the issue, and we thus strongly encourage systematic studies on this topic.[46,47] In this context, it is also important to gather empirical data on the opinions, attitudes, values, and beliefs of the stakeholders in particular neurotechnologies, as exemplified by the Asilomar Survey, a survey that analyzed such data by polling 145 researchers in the field of BCI technology.[48]

Finally, we have proposed the formation of an international task force based on a participatory model so that all stakeholders have the opportunity to deliberate the ethical and regulatory challenges of emerging neurotechnologies—like closed-loop devices for brain-computer interfacing.

Notes

1. Fairclough S. A closed-loop perspective on symbiotic human-computer interaction. In: Blankertz B, Jacucci G, Gamberini L, Spagnolli A, Freeman J, eds. *Symbiotic Interaction*. Berlin, Germany: Springer International; 2015:57–67.
2. Wilmshurst JM, Berg AT, Lagae L, Newton CR, Cross JH. The challenges and innovations for therapy in children with epilepsy. *Nature Reviews Neurology* 2014;10:249–60.
3. Bensmaia SJ, Miller LE. Restoring sensorimotor function through intracortical interfaces: Progress and looming challenges. *Nature Reviews Neuroscience* 2014;15:313–25.

4. Baranauskas G. What limits the performance of current invasive brain machine interfaces? *Frontiers in Systems Neuroscience* 2014;8:68.

5. Hochberg LR, Bacher D, Jarosiewicz B, Masse NY, Simeral JD, Vogel J, et al. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 2012;485:372–5.

6. Schneider MJ, Fins J, Wolpaw JR. Ethical issues in BCI research. In: Wolpaw JR, Wolpaw EW, eds. *Brain-Computer Interfaces: Principles and Practice*. Oxford: Oxford University Press; 2012:373–86. doi:10.1093/acprof:oso/9780195388855.001.0001.

7. Birbaumer N, Ghanayim N, Hinterberger T, Iversen I, Kotchoubey B, Kübler A, et al. A spelling device for the paralysed. *Nature* 1999;398:297–8.

8. Silvoni S, Ramos-Murguialday A, Cavinato M, Volpato C, Cisotto G, Turolla A, et al. Brain-computer interface in stroke: A review of progress. *Clinical EEG and Neuroscience* 2011;42:245–52.

9. Chatelle C, Chennu S, Noirhomme Q, Cruse D, Owen AM, Laureys S. Brain–computer interfacing in disorders of consciousness. *Brain Injury* 2012;26:1510–22.

10. Naci L, Monti MM, Cruse D, Kübler A, Sorger B, Goebel R, et al. Brain–computer interfaces for communication with nonresponsive patients. *Annals of Neurology* 2012;72:312–23.

11. Beauchamp TL, Childress JF. *Principles of Biomedical Ethics*. Oxford: Oxford University Press; 2001.

12. Mormann F, Andrzejak RG, Elger CE, Lehnertz K. Seizure prediction: The long and winding road. *Brain* 2007;130:314–33.

13. Nagaraj V, Lee ST, Krook-Magnuson E, Soltesz I, Benquet P, Irazoqui PP, et al. Future of seizure prediction and intervention: Closing the loop. *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society* 2015;32:194–206.

14. Falcone R, Castelfranchi C. The human in the loop of a delegated agent: The theory of adjustable social autonomy. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* 2001;31:406–18.

15. Russell S, Norvig P. *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall; 2013.

16. Hoppe C, Feldmann M, Blachut B, Surges R, Elger CE, Helmstaedter C. Novel techniques for automated seizure registration: Patients' wants and needs. *Epilepsy & Behavior* 2015;52, Part A:1–7.

17. Holmes NP, Snijders HJ, Spence C. Reaching with alien limbs: Visual exposure to prosthetic hands in a mirror biases proprioception without accompanying illusions of ownership. *Perception & Psychophysics* 2006;68:685–701.

18. Hoffmann M, Marques HG, Hernandez Arieta A, Sumioka H, Lungarella M, Pfeifer R. Body schema in robotics: A review. *IEEE Transactions on Autonomous Mental Development* 2010;2:304–24.

19. Given the presence of predictive algorithms in everyday devices like "smart" phones, most of us have experienced the unease (and sometimes the time-saving pleasure) that results from an algorithm inferring one's intention by autocorrecting text messages.

20. Barber KS, Goel A, Martin CE. Dynamic adaptive autonomy in multi-agent systems. *Journal of Experimental & Theoretical Artificial Intelligence* 2000;12:129–47.

21. Veruggio G, Operto F. Roboethics: Social and ethical implications of robotics. In: Siciliano B, Oussama K, eds. *Springer Handbook of Robotics*. Berlin: Springer-Verlag; 2008:1499–524.

22. Sharkey N. The ethical frontiers of robotics. *Science* 2008;322:1800–1.

23. Bostrom N. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press; 2014.

24. Murphy RR, Woods DD. Beyond Asimov: The three laws of responsible robotics. *IEEE Intelligent Systems* 2009;24:14–20.

25. Holz EM, Botrel L, Kaufmann T, Kübler A. Long-term independent brain-computer interface home use improves quality of life of a patient in the locked-in state: A case study. *Archives of Physical Medicine and Rehabilitation* 2015;96:S16–S26.

26. Liberati G, Pizzimenti A, Simione L, Riccio A, Schettini F, Inghilleri M, et al. Developing brain-computer interfaces from a user-centered perspective: Assessing the needs of persons with amyotrophic lateral sclerosis, caregivers, and professionals. *Applied Ergonomics* 2015;50:139–46.

27. Fins JJ. *Rights Come to Mind: Brain Injury, Ethics, and the Struggle for Consciousness*. Cambridge: Cambridge University Press; 2015.

28. Gilbert F. A threat to autonomy? The intrusion of predictive brain implants. *American Journal of Bioethics Neuroscience* 2015;6:4–11.

29. Lahr J, Schwartz C, Heimbach B, Aertsen A, Rickert J, Ball T. Invasive brain–machine interfaces: A survey of paralyzed patients' attitudes, knowledge and methods of information retrieval. *Journal of Neural Engineering* 2015;12:043001.

30. Mathews DJH. Deep brain stimulation, personal identity and policy. *International Review of Psychiatry* 2011;23:486–92.

31. Smeding HMM, Goudriaan AE, Foncke EMJ, Schuurman PR, Speelman JD, Schmand B. Pathological gambling after bilateral subthalamic nucleus stimulation in Parkinson disease. *Journal of Neurology, Neurosurgery & Psychiatry* 2007;78:517–19.

32. Gilbert F. The burden of normality: From "chronically ill" to "symptom free." New ethical challenges for deep brain stimulation postoperative treatment. *Journal of Medical Ethics* 2012;38:408–12.

33. As another somewhat related example, distress can occur when a patient at genetic risk for Huntington's disease is told that he or she is not carrying the mutation.

34. Schüpbach M, Gargiulo M, Welter ML, Mallet L, Béhar C, Houeto JL, et al. Neurosurgery in Parkinson disease: A distressed mind in a repaired body? *Neurology* 2006;66:1811–16.

35. Clausen J. Ethical brain stimulation—neuroethics of deep brain stimulation in research and clinical practice. *European Journal of Neuroscience* 2010;32:1152–62.

36. Kelley AS, Reid MC, Miller DH, Fins JJ, Lachs MS. Implantable cardioverter-defibrillator deactivation at the end of life: A physician survey. *American Heart Journal* 2009;157:702–8.

37. McLeod CJ, Boersma L, Okamura H, Friedman PA. The subcutaneous implantable cardioverter defibrillator: State-of-the-art review. *European Heart Journal* 2015;ehv507. [Epub ahead of print.] doi:10.1093/eurheartj/ehv507.

38. Maisel WHM. Safety issues involving medical devices: Implications of recent implantable cardioverter-defibrillator malfunctions. [Editorial]. *JAMA* 2005;294:955–8.

39. Vinck I, Laet CD, Stroobandt S, Brabandt HV. Legal and organizational aspects of remote cardiac monitoring: The example of implantable cardioverter defibrillators. *EP Europace* 2012;14:1230–5.

40. Shuster E. Fifty years later: The significance of the Nuremberg Code. *New England Journal of Medicine* 1997;337:1436–40.

41. Habermas J. Three normative models of democracy. *Constellations* 1994;1:1–10.

42. Chambers S. Deliberative democratic theory. *Annual Review of Political Science* 2003;6:307–26.

43. Landwehr C. Democratic and technocratic policy deliberation. *Critical Policy Studies* 2010;3:434–9.

44. Travis J. Germline editing dominates DNA summit. *Science* 2015;350:1299–300.

45. In Freiburg (the academic environment of authors PK, OM, and TB), for example, the interdisciplinary research consortium BrainLinks-BrainTools—which is sponsored as a "cluster of excellence" (*Exzellenzcluster*, in German) by the German Research Foundation (DFG)—includes science communication and public outreach as an integral part of its work. See BrainLinks-BrainTools. Science communication and public outreach; available at http://www.brainlinks-braintools.uni-freiburg.de/research/projects/reaching-out-science-communication-and-public-outreach/ (last accessed 2 Mar 2016).

46. Joffe S, Manocchia M, Weeks JC, Cleary PD. What do patients value in their hospital care? An empirical perspective on autonomy centred bioethics. *Journal of Medical Ethics* 2003;29:103–8.

47. Musschenga AW. Empirical ethics, context-sensitivity, and contextualism. *Journal of Medicine and Philosophy* 2005;30:467–90.

48. Nijboer F, Clausen J, Allison BZ, Haselager P. The Asilomar Survey: Stakeholders' opinions on ethical issues related to brain-computer interfacing. *Neuroethics* 2011;6:541–78.