

When Thinking is Doing: Responsibility for BCI-Mediated Action

Stephen Rainey , Hannah Maslen, and Julian Savulescu

The Oxford Uehiro Centre for Practical Ethics, University of Oxford

ABSTRACT

Technologies controlled directly by the brain are being developed, evolving based on insights gained from neuroscience, and rehabilitative medicine. Besides neuro-controlled prosthetics aimed at restoring function lost somehow, technologies controlled via brain-computer interfaces (BCIs) may also extend a user's horizon of action, freed from the need for bodily movement. Whilst BCI-mediated action ought to be, on the whole, treated as conventional action, law and policy ought to be amended to accommodate BCI action by broadening the definition of action as "willed bodily movement". Moreover, there are some dimensions of BCI mediated action that are significantly different to conventional cases. These relate to control. Specifically, to limits in both controllability of BCIs via neural states, and in foreseeability of outcomes from such actions. In some specific type of case, BCI-mediated action may be due to different ethical evaluation from conventional action.

KEYWORDS

neurotechnology; responsibility; control; neuroethics; BCI; disability; responsibility

INTRODUCTION

When assessing agents' moral responsibility for instances of action (and associated outcomes) we examine i) the degree of control they had over the relevant causal processes and ii) the degree of foreseeability of the morally relevant outcomes produced by instigating or inhibiting (or failing to instigate or inhibit) the relevant causal processes. We also attend to iii) the actor's intentions in instigating or inhibiting causal processes (Glannon 2016).

Consider the following example:

Standard Sam



During a heated argument, Sam punches Alex in the head. Immediately after the punch is issued, Sam claims that he didn't intend to punch Alex and that we shouldn't hold him responsible.

Despite Sam's claim, we would nevertheless hold Sam morally responsible for the harm caused by the punch, since Sam had control over his arm and should have foreseen that punching Alex would cause the harm that it did. Sam's moral responsibility would only be reduced if Sam lacked full control over his arm (perhaps he suffered from a momentary muscular spasm), or if he was not expected to foresee the harm that propelling his arm would cause (perhaps Alex had snuck

up beside Sam, who punched his arm in the air to stretch it out, unfortunately making contact with Alex's head). For most standard instances of punching, the agent's claim that they did not intend it would at most be taken to suggest that it was spontaneous rather than pre-meditated action, but would not mitigate their moral responsibility. The high degree of control that agents usually have over their arms makes most claims of involuntary action implausible.

This high degree of control over arm movement (especially highly targeted arm movement) makes attribution of moral responsibility fairly straightforward in standard cases. However, if an agent were to act via a neuroprosthetic device—a device controlled directly by the user's brain activity—assessing the agent's moral responsibility for that act would be more challenging. These challenges are generated by the somewhat reduced control the user has over the causal processes leading to an effect in the world, combined with the reduced foreseeability of the precise nature of the effect that eventuates.

In terms of control over the action, the concepts of executive and implementational dimensions of control are of relevance (Shepherd 2015). Executive control is a high-level dimension to do with general goal-setting. It can remain fairly unchanged by elements of context,

CONTACT Stephen Rainey  stephen.rainey@philosophy.ox.ac.uk  The Oxford Uehiro Centre for Practical Ethics, University of Oxford, Oxford, UK.

© 2020 The author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

whereas the implementational dimension must be more responsive to specific contextual constraints. A goal of “getting the cup” can remain invariant despite any environmental factors that will prompt implementational differences, such as altering reach, or grasp. What remains hard to discern is what indicates the relationship between executive and implementational dimensions of control. What, if anything, can be detected to signal that now implementation of a goal ought to begin. A plan can be held for some time, clearly held, with full intention to pursue it, and yet not be implemented. The point at which implementation does occur is not clearly signposted: I reach for the cup now, having had a general idea to so reach for some time, and nothing overtly has changed.

In terms of technology, the lack of a full understanding of this executive to implementational shift affects the extent to which a neuroprosthetic device can be developed that affords a user the highest levels of control. To achieve a high level of control over a prosthetic hand, for example, might require either (1) high decoding performance of motor tuning parameters (e.g. reach, grasp) from neural activity, or (2) the decoding of specific action goals (e.g. “get the cup”) to be implemented by semi-autonomous, or “smart” neuroprosthetic devices; or (3) a combination of (1) and (2). But what remains hard to make out is the implementational “trigger”, as executive goal-setting shades into implementation. This is a conceptual consideration that stands as an umbrella concern related to designs of neuro controlled devices. We will argue that there are knock-on effects for the use of specific devices. In terms of device use, there will be an upper limit on exactly how much control a user can have. This will be the case even if they are maximally competent with their device. This, and associated issues will be explored first with reference to brain-computer interfaces (BCIs), and a few illustrative examples.

Brain-computer interfaces can be used to replace or recover abilities lost through physical damage or bodily disease, as well as control a variety of non-therapeutic devices. Through recording and decoding brain signals, they can control assistive technologies like prosthetic limbs and wheelchairs, speech prostheses, software programs, or other devices like drones. BCIs offer the chance to interact with the external environment without moving the body at all, merely by realizing some neural activity. This has sometimes been referred to as “thought control” or “mind control” of devices. Despite this being a misleading way of putting things, because the realization of neural

activity is associated with particular thought patterns, learned via training, in some sense thinking is doing (Solon 2017; Whyte 2018; Revell 2018).

Unlike the case of Sam, claims regarding involuntariness or unforeseeability with respect to an action mediated by a BCI may not be dismissed so quickly. Consider the following:

BCI Bob

Bob uses a robotic arm that he controls by generating patterns of neural activity in his motor cortex. In order to control the arm, he must vividly imagine the movement he wants the arm to make. The device uses an algorithm to decode (make sense of) the neural activity and predict the intended movement on its basis. During a heated argument, Bob punches Alex in the head. Immediately after the punch is issued, Bob claims he didn’t intend to punch Alex and that we shouldn’t hold him responsible.

In this case the claim Bob makes, at least at first glance, is less straightforward to dismiss than it is in the Sam case. Here we can consider Bob’s action in terms of some distinctions among intentions. Distal intentions are states of being committed to perform certain (types of) actions in the future. Proximal intentions are understood as intentions to start (and keep) acting now while motor intentions are believed to inform the motor system (Bratman 1987; Mele 1992; Pacherie 2006). Did BCI Bob he have previous annoyance with Alex and so potentially a distal intention to harm him? In the heat of the moment, despite no distal intention, had he formed a proximal intention? Had his heated encounter prompted formation of a motor intention, or a vivid mental representation of punching Alex? We may be uncertain among these types of intention, as can be thought of in conventional cases. But additionally, in the BCI case, we have to consider whether Bob intended to realize the brain activity that triggered the device, whether the device decoded the activity as intended, or whether the device misfired in some other way out of Bob’s control. Even were these facts to be known, the ascription of responsibility to Bob would require careful analysis of the degree of control Bob has over which parts of the process and how foreseeable particular neural activity-behavior outcome pairs are. We have not only to consider executive-implementational control dimensions and intentions, but also the relations among some variety of technical considerations to do with brain activity recording, decoding, classifying, the triggering of devices, and the functioning of those devices. We argue that instances of BCI action may differ from standard action given features of BCI design and

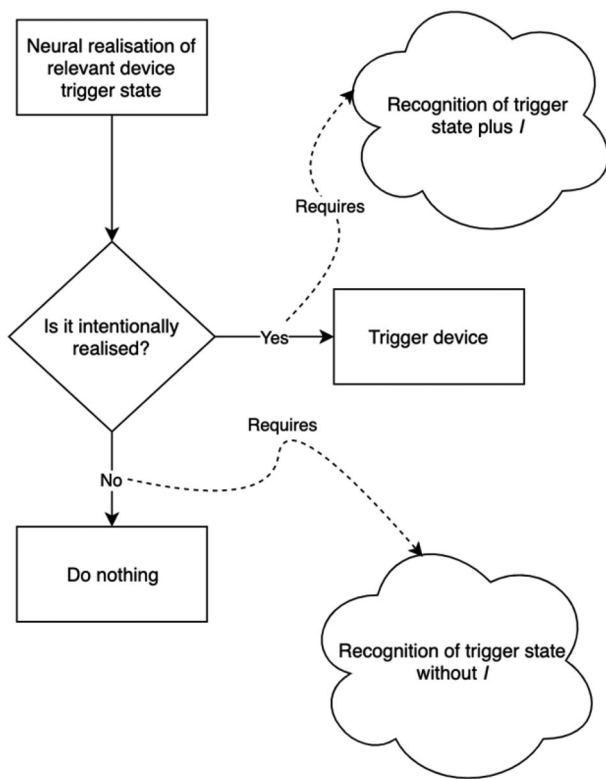


Figure 1 A schema for recognition of intentional device triggering through realizing a neural state.

use that relate to device control and the foreseeability of outcomes.

Figure 1 illustrates an outline for how intentional realization of a neural state as device trigger would need to be recognized to ensure full control of a device. The executive dimension of control over action may be ready to go, and a goal set. The translation of this into actual action, the implementational “I” factor, is not clearly identifiable. The “I” factor is the difference. What this *I* factor might consist in, neuroelectrically, or neuroanatomically, is unclear. Research has shown various details about the neuroelectric and other dimensions of consciously “wanting to move” (Desmurget and Sirigu 2012) or other planned activities (Snyder, Batista, and Andersen 2000), but this is only one sort of case. The point at which, a perhaps general plan or intention translates into action is broader than this. Empirically, patterns of activity in the medial frontal cortex of rats can predict waiting times before acting (Emmons et al. 2017). But this does not address the “I” factor in human imaginings and intendings as with, for example, having a plan to get a drink versus actually doing all it takes to get some specific drink, and to drink it. Between the having and the doing may be any number of factors in play, and any stretch of time.

As well as the difficulties of accounting for the decision to move from executive implementational control, there are also a variety of accounts of intention to consider. For instance, intentions may be about acting now or in the future. They may be about whether to act discretely or continuously, or to move in a specific way, or to refrain from so moving. Presumably, there are a number of factors that determine whether intention surfaces as action. There may be executive control-updating factors, meaning changed plans, or subconscious influences or checks on executive or implementational control. There may be similar influences on the variety of intentions that are formed which are not understood, and which fail to be integrated into BCI programming. Such factors would not easily be featured in BCI programming as they are not well understood, so could not be accounted for in a system. The full range of such factors requires a complete understanding of mental processes. Given the variety of the philosophical issues behind this, including how to define an action, intention, agency, and how to account for reasons and causes, this seems a long way off (See, e.g., Mele 2009, 1997; Searle 2001; Davidson 2001; Mele 1992). This is not to play down the rich interactions among philosophical accounts of action, cognitive science, the neurosciences, and so on. Mutually enriching dialogs are evident in a range of interdisciplinary articles and conferences. But as long as we lack a fuller understanding of the variety of dimensions indicated here, we will similarly lack a notion of what “complete” BCI control would be.

One might have a distal intention to make coffee throughout the activity of reading a paper. At some point, the paper is set aside, and coffee-making is undertaken. The distal intention “wanting to make coffee” here is rather nebulous, and coincides with the activity resulting from the proximal intention “to read a paper”. It is not clear that at some specifiable, concrete point in time the distal intention overrides the proximal, and implementational control shifts to serve one executive goal rather than the other. The coffee-making intention state is realized, but stands until some point at which it prompts action. This point is what the *I* factor might coincide with. The call to action seems to be “coffee-making-plus-*I*”. Its detection would require more than is presently, or perhaps in principle, possible to detect. Hence, there appears to be a ceiling on possible control of a neuro-controlled device.

In short, control of neuro-controlled devices will be predictably limited because despite knowing a lot, there is much unknown about the neural correlates of

intentional action. This puts a limit on the amount of control available to a device user at the level of device design, no matter how proficient a user may be with a device. Specific device use is thereby predictably limited in terms of potential user control. We will argue in the final section of this paper for a mitigating factor regarding responsibility ascriptions when predictably uncontrolled device use occurs. This will have to do with the nature of device use as necessary, or recreational. Before that, more needs to be said about BCI mediated action.

As well as the global constraint on possible control at the design level, there are practical, fine-grained limits on user control of specific BCI devices. Users of BCIs may have reduced control given that use of the device is:

1. Difficult and time-consuming to master
2. Sometimes involuntary, since the brain states that serve as the trigger might occur involuntarily
3. Only ever partial, as the BCI's processing partly determines the outputTaken together, these three limits on user device control have implications for the foreseeability of outcome from device use.
4. Outcomes of using the device are not always predictable

On the one hand, having limited control of devices seems to suggest device users ought to be considered less responsible for their actions mediated via BCIs. On the other hand, it is predictable that devices will be only partially controllable. Users of devices then seem to attract more responsibility for their BCI mediated actions as they ought to know they will have only partial control. At the extreme, this might indicate recklessness. This comes from what we know of the global design constraint on possible controllability, and the four factors mentioned about specific device control and outcome foreseeability. There seems to be a tension within accounting for responsibility where BCI mediated action arises, as contrasted with conventional action.

THE PROCESSES UNDERLYING BCI-MEDIATED ACTION

There are a variety of types of BCI, not with exactly the same characteristics. These include varying degrees of invasiveness of recording methods, BCIs that are active versus those that are passive, or others that are reactive (Steinert et al. 2018). Throughout the following, our examples will typically refer to active BCIs,

meaning devices that are controlled through deliberate realization of neural states by a user. These BCIs operate through the recording and decoding of neural activity. Part of this involves identifying relevant brain signals, processing them, and decoding them in order to realize them as control signals for a subsequent device, such as a wheelchair, speech prosthesis, or drone (Pan et al. 2017; Lin and Jiang 2015; Arjestan, Vali, and Faradji 2016; Bocquelet, Hueber, Girin, Chabardès, et al. 2016). One reason we might label this “thought control” is that significant training and mental discipline is required in order to be able to use these systems reliably. While a BCI can easily detect neural signals, the challenge comes from identifying relevant signals and processing them appropriately. For consumer devices, we can add to this the challenges of recording and processing in a non-research environment, with a moving head, and lots of potential user distraction. Among other things, input besides the user's intentions to operate the BCI controlled system include the system processing itself, and user inconsistency. Further constraints can come from the environment in which the system is operated.

Returning to Bob's unintended punch, there are a few ways in which this could have occurred. He may have mis-controlled his arm through just happening to realizing a brain signal as a control command. This could have occurred as a result of Bob involuntarily vividly imagining punching, for instance. Alternatively, the BCI may have mis-identified one brain signal as another and issued a control command where there was none. For example, in many ordinary circumstances we form a vague or hazy intention, but the precise call to action is mysterious to us. Were a device to be triggered by a hazy intention it would not necessarily mean the system was broken but only that, neuroanatomically, one signal was *close enough* to another to be processed as a command. This is an issue in identifying what signals are, or ought to be “relevant” for the system to decode. This is not just a signal discrimination issue, moreover, as illustrated in the discussion following [Figure 1](#) above.

Other factors that could lead to unintended system activations, or action consequences, would include mechanical problems with the device itself, or emergent conditions in the physical environment. What is of interest here is the user-generated control parameters for the technology and their role. Because these are brain signals, it will be helpful to discuss how some such signals are involuntary and others voluntary. This will have a bearing on how the intuition

from above—that Bob is not obviously as responsible as Sam—is developed.

Involuntary and Voluntary Neural Acts

The majority of neural activity is involuntary, and serves basically to keep us upright and breathing. Between this life-sustaining activity and explicitly intentional activity lies a range of phenomena that are more or less well understood. For example, involuntary recall of episodic memory can be distinguished from voluntary recall in terms of dorsal frontal region activation (Hall, Rubin, and Miles 2014). But the very fact that something as complicated as a memory can seem to flash before the mind uncalled tells us how little we understand control of the mind, and associated neural activity. We may undergo an involuntary recollection of some vivid memory. Conversely, we might understand well the neural story about memory retrieval, and yet we might be completely unable to recall some memory at will.

In another sense, however, examples of attempted remembering are evidence of some neural control, even when the memory sought remains elusive. Studies into non-sensory fringe experience, like “tip of the tongue” phenomena in which one mentally searches for a specific word, show that frontal brain areas are active during such experiences (Baars 2003, 9–10). Tip of the tongue phenomena include the neural act of searching for a word that can not quite be remembered. These acts are not aiming for doing anything other than remembering something. The brain regions active during these acts are mostly non-sensory. Hence, this evidence suggests that neural action, not aimed at sensory behavior, has measurable effects in the brain. Kirmayer and Gold suppose that this empirical evidence supports the idea that during neural action “... we *use* our brains” (2011, 317 emphasis in original). In other words, we voluntarily control neural activity in specific acts of trying to remember whether we remember or not. With this and involuntary memory recall taken together, we appear to have a kind of control over our neural activity. We can use our brains to try to remember, but also be subject to its activity when memory appears involuntarily.

Even this partial control can be seen as what makes BCI use possible at all. BCI use is an acquired skill that requires training (Wolpaw et al. 2002, 76ff). Training for using a neuro-controlled limb prosthesis, for example, can take “weeks”. Farahany notes that learning to control such a device via neuronal activity

requires the user “to learn a new language” (Farahany 2011, 10) which amounts to translating familiar intentions to act in this way or that into novel neural acts that themselves ground novel outcomes. Instead of trying to move my arm, for example, as I may have all my life, I now use another thought that triggers movement of the prosthesis.

We can exercise control over our brains, and we can act neurally, and so we can establish voluntariness about some types of neural activity. This must be so for BCI training to be possible. Yet the picture gained across contexts of involuntary and voluntary neural acts presents a complex picture of neural activity. What we can see is that neural activity has various dimensions that interact in producing overt behavior, as well as other neural activity. There are dimensions of these interactions that are voluntary, and that can be acquired by training, so we can say that it can be done poorly, or done well, just as with any other acquired skill. This provides us with a basis to posit a degree of freedom of action in terms of neural activity. We appear to be partially in control of some neural states, and more like subject to others. Where this kind of control is used as a means of controlling BCIs, we may have only partial control of the actions mediated by those BCIs.

BCI ACTION AND RESPONSIBILITY

Control of a neural device, triggered by patterns of voluntary neural activity developed through careful training, requires analysis in terms of responsibility. We need at least to establish whether we can justify BCI use when we know we will only ever achieve partial control of it, by dint of design as well as the practicalities of specific device use. Neural control will have to be sufficient to account for responsibility ascriptions where neuro-controlled devices are used.

Using the analogy of driving a car, if I am an untrained driver behind the wheel my driving may not be criticisable in the same ways as a licensed driver’s. For instance, putting the car in the wrong gear for the circumstances won’t really stick as a criticism—I don’t know about gears owing to my lack of training. If a trained driver brings about harm, that is criticisable in a more fine-grained way. Their foresight, gained via training, ought to be such that they are well disposed to prevent specific harms ever coming about in the first place. Failing to correct a bad gear change, to slow down in dangerous conditions, or to brake smoothly, are omissions for which a trained driver ought to be held particularly

responsible. Such a driver ought to have known which gear to have been in, ought to have changed into that gear, and ought to have been able to predict the problems arising from failing to do so.

The untrained driver does not know about appropriate gears. In having such knowledge available to them, and failing to act upon it, the trained driver omits to correct a worse scenario and is responsible for that omission in a way the untrained driver is not. In knowing that they are untrained, however, the untrained driver still ought to be held responsible for their actions as they ought to have been able to foresee problems would arise from getting behind the wheel untrained.

This sense of control seems fairly straightforward to understand for the case of a neuro-controlled device. Following appropriate training, neural activity is realizable such that the device will be triggered and will function according to a predictable scheme. A limb, for instance, will move as the user desires. But we may still ask: does the control here apply to the triggering of the device, or to the movement of the limb? On Farahany's account of neural triggering of devices (Farahany 2011), the “detectable, identifiable, neural artifacts associated with specific decisions to act” are the neural correlates of triggering the device. They are what is learned in “learning a new language” to control a device. These are precisely *not* the same as those associated with conventionally moving a limb (Steinert et al. 2018; Vallabhaneni, Wang, and He 2005). The “act” here may therefore be characterized as the triggering of the device rather than subsequent activity. This might suggest the user of a neuro-controlled device is more like the untrained than the trained driver.

In principle, control in the case of a neuro-controlled device may be akin to switching on a black box. A user might well know how to activate their device, and get it to do what they wish, but perhaps not really know how. Besides recording neural activity, the device will enact some kind of signal processing, which will involve altering that original input to some extent—amplification, filtering, signal extraction. BCI device users clearly have control of their device, but can they be said to have the sufficiently fine-grained sort of control we would expect of the trained car driver from above? This is not to suggest that one must know all the intricacies of a device's functioning if responsibility is to be ascribed. A highly skilled driver may know nothing of how an engine works. Rather, it is raising a question about the nature of the act associated with the use of a neuro-controlled

device: is the act complete in the triggering, in what subsequently occurs, or some sort of combination?

If I had a car which, once started, drove automatically giving me only steering ability this would be different to the standard case of driving. Certainly, I would have control in terms of each act of steering. I would have control in terms of being able to steer this way or that. But, can we call this “driving” and hold me to account as we would a trained driver in a typical car, complete with access to engine control, brakes, and the rest? The physical state of the car seems an important factor if we are to be able to distinguish between a coarser or finer-grained notion of control, not least in terms of the foreseeability of outcomes.

Driving is not simply an ensemble of discrete events, ordered in time. Driving is a skill that can be acquired, that can be reduced to simple, discrete actions. But the reduction would miss the point—driving as a patterned activity in itself. How this is seen affects how we can assess the foresee ability of outcomes from the perspective of the controller. In terms of the car example, this is like the distinction between the trained and untrained driver. In the case of BCI action, the themes are similar and will condition how we ought to ascribe responsibility to specific BCI-mediated acts.

Steinert et. al. (2018) note several “peculiarities” of BCI-mediated action including reference to something like that being developed here. For Steinert et. al., all BCI-mediated action is “non-basic”. This means that all BCI-mediated action is action done by doing something else—moving a prosthetic limb by trying to realize a neural pattern, for example. This is in contrast with the conventional moving of a limb by trying to move a limb. We suggest here that the nature of an action can be questioned if the actor is engaged primarily in controlling a device, rather than primarily controlling the action the device realizes.

RESPONSIBILITY DIFFERENCES BETWEEN BCI ACTION AND CONVENTIONAL ACTION

Conventionally, if I catch a thrown ball my action is the catching of the ball (or the trying to catch the ball). The conventional catch is catch-focused. The BCI-mediated catch might be better described as device-control focused. BCI-mediated actions *relate* to their desired outcomes but are perhaps not as closely identifiable with them as in conventional action. On this analysis, the acts the BCI-user carries out are the basic actions that constitute the overall BCI-mediated

action. These are neural acts. Moving a BCI-controlled prosthetic arm can take a lot of concentration, and involve repeating to oneself something like “right, right, right” in order to effect movement (Khatchadourian 2018). Conceptually, this seems more like the case of trying to catch a ball with a net, and so more akin to tool use.

In conventional cases of catching we do not typically fixate on our limbs, but the object to be caught. We do not typically have to will our arms or hands to move in specific ways. If we did, we would probably catch less often than we otherwise do. In terms of BCIs we have to think in terms of issuing “go-commands” for specific movements, which is unlike conventional action. This shows a complexity in the causal chain not present in conventional action, where we implement a process (e.g., of catching, rather than sets of arms movements resulting in a catch).

Recalling the difference between executive and implementational control, and the technical means of achieving high levels of control over a neuro-controlled device, this can be clarified further. It was suggested above that to achieve a high level of control over a prosthetic hand could require (1) high decoding performance of motor tuning from neural activity, or (2) the decoding of specific action goals to be implemented by a “smart” neuroprosthetic device; or (3) a combination of (1) and (2). In conventional cases of catching, executive control sets a goal of catching, and this is implemented as per years of experience with a spatio-temporal world of throws and catches.

With the technological case, implementational control is ceded to the decoding of neural activity. It may further be taken over by a “smart” limb, bestowed with sensors to modify movement in response to contextual factors (Hochberg et al. 2012). The user has no control over the decoding itself. Where a smart device is present, executive control too may interpreted by the decoding operation of the system. Control over goals is thus diminished. Questions over the nature of the action here are prompted, as well as responsibility for outcomes from those actions. For example, the actions themselves may be akin to tasks given to the device by the user, like an employee, and there may be a responsibility gap to be accounted for (Stephen Rainey 2018; Kellmeyer et al. 2016; Klein et al. 2015).

Triggering and Control

Neural control of devices involves a lesser degree of control over the actions they mediate. This is not least

because neuro controlled devices might be triggered by involuntary neural activity. This is on the one hand a way of thinking about a type of brain machine interface decoding error (Milekovic et al. 2012, 2013). These types of errors refer to effectors moving in unintended ways from a user perspective. Another example comes from developments in creating a speech neuroprosthesis operated by means of covert speech.

Neuroanatomically, covert speech is similar enough to overt speech to permit its use as a trigger for a voice synthesizer (Bocquelet, Hueber, Girin, Savariaux, et al. 2016) But the mechanism of realizing covert speech (e.g. “shouting” a word inside one’s head) might be similar enough to some forms of inner speech not intended for externalization so as to trigger the speech synthesizer unintentionally (S Rainey et al. 2019 in press). Here, the “speaker” would say things they did not wish to, through intentional cognitive activity, but with an unintended outcome. A speech neuroprosthesis would need to be able to detect a difference between “covert speech” and “covert speech plus I” in order to mitigate this completely. But what “I” is remains unknown.

It is a small step to take to think of sensitive BCIs effecting activity not intended at all, as opposed to effected badly. At the very least, if this is a small possibility with present technology, probable developments in future devices will make such accidental triggering more likely.

Given challenges in signal acquisition, caused by electrode placement, demands of spatio-temporal resolution, and the dynamics of brain signals, decoding and classification of signals is likely to become more prediction-based, and to use more artificial intelligence. The convergence of predictive neural decoding strategies with BCI technologies has implications for motor prostheses (Truccolo, Hochberg, and Donoghue 2010). Such developments are likely to exacerbate the potential control deficit we are discussing. This is hardly science fiction. Rather, it seems an obvious next step in developing technologies that are robust in decoding fast-changing input.

Predictive classification would address issues already present in the literature on neuroprosthetics for speech. Moreover, in a further bid to improve signal acquisition in the case of speech neuroprosthetics, data fusion techniques using multiple sensor inputs could serve well (Denby et al. 2010). The complexities of such fusion are thought to be handled well via use of artificial intelligence approaches (Glaser et al. 2017). An increase in predictive decoding will serve to mitigate issues of signal acquisition, as well as increase

processing efficiency. The convergence of predictive decoding and artificial intelligence in BCI technology will produce issues in terms of control due for consideration now.

In the case of conventional action, even if I want to throw an object in my hand and am vividly imagining doing so, I can easily refrain. The actual throwing of the object in the conventional case is controlled by more than the mere realization of a brain state. It requires onward neural activation to move muscles. This is further conditioned by the possibility of subsequent intentional modification. For example, I may try to catch a ball. While in the process, I see I may burn myself and so I modify my action. This complex picture of action is nonetheless the familiar, conventional one. It amounts to implementing and action, not issuing a series of go commands for movements, via realizing sets of brain states.

In some sense, of course, the movement of an arm is triggered by the realization of a brain state. But the conventional act sees brain states realized in the implementation of the action, not as its basis. What's more, conventional action requires the realization of an additional brain state (a decision to act) which does not necessarily accompany an imagined action. In some cases of BCI-mediated action, the mere realization of a brain state may be sufficient to trigger the BCI, but in the absence of a decision to initiate an actual movement. The nature of the device's own processing may be such that they are triggered before the user's decision to act is realized. This is one way of coming at the global control constraint mentioned at the outset. Without detailed knowledge of the neural differences between planned, intended, desired, merely imagined (etc.), courses of action, there will be a ceiling in terms of possible control over BCI devices. If we take as an example a BCI-mediated speech device the issue can be illustrated clearly.

Neural speech prostheses promise to record the neural signals associated with vividly imagined, but un verbalised, covert speech from which can be decoded into overt speech features (Bocquelet, Hueber, Girin, Chabardès, et al. 2016). Software processing of the signals, using neural net computing for example, allows reconstruction of acoustic features to represent the covert speech (Bocquelet, Hueber, Girin, Savariaux, et al. 2016). Together, this recording, processing, and reconstruction enables a system to externalize covert speech (Chakrabarti et al. 2015; Mugler et al. 2014). The technology has applications in various medical contexts, such as in cases of aphasia, locked-in syndrome, and speech pathologies where motor function is compromised but neural control is not.

It is easy to imagine that, from the recording of brain signals at least, a BCI could be overly sensitive in decoding those signals. Without a veto control that allowed a speech BCI user to stop outputs happening, this could result in an involuntary externalization of more covert speech than the speech BCI-user might desire (Steinert et al. 2018; Clausen et al. 2017). Any activity that was neuroanatomically covert language-like could be externalized. Veto control might amount to letting the “speaker” hear the proposed output before it is actually externalized as synthetic speech. While this would slow things down, likely to below that rate of normal conversation, we know that not every instance of imagining speech coincides with a desire to externalize that speech. Thus, veto control would be of value. A system that was triggered by the mere presence of language-like brain signals might be too sensitive. We would rightly question ascriptions of responsibility to the “speaker” for the content of speech externalized by such a system. We know enough about how thought and language work that we would not think of the mere occurrence of imagined speech as equivalent to conventional speaking.

In the BCI-mediated action case, where the brain state might be the triggering state, the kind of veto control possible for the speech system may not be available. Refraining from doing something may not clearly arise as a possibility from the actor perspective. As Glannon notes, “It is not clear how a subject who intends to perform a movement that is predictable on the basis of neural EEG signals could change her mind, cancel the intention and refrain from performing it.” (Glannon 2016). Even if a veto control of some kind were envisaged for physical action, perhaps as a verbal description of the intended action could be played for the user before actual execution of the movement, this would be problematic. The use of BCI controlled limbs can be very taxing, and so adding further steps to the process might be to overburden the user. Moreover, this could rule out acting quickly, and so preclude urgent action for users of the technology.

These discussions should serve to substantiate the claims made in the opening paragraphs, relating to the use of devices and their outcomes in terms of BCI action. As well as noting a ceiling on possible control over a neuro controlled device per se, we claimed that the use of any specific device is:

1. Difficult and time-consuming to master
2. Sometimes involuntary, since the brain states that serve as the trigger might occur involuntarily

3. Only ever partial, as the BCI's processing partly determines the output and that the outcomes of using the device are:
4. Not always predictable

In Section “BCI action and responsibility”, we discussed distinctions between more or less fine-grained control especially as they relate to training. In doing this, we established the case for neuro control of devices to be no different in principle, establishing 1. Section “The Processes Underlying BCI-mediated Action” discussed types of neural activity, including that which is involuntary. This was bolstered with further discussion of how BCI devices could be triggered by neural activity not intended to result in action. These points served to substantiate 2. The partial control claim in 3 was grounded in discussion of BCI action as “non-basic”, and reliant upon system features over and above the intentions of the user of a BCI device. The neural basis of intention, and differences among voluntary and involuntary neural states, moreover, is not clearly understood. These same points, and those raised especially in Sections “Responsibility differences between BCI action and conventional action” and “Triggering and Control” spoke to the unpredictability of BCI device-mediated action, as illustrated with the analogies with car driving from Section “BCI action and responsibility”. Evidence for the possibility of a responsibility gap was developed, especially where the types of functioning of devices challenge executive and implementational control. Overall, it seems apparent that there are morally relevant distinctions to be drawn between BCI action and conventional action.

The tension referenced at the outset has been established. This tension was between whether use of a device we know to be only under partial control ought to attract more or less responsibility for the user. A brief exploration of some contexts of BCI action, distinguishing between “necessary” and “recreational” use, will serve to clarify things. The upshot will provide conditions for degrees of responsibility for BCI users in different circumstances.

THE MORAL RELEVANCE OF NECESSARY USE

“Necessary use” of a BCI is that use predicated on an actual need. Necessary use of a device means that an actor could not act in some given way but for their device. For instance, paralysis-prompted use of a BCI-controlled speech system. By contrast, recreational use is use of a device out of desire. Perhaps a BCI operated drone is used from sheer interest in the

technology. Maybe a BCI operated vehicle is used despite having no impairment to walking. We explore the moral relevance of necessary use versus recreational use with the following two imagined types of action mediated via BCI controlled devices:

Type A: Necessary use case: Necessary Nina

A significant physical disability dictates that in the absence of a BCI mediated device Nina will not be able to live her life as she would like to. Nina relies upon a BCI controlled wheelchair for her everyday life.

Type B: Recreational use case: Recreational Rick

Rick is excited about BCI technology and wishes to enjoy himself by using a variety of BCI controlled devices for fun, and to carry out tasks. He could find other ways to enjoy himself, and could carry out the tasks without relying on BCI controlled devices.

Cases of type A are those where people with disabilities use assistive technologies, for example. As technology grows more advanced, it seems a good use of these advances to make more capability-restoring devices, and to make them more seamless to control. Cases of type B may seem less familiar, but are nonetheless a growing area as BCI technology moves into the consumer market. Already, apart from computer game controllers and neurofeedback devices, more exotic BCI devices are becoming available (Roelfsema, Denys, and Klink 2018).

Let us say that Nina, from the type A scenario above, accidentally drives a BCI controlled wheelchair into a bystander, causing some injury. It might be argued that this sort of BCI mediated action is quite different from that of a conventional sort. The “act” in this context is a mental act—the triggering and/or control of a device by realizing the neural correlates of particular trained thought activity. On the face of it, this is different from conventional acting in which bodily movements, the outcomes of decisions, constitute the action. In the latter case, we can ascribe responsibility for actions and outcomes quite straightforwardly. In the former, things appear more obscure.

If Nina's act is really the result of a mental phenomenon, perhaps we ought not to ascribe to her responsibility for the accidental consequences in the world. To do so would appear to place constraints upon Nina's mental activity, rather than physical as in conventional cases of action. This seems a radical step—bringing the hitherto private, subjective realm of the mental into the public, objective realm of praise, blame, and reciprocity. Indeed, in the legal context, Steinert et al. (2018) argue, this is so far the case. After all, Nina did not intend to drive into anyone, and certainly she did not move her body so as to

realize the eventual outcome. There is no *mens rea* or *actus rea*, in this case. There may be no clear and adequate legal provision capable of handling ascriptions of responsibility for BCI-mediated action.

BCI-mediated action may not qualify as “action” in a legal sense as legal definition of action involves reference to willed movement of the body (Munoz-Conde and Chiesa 2006). This physical criterion may be missing in BCI-mediated action. As such, no BCI action might be considered legally relevant. This seems a rather absurd potentiality. To draw this out, we can re-think the scenario in type B terms.

If Rick is using a BCI controlled wheelchair *recreationally* and accidentally injures someone we might immediately think this is worse in comparison to Nina’s accident. It might seem worse because there is no need for Rick to be using the wheelchair at all, and so the risk he took in so using it was reckless. If, as was just suggested, BCI mediated action does not count as action this intuition would be baseless. Rather than that, it seems plausible that BCI-mediated action ought to be considered as something to which responsibilities, consequences, intentions and so on, can attach. Fleshing out this intuition would therefore give a basis for updating legal definitions of action in order to incorporate these emerging technologies and the action they enable.

Accidental Outcomes and Failures of Control

The control of a BCI system will always have the possibility of accidental outcomes. Even if small, the unpredictability of control is itself predictable owing to factors to do with recording, decoding, user distraction, and so on. If Nina routinely drives her wheelchair into people, each time accidentally, we would be in a position to say she is unskilled at using the chair. But we can ask further, what does “accidental” mean here? It could mean that:

1. Nina intended to move the wheelchair, saw the bystander but could not correct the chain of events
2. The BCI malfunctioned and took the incorrect path
3. She did not notice the bystander because her attention was consumed with trying to control the wheelchair
4. She failed to take adequate account of the chances of hitting the bystander if she moved he wheelchair

Cases 1 and 2 seem to leave Nina blameless. In 1, the system is too unresponsive, or the bystander was simply too close. In 2, again, simple error. In the case

of 4, and possibly 3, it does seem that Nina is responsible for the harmful outcomes. We might be inclined to provide Nina with more or better training in case 3. In case 4, we might be inclined to modify her software to make some evasive functions automatic.

In the case of Rick however, were he routinely to accidentally drive into people we might want to simply stop him using the technology. His blameworthiness for each accident is worse than Nina’s. There seems to be this moral difference between type A and type B cases. Where a person cannot act but for their BCI device, we argue that ascriptions of responsibility ought to be nuanced by an understanding that a given context may not be optimally accessible to that person. Were I to hit you with my prosthetic limb through an aberrant mis-control, while it is my fault, I am less blameworthy than were I to do similarly but in a case where I operate that prosthesis simply recreationally. If I just *decide* to use a BCI device despite it not being necessary, the levels of blameworthiness my actions ought to attract should be considered on a different scale to those for a user who cannot act in some way but for their BCI.

BCI-mediated action ought to be, on the whole, treated as conventional action, perhaps as instances of complicated tool use. Deviations from this general approach might come in small degrees, motivated by the sense in which disability can prompt “reasonable adjustment” in evaluating action. Nevertheless, there are morally relevant distinctions to be drawn where recreational, or type B cases, of BCI-mediated actions arise.

Differences in responsibility-ascription between BCI-mediated action cases of type A and type B might be accounted for in terms of the rights of persons with disabilities. The UN has the Convention on the Rights of Persons with Disabilities (CRPD), while the EU includes text in its charter of fundamental rights, relating to the necessary recognition of differential needs of persons with disability such that societal inclusion is maximized. In UK domestic employment law, the concept of an employer’s duty to make “reasonable adjustment” does this work quite straightforwardly: “The duty to make reasonable adjustments aims to make sure that as a disabled person, you have, as far as is reasonable, the same access to everything that is involved in getting and doing a job as a non-disabled person.”¹ The idea is that in the controlled environment of a workplace, measures ought to be taken to ensure that that environment accommodates differences as they emerge from the result of disability.

¹<https://www.equalityhumanrights.com/en/multipage-guide/employment-workplace-adjustments>

We could suggest extending this concept behind this duty to include contexts beyond the workplace. Given uncontrolled environments beyond the workplace, we could say,

Where a disabled person cannot act in some way but for a BCI device, a reasonable adjustment in considering their BCI-mediated action ought to be made in order that appreciation of action as being enacted in a non-optimal environment is acknowledged.

What we are talking about here is the causing of harm through predictably unpredictable outcomes from the necessary use of BCI devices. As such this covers a variety of actions, some possible now with existing technologies, others more notional and future-facing. The principle ought to apply to cases where control is not certain, but use required somehow. What is excluded from the discussion, as a different though interesting point, is criminal activity. Crimes committed via BCI should presumably be treated as any other crime, as they are predicated on criminal intent rather than deficits in control. Nevertheless, puzzles may remain regarding how to make sense of BCI criminality (Steinert et al. 2018, sec. 5).

Although this paper has been discussing BCI centrally, the text here uses a narrow framing in terms of BCI. It may be that there is no reason not to expand the idea here to other sorts of assistive technologies, such as blink-controlled devices, as these are more widespread as just as apt for being difficult or unpredictable to control.

CONCLUSIONS

The actions mediated by BCI devices seem, by and large, similar to conventional actions. It seems clearly not the case that, because the “act” involved in triggering and guiding BCI devices is mental, that actors are automatically exculpable for actions. Ethically relevant distinctions between conventional and BCI-mediated actions revolve around issues of control and the foreseeability of outcomes. Further dimensions of importance cluster around necessary versus recreational uses of BCI technologies. Analysis of these motivate a modification of the law in which action is “willed bodily movement”. The idea of “reasonable adjustment” can play a useful role in this.

Where use of a device is required for some action for reasons of disability, BCI-mediated action that deviates from standard norms or that leads to some kind of harm ought to be accommodated. Exculpations ought to be forthcoming from fellow actors in just the same way they ought to should a

person with limited mobility walk slowly in a rushing crowd. Where there is a clear reason for the use of a BCI technology for assistance or owing to disability, it appears easy to assimilate this kind of action.

Cases of type B, increasingly enabled by commercial and consumer appetite for BCI technology, require further conceptual and legal preparation. Where BCI use arises recreationally, without a reason deriving from a case of disability or something similar, it ought to attract more scrutiny than where such a reason is present. This case is like being an unqualified driver but getting behind the wheel of a car. A recreational adoption of an action strategy that can foreseeably produce harms makes the action performed via that strategy potentially ethically hazardous. It does this through general issues of potentially failure-prone tool use, and BCI specific issues of control, such as BCI-mediated action as non-basic.

It was noted that the possibility of full control of BCI devices will be limited because of gaps in knowledge concerning the neural story of intentional action (versus mere imaginings, or mental planning, for example.) This basic design constraint on the possibility of control has knock-on effects for the use of specific devices. Control will only ever be partial. This stands as an umbrella concern for neuro controlled technologies. There will always be a cap on the amount of control available to a user, even if they are maximally proficient in its use.

In terms of specific device use, BCI-mediated action via a neural controlled device may be difficult and time-consuming to master, prone to involuntary triggering, and somewhat unpredictable in terms of outcomes. This unpredictability of outcomes itself may be foreseeable, that is, we could guess that using a device to act in certain ways may be prone to failures of varying degrees. This is what is captured in the “reasonable adjustment” notion as applied to essential use of BCIs for type A cases. There is more room for exculpability in such cases, or at least latitude in ascriptions of blame, because the user in some sense *has to* use the potentially failure-prone device. This is not about patronizing those who may have to act using assistive technologies. It is a consequence drawn from the nature of a control deficit with which necessary mediation of action in devices presents such users. This is also why there is scope to put more scrutiny, and apply firmer ascriptions of blame, for harms that arise from type B cases. Such use is predictably riskier than conventional action, and so can be treated more robustly than conventional action.

BCI-mediated action ought to be, on the whole, treated as conventional action. Law and policy ought to be amended to accommodate this by broadening the definition of action as “willed bodily movement”. Where BCI-mediated action may be due different evaluation from the conventional case may be motivated by the sense in which disability can prompt “reasonable adjustment” in evaluating action. Where BCI-mediated action results in harms of some kind, this may be in some sense excusable for those who could not act at all except for with a BCI driven device. But BCI mediated action through neuro controlled devices that is recreational ought to present more moral jeopardy for actors than conventional action.

FUNDING

The authors gratefully acknowledge funding from the BrainCom Project, Horizon 2020 Framework Programme (732032) and Wellcome Trust grant no. WT104848/Z/14/Z.

ORCID

Stephen Rainey  <http://orcid.org/0000-0002-5540-6046>

REFERENCES

- Arjestan, M. A., M. Vali, and F. Faradji. 2016. Brain computer interface design and implementation to identify overt and covert speech. In *2016 23rd Iranian Conference on Biomedical Engineering and 2016 1st International Iranian Conference on Biomedical Engineering (ICBME)*, 59–63. doi: [10.1109/ICBME.2016.7890929](https://doi.org/10.1109/ICBME.2016.7890929).
- Baars, B. 2003. How brain reveals mind neural studies support the fundamental role of conscious experience. *Journal of Consciousness Studies* 10(9–10): 100–114. <http://www.ingentaconnect.com/content/imp/jcs/2003/0000010/F0020009/art00008>.
- Bocquelet, F., T. Hueber, L. Girin, S. Chabardès, and B. Yvert. 2016. Key considerations in designing a speech brain-computer interface. *Journal of Physiology-Paris* 110(4): 392–401. doi: [10.1016/j.jphysparis.2017.07.002](https://doi.org/10.1016/j.jphysparis.2017.07.002).
- Bocquelet, F., T. Hueber, L. Girin, C. Savariaux, and B. Yvert. 2016. Real-time control of an articulatory-based speech synthesizer for brain computer interfaces. *PLOS Computational Biology* 12(11): e1005119. doi: [10.1371/journal.pcbi.1005119](https://doi.org/10.1371/journal.pcbi.1005119).
- Bratman, M. 1987. *Intention, plans, and practical reason*. Center for the Study of Language and Information. Cambridge, MA: Harvard University Press.
- Chakrabarti, S., H. M. Sandberg, J. S. Brumberg, and D. J. Krusienski. 2015. Progress in speech decoding from the electrocorticogram. *Biomedical Engineering Letters* 5(1): 10–21. doi: [10.1007/s13534-015-0175-1](https://doi.org/10.1007/s13534-015-0175-1).
- Clausen, J., E. Fetz, J. Donoghue, et al. 2017. Help, hope, and hype: Ethical dimensions of neuroprosthetics. *Science* 356(6345): 1338–1339. doi: [10.1126/science.aam7731](https://doi.org/10.1126/science.aam7731).
- Davidson, D. 2001. *Essays on actions and events: Philosophical essays*. Oxford, United Kingdom: Oxford University Press, UK.
- Denby, B., T. Schultz, K. Honda, T. Hueber, J. M. Gilbert, and J. S. Brumberg. 2010. Silent speech interfaces. *Speech Communication* 52(4): 270–287. doi: [10.1016/j.specom.2009.08.002](https://doi.org/10.1016/j.specom.2009.08.002).
- Desmurget, M., and A. Sirigu. 2012. Conscious motor intention emerges in the inferior parietal lobule. *Current Opinion in Neurobiology* 22(6): 1004–1011. doi: [10.1016/j.conb.2012.06.006](https://doi.org/10.1016/j.conb.2012.06.006).
- Emmons, E. B., B. J. De Corte, Y. Kim, K. L. Parker, M. S. Matell, and N. S. Narayanan. 2017. Rodent medial frontal control of temporal processing in the dorsomedial striatum. *The Journal of Neuroscience* 37(36): 8718–8733. doi: [10.1523/JNEUROSCI.1376-17.2017](https://doi.org/10.1523/JNEUROSCI.1376-17.2017).
- Farahany, N. A. 2012. A neurological foundation for freedom. *Stanford Technology Law Review* 2012(4): 1.
- Glannon, W. 2016. Ethical issues in neuroprosthetics. *Journal of Neural Engineering* 13(2): 021002. doi: [10.1088/1741-2560/13/2/021002](https://doi.org/10.1088/1741-2560/13/2/021002).
- Glaser, J. I., R. H. Chowdhury, M. G. Perich, E. M. Lee, and P. K. Konrad. 2017. Machine learning for neural decoding. *ArXiv:1708.00909 [Cs, q-Bio, Stat]*, August. <http://arxiv.org/abs/1708.00909>.
- Hall, S. A., D. C. Rubin, A. Miles, et al. 2014. The neural basis of involuntary episodic memories. *Journal of Cognitive Neuroscience* 26(10): 2385–2399. doi: [10.1162/jocn_a_00633](https://doi.org/10.1162/jocn_a_00633).
- Hochberg, L. R., D. Bacher, B. Jarosiewicz, et al. 2012. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* 485(7398): 372–375. doi: [10.1038/nature11076](https://doi.org/10.1038/nature11076).
- Kellmeyer, P., T. Cochrane, O. Müller, et al. 2016. The effects of closed-loop medical devices on the autonomy and accountability of persons and systems. *Cambridge Quarterly of Healthcare Ethics* 25(4): 623–633. doi: [10.1017/S0963180116000359](https://doi.org/10.1017/S0963180116000359).
- Khatchadourian, R. 2018. How to control a machine with your brain. Accessed November 19, 2018. <https://www.newyorker.com/magazine/2018/11/26/how-to-control-a-machine-with-your-brain>
- Kirmayer, L. J., and I. Gold. 2011. Re-socializing psychiatry: critical neuroscience and the limits of reductionism. In *Critical Neuroscience*, ed. Suparna Choudhury and Jan Slaby, 305–330. Oxford, UK: Wiley-Blackwell.
- Klein, E., T. Brown, M. Sample, A. R. Truitt, and S. Goering. 2015. Engineering the brain: Ethical issues and the introduction of neural devices. *Hastings Center Report* 45(6): 26–35. doi: [10.1002/hast.515](https://doi.org/10.1002/hast.515).
- Lin, J.-S., and Z.-Y. Jiang. 2015. Implementing remote presence using quadcopter control by a Non-invasive BCI device. *Computer Science and Information Technology* 3(4): 122–126. doi: [10.13189/csit.2015.030405](https://doi.org/10.13189/csit.2015.030405).
- Mele, A. R. 1992. Acting for reasons and acting intentionally. *Pacific Philosophical Quarterly* 73(4): 355–374. doi: [10.1111/j.1468-0114.1992.tb00344.x](https://doi.org/10.1111/j.1468-0114.1992.tb00344.x).
- Mele, A. R. 1997. *The philosophy of action*. Oxford, UK: Oxford University Press.

- Mele, A. R. 2009. *Effective intentions: The power of conscious will*. Oxford, UK: Oxford University Press.
- Milekovic, T., T. Ball, A. Schulze-Bonhage, A. Aertsen, and C. Mehring. 2012. Error-related electrocorticographic activity in humans during continuous movements. *Journal of Neural Engineering* 9(2): 026007. doi: [10.1088/1741-2560/9/2/026007](https://doi.org/10.1088/1741-2560/9/2/026007).
- Milekovic, T., T. Ball, A. Schulze-Bonhage, A. Aertsen, and C. Mehring. 2013. Detection of error related neuronal responses recorded by electrocorticography in humans during continuous movements. *PLoS One* 8(2): e55235. doi: [10.1371/journal.pone.0055235](https://doi.org/10.1371/journal.pone.0055235).
- Mugler, E. M., J. L. Patton, R. D. Flint, et al. 2014. Direct classification of all American English phonemes using signals from functional speech motor cortex. *Journal of Neural Engineering* 11(3): 035015. doi: [10.1088/1741-2560/11/3/035015](https://doi.org/10.1088/1741-2560/11/3/035015).
- Munoz-Conde, F., and L. E. Chiesa. 2006. The act requirement as a basic concept of criminal law. *Cardozo Law Review* 28: 2461.
- Pacherie, E. 2006. Toward a dynamic theory of intentions. In *Does consciousness cause behavior?* 145–167. Cambridge, MA: MIT Press.
- Pan, P., G. Tan, P. Wai, and A. Aung. 2017. Evaluation of consumer-grade EEG headsets for BCI drone control. *IRC conference on science, engineering, and technology*. Available at: <https://oar.a-star.edu.sg/jspui/handle/123456789/2149>
- Rainey, S., H. Maslen, P. Megevand, L. Arnal, E. Fournier, and B. Yvert. 2019. Neuroprosthetic speech: The ethical significance of accuracy, control and pragmatics. *Cambridge Quarterly of Healthcare Ethics* 28(4):657–670. doi: [10.1017/S0963180119000604](https://doi.org/10.1017/S0963180119000604).
- Rainey, S. 2018. A steadying hand?: Ascribing speech acts to users of predictive speech assistive technologies. *Journal of Law and Medicine* 26(1): 44–53.
- Revell, T. 2018. Mind-reading devices can now access your thoughts and dreams using AI. *New Scientist*, 26 September 2018. <https://www.newscientist.com/article/mg23931972-500-mind-reading-devices-can-now-access-your-thoughts-and-dreams-using-ai/>
- Roelfsema, P. R., D. Denys, and P. C. Klink. 2018. Mind reading and writing: The future of neurotechnology. *Trends in Cognitive Sciences* 22(7): 598–610. doi: [10.1016/j.tics.2018.04.001](https://doi.org/10.1016/j.tics.2018.04.001).
- Searle, J. R. 2001. *Rationality in Action*. Jean Nicod Lectures. Cambridge, MA: MIT Press.
- Shepherd, J. 2015. Conscious control over action. *Mind & Language* 30(3): 320–344. doi: [10.1111/mila.12082](https://doi.org/10.1111/mila.12082).
- Snyder, L. H., A. P. Batista, and R. A. Andersen. 2000. Intention-related activity in the posterior parietal cortex: A review. *Vision Research* 40(10–12): 1433–1441. doi: [10.1016/S0042-6989\(00\)00052-3](https://doi.org/10.1016/S0042-6989(00)00052-3).
- Solon, O. 2017. Facebook has 60 people working on how to read your mind. *The Guardian*, 19 April 2017, sec. Technology. <https://www.theguardian.com/technology/2017/apr/19/facebook-mind-reading-technology-f8>.
- Steinert, S., C. Bublitz, R. Jox, and O. Friedrich. 2018. Doing things with thoughts: Brain-computer interfaces and disembodied agency. *Philosophy & Technology* 32(3): 457–482. doi: [10.1007/s13347-018-0308-4](https://doi.org/10.1007/s13347-018-0308-4).
- Truccolo, W., L. R. Hochberg, and J. P. Donoghue. 2010. Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nature Neuroscience* 13(1): 105–111. doi: [10.1038/nn.2455](https://doi.org/10.1038/nn.2455).
- Vallabhaneni, A., T. Wang, and B. He. 2005. Brain–computer interface. In *Neural engineering*, 85–121. Boston, MA: Springer.
- Whyte, C. 2018. Mind-reading headset lets you google just with your thoughts. *New Scientist*, 4 April 2018. <https://www.newscientist.com/article/mg23731723-300-mind-reading-headset-lets-you-google-just-with-your-thoughts/>.
- Wolpaw, J. R., N. Birbaumer, D. J. McFarland, G. Pfurtscheller, and T. M. Vaughan. 2002. Brain–computer interfaces for communication and control. *Clinical Neurophysiology* 113(6): 767–791. doi: [10.1016/S1388-2457\(02\)00057-3](https://doi.org/10.1016/S1388-2457(02)00057-3).