Modulating Agency:
the Moral & Aesthetic Import of Closed-Loop Deep-Brain Stimulation

Timothy E. Brown

A dissertation

submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

University of Washington

2019

Reading Committee:

Sara L. Goering , Chair

Andrea I Woody

Carina Fourie

Program Authorized to Offer Degree:

Philosophy

University of Washington

Abstract

Modulating Agency:
the Moral & Aesthetic Import of Closed-Loop Deep-Brain Stimulation

Timothy E. Brown

Chair of the Supervisory Committee:
Sara L. Goering
Department of Biochemistry

Deep-Brain Stimulation (DBS) is an FDA-approved treatment for symptoms of motor disorders—with experimental use for psychiatric disorders. DBS, however, causes a variety of side effects. The next generation of DBS systems—Closed-Loop DBS (CL-DBS)—will be able to record users' neural activity in real time and adjust stimulation in order to meet users' needs or demands. Moral philosophers question DBS's influence on users' experiences of authenticity, identity, and autonomous agency. Many of characterizations of (CL-)DBS, however, may not make sense of how DBS complicates, rather than simply impedes or bolsters, users' abilities to exercise agency. This dissertation is a collection of three papers that propose frameworks for understanding how DBS users form relationships with their implants in order to express agency,

how DBS has an impact on the aesthetic dimensions of the experience of agency, and how CL-DBS may prove trustworthy for users without a number of safeguards in place.

Table of Contents

Acknowledgements

Dissertations are the culmination of years of hard work and hardship. But contrary to what bouts of depression and anxiety tell us: we do not work alone, and we do not suffer alone. I could not have completed this dissertation without the kindness, wisdom, patience, understanding, and—crucially—participation of so many friends and colleagues during my graduate career. I'll take this moment to thank as many of them as I can.

First of all, I want to thank my family—especially my mother Carmen and my sister Tugoa—for giving me a perspective from which to speak and the strength to stand behind my convictions.

I'm also indebted to the excellent scholars who comprise UW's Neuroethics group—Michelle Pham, Paul Tubig, Darcy McCusker, Erika Versalovic, Marion Boulicault, Dr. Kate McDuffie, Dr. Laura Specker Sullivan, and Dr. Eran Klein—for squinting at early drafts of each of these papers and trying their best to imagine what each could be.

Through the Center for Neurotechnology's (CNT) generous support, I had the privilege of being able to work with engineers and scientists across a number of fields and on a number of projects. I thank Dr. Howard Chizeck, co-PI of the Biorobotics Laboratory (BRL), for welcoming me into his lab, and being a champion for engagements between humanists and engineers. I thank my labmates within the BRL and colleagues within CNT—Dr. Tamara Bonaci, Dr. Jeff Herron, Dr. Margaret Thompson, Dr. Katherine Pratt, Dr. Nile Wilson, and Dr. Devapratim Sarma—for being co-conspirators and comrades, for their deep commitments to our interdisciplinary work together, and for being the geeky friends I needed when I felt alone.

I also thank all of the students I've mentored who gave me the courage to explore areas of research that seemed strange at first but feel so natural now: Sami Ayele, Hale Soloff, Alex Franke, Mikalia Jackson, Tatianah Summers, Phoenix Meadowlark, and Hannah Martens.

Without Hannah or Phoenix, in particular, many of the themes of this dissertation would not have been as developed as they are now.

I am forever grateful to my committee chair, advisor, and mentor Dr. Sara Goering for investing her time, effort, and wisdom in me. Without her guidance, enthusiasm, and advocacy, I would not have made it to this point. Thank you, Sara, for making space for me to be myself.

Finally, I thank my long-term partner, Sandra Chaidez, for reminding me of who I am, what I stand for, and what's important—not just in intellectual matters, but in life more broadly.

Introduction

Therapies that apply electrical stimulation to the central nervous system are becoming a more common part of medical practice. One such technology, Deep–Brain Stimulation (DBS), is an FDA-approved treatment for the symptoms of Parkinson’s Disease (PD), Essential Tremor (ET), and Epilepsy—with experimental use for psychiatric conditions like Major Depressive Disorder (MDD), Obsessive-Compulsive Disorder (OCD). These systems consist of a signal generator, often implanted in the user’s chest, that sends electrical impulses to electrodes implanted in carefully-selected areas of the user’s brain.

DBS, however, comes with a number of drawbacks. Stimulation can cause side-effects, depending on where the stimulation electrodes are placed, the stimulation parameters used, and the user’s neurophysiology. Some users experience psychological side-effects: for example, a small subset of users with PD develop impulsivity and risk-taking behaviors (i.e., symptoms of mania). Further, these systems use a battery contained in the stimulator, and once the battery is depleted—between two and five years—the entire stimulator unit must be replaced through a minor surgery. DBS may also become less effective over time—DBS users with PD, for example, require increased voltages to receive the same levels of effectiveness.

A new generation of Deep–Brain Stimulation systems offer a possible solution to many of the aforementioned problems. Where the current generation of DBS applies stimulation constantly, in an “open-loop” way, next-generation DBS systems would record the user’s neural activity in real time and use it to adjust the level of the stimulator in order to fit their needs or demands. These closed-loop DBS (CL-DBS) systems could potentially lessen the impact of side-effects, reduce the number of battery-replacement surgeries required, and prevent the user from building a tolerance to stimulation.

Some users, however, report a number of what some commenters call *psychosocial maladaptations*: unwelcomed changes to their sense of self and interpersonal relationships post-op. Several commenters worry that these reports demonstrate how DBS can pose a threat to various features of users' selfhood. Felicitas Kraemer argues that DBS can help users feel authentic to or make them feel alienated from their "true selves." François Baylis argues that personal identity is too dynamic, narrative, & relational for DBS to threaten identity. Instead, DBS poses a threat to user agency, or (more specifically) their ability to contribute to or be the author of their own personal narratives. Catriona Mackenzie and Mary Walker, in line with Baylis, note that DBS may impair autonomy competence—or the abilities and skills users need for deliberation and action.

Views like these, however, characterize DBS as an external force that may strip them of or bestow them with powers of agency. Such views may not be as useful in characterizing more common experiences of living with, making use of, and struggling with DBS—especially as DBS becomes a more common treatment for psychiatric conditions, where psychological and social changes are the goal of treatment. These views, also, may not be as useful in anticipating how users will interact with, or even have difficulty using, next-generation CL-DBS systems. In other words, I worry that many views about the impact of DBS on self-hood are not as capable of making sense of everyday experiences of using DBS—experiences that may also be difficult for users.

I insist, then, that we ask three sets of questions about the experience of using DBS. First, what are the logistics of the relationship between DBS systems and users? How do/can users navigate them? Second, how does DBS change the experience of being an agent? What shapes the experience? Third, and finally, what will happen when CL-DBS is able to record neural data, make inferences from these data, and make decisions about how to adjust stimulation as a result? Will users be able to trust CL-DBS to make decisions on their behalf? This dissertation is comprised of three papers on DBS meant to address the above questions.

In the first, *Building Intricate Partnerships with Neurotechnology: Deep–Brain Stimulation and Relational Agency*, I argue that some DBS users develop relationships with DBS in order to exercise agency. As such, we are better able to understand the issues DBS users face when we home in on the intricacies of their relationship with DBS and the agency they share as a result. At the end, I propose the concept of an "agential competence" as an attempt to codify strategies DBS users (can) employ to build these relationships.

In the second, *Deep–Brain Stimulators and the Aesthetics of Agency*, I try to make sense of users' difficult first-personal experiences of using DBS—of, for example, feeling inauthentic to themselves, or of not knowing if they're the source of their own actions. In *Building Intricate Partnerships* I reprise a number of views that deny that users have an "authentic self" that stimulation can disrupt or uncover. In this piece, however, I argue that we still need a means of taking users seriously when they report feeling inauthentic. On my view, it is useful to describe experiences like these in aesthetic terms, or in terms of how DBS impacts the creative act of being an agent and the user's appreciation of their own experience of being an agent.

Finally, in *Jumping Through Loops: Closed–Loop Control of Deep–Brain Stimulation and Its Impact on User Agency*, I take a closer look at the constituent components of CL-DBS—neural recording, data classification, and stimulation triggering. Some commenters take advantage of the turn of phrase "closed-loop," arguing that the CL-DBS user's position inside or outside the system's control loop will have implications for how they exercise agency and how responsible they are for actions undertaken with DBS' assistance. I instead, argue that the design of each of these components of a closed-loop system can change the user's decisional loop in ways that reflect the assumptions, biases, and decisions made by a variety of stakeholders and institutions. The way each component is designed, I argue, comes with deep implications for whether or not users can *trust* these systems to support their agency.

Building Intricate Partnerships with Neurotechnology: Deep-Brain Stimulation and Relational Agency

Abstract

Deep-Brain Stimulation (DBS) is an FDA-approved treatment for symptoms of motor disorders—with experimental use for psychiatric disorders. DBS, however, causes a variety of side effects. Moral philosophers question DBS's influence on users' experiences of authenticity, identity, and/or autonomy. These characterizations of DBS, however, may not make sense of how DBS complicates, rather than simply impedes or bolsters, users' abilities to exercise agency. Empirical work exploring DBS users' lived-experiences and feminist accounts of relational autonomy demonstrate that the issues users face are better characterized in terms of the user's relationship to their stimulator, that is, in terms of "relational agency."

1. Introduction

Deep-Brain Stimulation (DBS) is an FDA-approved treatment for the symptoms of Parkinson's disease (PD), essential tremor (ET), dystonia, and epilepsy—with experimental use for mood disorders. DBS systems consist of a signal generator, typically implanted in the user's chest, that sends impulses to electrodes implanted in select areas of the user's brain. These signals change the activity of areas of the brain associated with unwanted symptoms. Several research groups have begun trials to use DBS as a treatment for psychiatric disorders (e.g., major depression [Mayberg et al. 2005; Mayberg 2008; Mayberg et al.2016; Widge et al. 2015] and obsessive-compulsive disorder [Alonso et al. 2015]). DBS, however, comes with a number of physiological and psychological side effects. People using DBS to manage motor disorders, for example, may experience paresthesias (tingling sensations) in their limbs, weak or slurred speech, or feelings of euphoria (Kuncel et al. 2006; Smeding et al. 2007; Hariz et al. 2008).

Some users also report more severe neuropsychiatric problems post-surgery (Voon et al. 2006). A small subset of people with PD develop impulsivity and risk-taking behaviors (e.g., gambling) as a side effect of stimulation of the subthalamic nucleus (Smeding et al. 2007). Changes like these can result in lifestyle and social difficulties: an altered self-image, feelings of self-estrangement, and a lack of motivation (Schüpbach et al. 2006; Agid et al. 2006; Kraemer 2011b). In one study, people using DBS for PD reported a variety of self-impressions, including that DBS makes them "feel like a 'Robocop'" or like they're "forced to live in a body that's out of control" (Agid et al. 2006, 412). Others noticed changes in their interpersonal relationships. One woman acting as a caretaker for her husband with PD explains, "since the operation, [my husband] wants to live like a young man … I prefer him as he was before, always nice and quiet" (412). The nature of and appropriate response to these so-termed "psychosocial" changes is a matter of on-going debate (Mecacci and Haselager 2014; Haan et al. 2017), but a recent pilot study suggests that concurrent education sessions may help reduce some of the anxiety and depression in some patients (Dos Santos et al. 2017).

Some commentators frame these psychosocial changes in terms of DBS's effects on selfhood: DBS may make the user feel inauthentic (Kraemer 2011b, 2013), shift important features of their identity (Glannon 2009), or diminish their autonomy (Mackenzie and Walker 2015). These views characterize DBS as an external force that can confound or bolster these features of selfhood. Such views may, however, have difficulty parsing cases where users have trouble distinguishing between their own behaviors and those behaviors influenced or instigated by stimulation. For example, some people using DBS for depression and obsessive-compulsive disorder have reported problems determining if their actions are "their own" or if they are the result of stimulation's side effects. In moments like these, I maintain, DBS users experience confusion with respect to their sense and exercise of agency. In fact, some DBS users are actually put in a position where they develop intricate partnerships with DBS in order to exercise agency.

That is not to say we should not characterize these negative outcomes as cases of felt-inauthenticity, identity change, or autonomy encroachment—these modes of explaining DBS users' experiences are all useful in certain contexts, and users' experiences vary widely enough that it is possible that they benefit from multiple modes of description at once. My view, however, is that we are better able to anticipate and address the issues DBS users face, such as the aforementioned feelings of ambiguity, when we home in on the intricacies of the users' relationships to their stimulators and the agency they share as a result. The ability to make sense of these partnerships will be vital as DBS becomes a more common treatment for a wider variety of conditions and as advanced, artificially-intelligent DBS systems become more commonplace.

I begin by considering the aforementioned three views of how DBS changes selfhood and identifying the potential gaps in them. I propose that each of these views fall short of capturing users' experiences of struggling against or alongside their neurostimulators. These experiences, I maintain, are the product of a relationship some users form with their DBS systems. I present two cases where DBS users run into difficulties in these relationships—the aforementioned agential ambiguities people navigate when they use DBS for psychiatric disorders—and the tradeoffs people negotiate when they use DBS for motor disorders. I suggest that we can make sense of these cases by making use of a notion of "relational agency"—according to which DBS users' exercises of agency may be supported, encumbered, and otherwise "intertwined with" their DBS systems (Goering et al. 2017, 67). Finally, I propose the concept of an "agential competence"—building on Diana Meyers's notion of autonomy competences—as a first attempt to codify strategies DBS users employ to "make do."

2. Three views assessing the psychosocial effects of DBS

One way to describe the psychological and social problems some DBS users experience centers users' feelings of inauthenticity or alienation (Kraemer 2011a). Although people with conditions like PD and ET may encounter profound feelings of alienation or inauthenticity because of their conditions, Kraemer argues that DBS users may, similarly, feel authentic or inauthentic, alienated or more at home with themselves, as a result of using DBS. On her view, the experience of using DBS is a complicated process whereby users come to grips with how authentic (or inauthentic) and at home (or alienated) they feel using DBS around their friends and family. To understand why some people have trouble adjusting to life with DBS, it is important attend to how each user undertakes a process of "recognizing, exploring, and enacting what they regarded as their 'true selves'" (Kraemer 2011b, 496).

Take, for example, a person with a severe tremor who finds it difficult to eat without making a mess. Such a person might feel embarrassed eating at a public restaurant, and (as a result) they may feel less authentic to themselves than they did before the onset of their disorder. They may, further, feel alienated from their dinner partners and from onlookers. For people who feel this way, DBS may seem like a means of recovering lost ground if it successfully eliminates the tremor. But, if the DBS user experiences psychosocial side effects (e.g., increased impulsivity), they may feel even more alien or unrecognizable to themselves post-op than they did pre-op. In her analysis of the now-canonical Schüpbach cases (Agid et al. 2006; Schüpbach et al. 2006), Kraemer (2011b) argues that some feelings of alienation are signs of progress (e.g., when a person feels alienated from their life with a movement disorder) and some feelings of authenticity cause problems (e.g., when DBS causes changes that users embrace but their loved ones are unable to accept).

While the language of authenticity might be useful for capturing salient features of DBS users' testimonies, it does suggest a problematic metaphysics of authentic selfhood. That is, if we take people literally when they say they feel "less like themselves," we may also (perhaps tacitly) accept that they feel authentic to or alienated from some predetermined "self." There are, of course, reasons to be skeptical of this claim: to say that a person has an "authentic self" is to draw a boundary around a set of attitudes, behaviors, and actions that constitutes a person's self, but it would be difficult to draw such a boundary given how these behaviors change over time and across contexts. We accept, for instance, that we can act in completely different ways around different groups of people—who we are around our families is often different from who we are around our long-term friends. We would be hard-pressed to say which of these "selves" is the "real" one—that is, insofar as one "self" takes precedence over all the others.[1]

Similarly, Françoise Baylis (2011) argues that DBS is not a threat to identity in the ways indicated by the authenticity view. Baylis, instead, takes on a narrative view of identity—where a person's identity consists of the on-going story they construct about themselves. Self-narrative, she argues following Schechtman (2010), is too dynamic to talk about in terms of authenticity: who a person is changes with their lived experiences, the narratives they take to be constitutive of their identity, and the constraints placed on their self-narratives through their relationships with others. To this last point, Baylis asks us to consider the tension between how we describe

---

[1] Perhaps there are even metaphysical and (by extension) epistemological difficulties when we presume that we have selves that are unified enough to call "authentic" in the first place. Specker Sullivan (2019) explores this possibility and its implications.

ourselves and how others describe us: we cannot describe ourselves any way we want, and others cannot describe us any way they want. There is, in Baylis's terms, an "equilibrium constraint" on our self-narratives, whereby we negotiate with others to arrive at descriptions of ourselves that all parties involved can accept. "Relational identity," then, "is a dialectical process aimed at achieving equilibrium—some kind of temporary and temporizing balance between self-ascription and ascription by others" (Baylis 2011, 518). We might imagine that Baylis asks: what sense does it make to say a person has an "authentic self" when selves are so dynamic and constantly up for negotiation?

As such, Baylis takes a hard stance against the authenticity view: identity simply does not work this way, and to operate under the assumption that it does is to miss the details of how self-ascription works in its context: "it makes no sense to describe any particular event or experience (including DBS) as threatening. There is no pre-set autobiographical narrative (i.e., life story) that has been thwarted by some unfortunate event or experience. There just is the life story as it unfolds" (Baylis 2011, 516). That is, when DBS causes a change that makes its user feel out of sorts, DBS does not shift the user away from their authentic self. Instead, DBS simply causes yet another change in the user's ongoing autobiographical narrative. The DBS user may feel uncertain or upset about the self-narratives available to them, unsatisfied with the narratives their friends and family produce on their behalf, or uncomfortable "performing" any of these self-narratives (Baylis 2012). DBS implantation, Baylis argues, is not unlike other events (e.g., accidents, relationship loss, etc.) that have an impact on the stories we tell about ourselves.

There are cases, however, where Baylis admits that DBS constitutes a particular kind of threat to its user's identity on her view: "when DBS undermines agency to such an extent that the person is no longer able to meaningfully contribute to the authoring of her own life (i.e., to contribute to the cyclical and iterative process of projecting, defending and revising a self-narrative). Without the ability to contribute to the process of identity formation, a person loses the ability to hold on to a sense of self" (Baylis 2011, 525). It is one thing, Baylis argues, for DBS to cause unwanted changes to a person that are then reflected in their self-narratives; it is another entirely for DBS to undermine a person's ability to make choices or form a self-narrative. Think, for example, of cases where people using DBS for PD develop gambling or risk taking behaviors. In such cases, users are rendered less able to make their own choices—they lose some of their control over their own behaviors, even as the system gives them relief from some symptoms of their movement disorder. These encroachments on the user's agency are the problem, not the mere fact that their self-narratives have changed in response to some important event. "When direct brain manipulation explains a belief or behavior there is reason to think of this as a serious threat to agency," she summarizes, "which in some instances may give rise to a threat to identity" (Baylis 2011, 524). As such, she concludes, "[these] disruptions [are] best examined through the *lens of agency*" (Baylis 2011, 514, emphasis added).[2]

Mackenzie and Walker (2015) extend Baylis's argument against the authenticity-view. They argue that "the appeal to authenticity … blurs the distinction between narrative identity and autonomy" (386): not only does the authenticity view obscure details about how DBS users (re)form, negotiate, and perform their identities, authenticity-talk also obscures the details of how DBS can impair or bolster the user's ability to act autonomously. To illustrate, they appeal to Walter Glannon's (2009) analysis of the aforementioned "Dutch patient," who faces a choice between keeping his DBS system turned on but needing to stay in a psychiatric facility to monitor stimulation-induced mania, or turning the system off but needing to stay in a nursing

---

[2] Also see Goddard (2017), who also calls for framing the users' experiences of DBS through "the lens of agency."

facility to receive care given his severely limited mobility. Glannon asks, “How much disruption can one’s life narrative accommodate without threatening the integrity of the whole?” (291). Mackenzie and Walker (2015), in response, worry that “narrative reinvention of the self may not be possible” (382) when DBS “[disrupts] a person’s autonomy competence to such an extent that he is unable to engage in narrative self-revision” (390). The salient issue for DBS—on their view—is autonomy competency.

That is to say, Mackenzie and Walker (2015) adopt and apply a relational, competency view of autonomy, on which autonomous agency “is a competence involving a complex repertoire or suite of reflective skills” (387). Among these skills are: the introspective skills necessary for knowing what we feel, need, and desire; imaginative skills that allow us to foresee possibilities for ourselves; communicative skills that make it possible to seek out and make sense of other people’s testimonies and advice; volitional skills required to make decisions and follow through on them; and the interpersonal skills that help us form and sustain relationships (Meyers 2000, 166). On this view, in order to be autonomous we “must be able to pose and answer the question ‘What do I really want, need, care about, value, etcetera?’; [we] must be able to act on the answer; and we must be able to correct [ourselves] when [we] get the answer wrong” (Meyers 1989, 52-53). We cultivate our competencies in order to answer to this question for ourselves and act upon it. Further, crucially, our autonomy skills are developed, maintained, and exercised in the context of significant interpersonal relationships (Mackenzie and Walker 2015, 388). As Mackenzie and Walker (2015) apply this account, DBS is a threat insofar as it disrupts the user’s autonomy competencies, rendering them less able (or unable to) employ these skills. In contrasting cases, when treatment with DBS is successful, DBS can help users regain the “volitional, emotional, motivational, imaginative, and critically reflective capacities necessary for autonomous deliberation and action" that were impaired by physical and mental illness (Mackenzie and Walker 2015, 390).

Some of Mackenzie and Walker's (2015) conclusions, however, seem out of step with the lived-experiences of the DBS users they consider. Is the Dutch patient well and truly "unable to engage in narrative self-revision" (390)? After all, given the choice between severe mania upon receiving DBS and debilitating Parkinsonian tremor without DBS, he was able to make the choice: he chose to use DBS, despite the onset of severe mania. That is, he chose to bring about a future for himself free of severe tremor, and we should not discount the importance or impact of that choice for him. Further, we cannot presume this is the end of the Dutch patient's ability to exercise agency. Encapsulated within these conclusions about the Dutch patient's decision—that he "sacrificed his mental competence and autonomy" (Kraemer 2013, 758), that he "is no longer able to contribute meaningfully to the authoring of [his] life" (Baylis 2011, 525; also quoted in Mackenzie and Walker 2015, 382)—is a misconception about the lives of people with psychiatric disorders. We should ask: does the onset of a person's mania mark the end of their ability to "contribute meaningfully to the authoring of [their] life" (Baylis 2011, 525) through their capacity to act as agents (Mackenzie and Walker 2015, 390)? I argue that it does not. Instead, it marks the beginning of psychiatric treatment, more choices, and more decisions.

I agree with the move to frame cases like these in terms of DBS's impact on the competencies required for autonomous agency, but it is not clear that even the more severe side effects in these cases are so severe that autonomous agency is impaired or no longer possible. In order to be autonomous, we must be able to reflect on the options available to us, make a voluntary decision about what we want to do, follow through on that decision, and reflect on the consequences. For the Dutch patient, these competencies were made much more difficult to exercise: his mania might have constrained his autonomous decision-making abilities, but PD placed significant constraints on what he could choose. Even with these limitations, however, he retains the capacities necessary for agency—a person who intends and acts, even within significant constraints. In a way, it is less apt to say that the Dutch patient is *impaired* with respect to his autonomy competency: even the most competent of us would have a difficult time making the same choice. I argue it is more apt to say that the Dutch patient's competencies are *complicated* by DBS in a way that makes agency a struggle. We certainly ought to be able to make sense of how DBS impacts a user's ability to do the work required of being an agent, but we also ought to be able to make sense of how DBS users are able to make use of or work around their neurostimulators in order to take on the work agency requires. In the next section, I'll present two cases that showcase several salient features of users' struggles with DBS.

3. The struggle of using DBS

A common conception of agency might maintain that a person expresses agency when their actions are initiated by the right mental states (Schlosser 2015). A more robust formulation of this conception recognizes that exercising agency is itself a skill—one cultivated through practice with and with help from others (Meyers 1989, 2000). That is to say, it can take work to figure out how to form an intention, work up a plan, and carry it out. After all, there are better and worse ways to decide between competing intentions, form plans, and drum up the motivation to act. Sometimes exercising agency is a struggle in and of itself, even if we table the possibility that some external forces can threaten autonomy. Motor and psychiatric disorders complicate that struggle, and DBS complicates that struggle further when it changes how its users are inclined to form intentions and act on them. The Dutch patient's experience is an extreme example of how DBS impacts agency—insofar as stimulation-induced mania changed how he formed intentions, insofar as his severe PD symptoms limited his available choices, and insofar as the juxtaposition of the two states was a challenge unto itself. DBS, however, can complicate even the average user's life: users will still need to choose stimulation settings, decide if and when to turn the stimulator on or off, determine if there are side effects, and so on.

Each of the views we've considered so far characterizes parts of the struggle DBS users face: when they feel inauthentic, are unable to contribute to their self-narratives, or are unable to make use of or further develop their autonomy competencies. Both Baylis (2011) and Mackenzie (2015), further, underscore the role interpersonal relationships play in these struggles. What is missing, I maintain, is an account of the mechanics of the relationships users form with the DBS system itself as they struggle against, alongside, and in tandem with it. One crucial feature of this relationship lies in how users come to rely on DBS in order to implement their intentions—despite how the system may not support the user as well in some contexts. Similarly, we rely on other people to support us agentially, even if it is possible for them to fail us. As a poignant illustration, Alisa Bierria (2014) argues that we rely on others to recognize the intentions behind our actions, even if they do not agree with or endorse them—"their recognition of our intention becomes a part of the social choreography of agency" (131). When others in our context do not recognize or validate our intentions, they may misinterpret our actions and constrain our actions as a result. But even under external constraints, it is still possible to exercise agency, and we are able to recognize these exercises when we take up "a reimagining of agency that foregrounds the experience of 'struggle' when trying to accomplish an end and emphasizes the role of improvisation and making do when there is a breakdown of an agentic plan" (Bierria 2014, 137). Bierria, here, borrows this last insight from Susan Wendell's (1989) exploration of what feminist theory can learn from the lived-experiences of people with disabilities, and she does so in order to "trouble an idealized notion of agency that hinges on efficacy" (Bierria 2014, 142–43).

Using DBS (in some cases) requires just this kind of improvisation and making do under the constraints of systems users would like to be able to rely on. That is, when DBS does not work the way the user expects and there is a breakdown of their agentic plan, the user will need to improvise (perhaps devising a way to work around the problem) if they want to carry out an agentic plan at all. But, where Bierria retreats from the relational for the sake of making sense of how agency is possible without uptake from (or in defiance of) oppressive others, I embrace the relational for the sake of making sense of the lived-experiences of people using neuromodulation. Instead of asking if (and how) DBS poses a threat to the user's agency, I ask: what are the mechanics of DBS users' confusions, negotiations, and collaborations with their neurostimulators?—and how do users express agency throughout? For the remainder of this section, I'll present two of the many possible challenges DBS users could face as they struggle with, against, and alongside their stimulators. For the first, users find it difficult to determine the source of their behaviors. For the second, users must negotiate make the tradeoff between DBS's side effects and the effects of the disorder treated.

3.1 Agency ambiguities and psychiatric disorders

In a focus group conducted with people participating in a research study focused on using DBS to treat their treatment-resistant clinical depression, Klein et al. (2016) found that a number of participants expressed worries over what the authors termed *ambiguous agency*—or the possibility "that individuals with DBS might lose track of their sense of agency if it became unclear whether they or the device was responsible for their feeling or acting in a certain way" (4). One participant describes it:

> I've begun to wonder what's me and what's the depression, and what's the stimulator. I mean, for example, I can be fine, and then all of a sudden … I'll say something that is insensitive or just misread a person entirely, say something that either makes ME look like a fool, or, hurts them, or, something along that line. I can't really tell the difference. There are three things—there's me, as I was, or think I was; and there's the depression, and then there's depression, and the device and … it blurs to the point where I'm not sure, frankly, who I am. (Goering et al. 2017, 63)

Another participant framed this concern in terms of how others think of their DBS use, "So, there's people in my family that sometimes question … how much of it is me any more and how much of me is … bein' programmed. So that's been a … that's a hard thing to deal with sometimes, but, you know, like I said, I'll take it over what I had" (Goering et al. 2017, 5). A third participant tried to imagine what their behaviors would have been without DBS, "There are parts of this where you just wonder how much is *you* any more, and you wonder kind of, 'How much of it is my thought pattern? How would I deal with this if I didn't have the stimulation system?' You kind of feel artificial" (Goering et al. 2017, 5). In a later analysis of these focus group data, Goering et al. (2017) report one user's difficulty determining why she could not cry at a funeral: her life's circumstances, her drug prescriptions, or her stimulator's settings (63).

In a series of interviews of people using DBS for OCD, Haan et al. (2017) found that a number of participants "had to get used to themselves, their new lives, and their new behaviour" (14). Many of their experiences share much in common with the participants in Klein et al.'s study. According to De Haan et al. (2013), "the majority of the participants (12) had been thinking about the question of how to determine whether the changes that accompanied the DBS treatment were side-effects or rather reflections that they were becoming more themselves" (17). One interviewee remarks on how their interpersonal behaviors have changed: "From always making up excuses, never being able to keep appointments …. And then I just say: 'Yes, sounds fun'—and then I think 'Is that me?' A bit strange, really, I need to get used to myself too" (14). Another reflects on how they've lost friends after using their implant, "I think that I just, yes, stand up for myself more. And for some people, that is just not comfortable; the relationship [with them] just changes because of that. … What I also have, and that also bothers me a lot, is that my interest in the people around me has decreased significantly. … I am still not used to it" (15).

We might characterize these users' remarks as indications that DBS changed their *sense of agency*—"the sense that [they are] the one who is causing or generating an action" (Gallagher 2000, 15). This sense of agency typically includes and understanding of ownership or "'mineness' for one's body and actions" (Gallagher 2018). The impact of DBS on the user's sense of agency might be even more profound so as to disrupt or supplant the user's agency when they act in ways they cannot claim as their own (Lipsman and Glannon 2013). We must remember, however, that if a person uses DBS to treat their depression or OCD, then they necessarily mean to change features of their agency—they, after all, want to act in a way less impeded by their depressive episodes or their obsessions and compulsions. That is, to treat a person's psychiatric disorder is to prevent unwanted intentions and actions—it is an intervention on their agency. And, if we start from an enactive, embodied perspective, as de Hann et al. (2013) do, we start by considering "the whole complex system of a person interacting with the landscape of affordances" (12). These systems, if they work, are meant to change crucial affordances available to users (Haan et al. 2013, 2017; Gallagher 2018). But those changes can only happen in light of substantial changes to how they exercise their agency.

Even still, such users, as the participants in Klein et al.'s (2016) and De Haan et al.'s (2017) respective studies demonstrate, may not feel they are the sole source of their own actions, even if they previously intended to bring about and currently endorse these changes. Even if the user's new behavior is precisely what they wanted, currently want, and need for their own continued survival, the very fact that neuromodulation is the cause of their personality changes may challenge their beliefs about the relationship between their brains and their behaviors (Mecacci and Haselager 2014), or change the way they approach situations. As such, these feelings of ambiguous agency could have an effect on users' actual exercises of agency in several ways. Users may be less inclined to hold and act on intentions if they cannot be sure they are the source of those intentions. Further, they may be less willing to take responsibility for actions resulting from those intentions if they don't feel that their actions are their own (Klaming and Haselager 2013; Sharp and Wasserman 2016). But lastly, these ambiguities make it difficult to determine what to do next.

To underscore the problem raised at the end of the last section: where DBS gives the user the ability to act in ways their psychiatric disorder would not allow, the salient question may not be, "Does this system threaten the user's ability to contribute to their identity-constituting narratives?" The answer to this question may very well be "yes" for as long as the user experiences these difficulties. But such DBS users are not rendered powerless. To the contrary, the experience of agential ambiguity is a challenge to the user's agency—a challenge they must address in order to see therapeutic benefits. The woman who found herself unable to cry at a funeral also found it difficult to determine what kept her from crying and what to do about it. She is not rendered powerless, but her agential power is made complex through neurostimulation. Even if the user would "take it over what [they] had" (Klein et al. 2016, 5), these complications to users' sense of agency make up a terrain that takes skill to navigate.

3.2 The tradeoffs of DBS for Parkinson's disease

People with motor disorders like PD face analogous complications to their agency, but their struggle might more often require making sense of how DBS limits or extends their control over their bodies. Consider the case of Helmut Dubiel, a professor of philosophy who lived with

PD and used a DBS system to manage his symptoms. Stimulation came with a variety of side-effects: a severe speech impairment, exhaustion, and depression. He writes: “My worst post-operative symptom, which remains unchanged to this day, is a speech disturbance: my volume is too low, and my articulation is poor, slurred” (Dubiel 2009, 94). This side-effect made it difficult for Dubiel to work as a professor—to give lectures or speak with others. Further, he tried to keep his diagnosis private—it was a secret that he believed would, at least in part, “discredit [him] in the eyes of the world” if revealed (94). But PD is a degenerative disease; its symptoms worsen over time: he could only conceal those symptoms for so long before they became too apparent. DBS might have given him another means of managing (and thus concealing) his symptoms, but it also complicated Dubiel’s already-complicated expressions of agency.

Dubiel was only able to get a handle on his stimulation-induced speech impairment once he sought out the help of a neurologist and representative from the device’s manufacturer. They determined that stimulation was, indeed, the cause of his difficulties speaking. “It was as if I were channeling a spirit,” he says, describing the moment the DBS was turned off, “that very second my voice returned, sonorous and clearly enunciated, only slightly hoarse” (Dubiel 2009, 118). He received a *patient programmer*, a handheld device that allowed him to turn his device on and off. This, then, gave him the ability to choose between treating his motor symptoms and speaking without an impediment.

As such, Dubiel could exercise his agency and make choices about his treatment, even as the device constrained his available choices. Baylis (2011), reflecting on her conversations with Dubiel, underscores the system’s positive impact on Dubiel’s outlook: “While having to choose between talking and walking is less than ideal,” she writes, “Dubiel’s restored ability (within limited parameters) to control his body in social situations restores his pride and professional authority" (521). She characterizes Dubiel's struggle with his DBS (and with PD in general) in terms of his ability to construct his identity-constituting narratives given the constraints of PD. "Dubiel is the person at the intersection of who he wants to be, and who others will minimally let him be," she concludes (521). This description, while useful, leaves us with a blind spot. DBS itself is also constraint on who Dubiel wants to be. As such, he is also what his DBS will allow him to be through his successive attempts to navigate the device's settings, their impact on his behaviors, his (and others') reflections on the device, and so on.

4. Neurostimulators as relational agents

The views we considered earlier could lead us to say that neurostimulators are inert devices that can cause changes in users that confound or bolster key features of their selfhood. DBS is considered a threat when it undermines a feature of the self the user wants to preserve. The examples we have considered, however, suggest that interactions between DBS and their users are more negotiative and collaborative than previously acknowledged. They frame users' interactions with their stimulations in relational terms. We only need fill out the details of this relationship.

Let's grant, as Baylis (2011) does, that our narrative identities are constituted relationally through an "iterative process involving 'self'-perception, 'self'-projection, 'other'- perception, and 'other'-reaction" (514). Further, let's grant, as Mackenzie and Walker (2015) do (following Meyers [1989]), that our autonomy competencies are cultivated in part by and through our relationships with others. It is not a stretch to say that parts of ourselves are cultivated through our interactions with technology, constituted through our iterative relations to devices.[3] Neither

---

[3] Others, Andy Clark (2008) and Andy Clark and David Chalmers (1998) most notably, go further to say that cognition can extend into the technologies we use, if those technologies stand in the right relationship to us.

is it a stretch to say that our autonomy competencies are cultivated by and through our relationships with devices. We need only give a more robust framework for thinking about the details of these relationships.

First, consider how relationships work between humans given how we create identities that we enact on the world. Hilde Lindemann (2014), in her *Holding and Letting Go*, defends a relational view of identity that underscores how others (family and friends most often) play an active role in constructing our identities as well as with supporting and resisting our attempts to express them. This process of "ordinary identity work" (94) involves a kind of performance, making use of a shared corpus of narratives held about one another as the springboard for improvised action. We build as well as enact our identities socially. "It's an act," she maintains, "a kind of theatrical performance" (98) wherein performers use their knowledge of themselves, their fellow actors, and the social context to determine how to express themselves. Further, stage actors perform for audiences, and our expressions of personhood are similar: "as we express ourselves, others will in turn have to be impressed by us in some way" (98). Even further, each performer relies on their fellow performers to respond appropriately: "[other actors'] job is to perform the 'response' moment of personhood, which they do by taking up the dual role of audience and performer" (99). By following the rules recognized within the social context the actors are in, and acknowledging when their fellow actors act perform appropriately, each actor performs a normative role. We hold each other in our respective identities as we perform (or express) them together in a variety of contexts.

We must remember, however, that performance is difficult, and sometimes our fellow actors could do a better job of giving us the uptake we require to perform our identities well. Bierria (2014) demonstrates how our fellow actors can oppress us through their misinterpretations of our actions (and the intentions behind them). Lindemann (2014), similarly, notes that when our fellow actors are inexperienced or unfamiliar with us, we may have to "engage in what has been called the 'hidden labor' of manipulating the others into giving her the proper response" (101). If a more experienced actor is paired with a less experienced one, they may need to "carry" the performance for their partner—by making it obvious where their partner should be in the script, exaggerating their performance so that their partner emotes appropriately, and so on. She gives the example of Jackie Leach Scully's (2010) experiences traveling with a hearing impairment: "she has to 'perform deafness'—for instance, by cupping her hand behind her ear and looking quizzical—if she's to convince flight attendants that she needs special notification of in-flight announcements" (101). Scully had to exaggerate her own performance to pick up the slack from of the flight attendant—who is not accustomed to performing with people with hearing impairments.

Similarly, DBS users may need to "pick up the slack" for their DBS systems when they are not optimized. Dubiel (2009), for instance, wanted the opposite of what Scully did: he wanted to keep his symptoms hidden from view. He found this difficult to do not only because of the expectations of his friends, family, students, and colleagues. He also struggled against, with, and alongside his Parkinsonian symptoms (first and foremost), his neurostimulator, and the medical professionals who were responsible for maintaining his neurostimulator system. The outcome of Dubiel's attempts to conceal his disease was contingent on his ability to improvise over his Parkinsonian symptoms, negotiate with the device's manufacturers, and control the neurostimulator's settings (once given the means to). Dubiel had to *carry the performance* despite his stimulator's inefficiency. Similarly, the participants in both Klein et al.'s (2016) and De Haan et al.'s (2017) respective studies describe being uncertain about being the source of their own actions, needing to get used to who they are while using DBS, and having to convince others to accept their new behavior. If DBS systems play some role in the mechanisms of their users' ability to form intentions and carry them out—and surely, in the case of psychiatric disorders, they must—what role can we say DBS plays in how they perform their identity? Is DBS a prop or a character?

We can begin to answer to these questions if we attend to the possible status conferred upon DBS systems through users' interactions with them. Although the DBS systems may appear inert, perhaps they gain a kind of conferred animacy, given how it intimately they interact with the users in ways that impact their agency. In her "Vital Wheels: Disability, Relationality, and the Queer Animacy of Vibrant Things," Julia Watts Belser (2016) argues that wheelchairs—thought normally to be "dead matter" by the able-bodied—take on a kind of "animacy" through interactions with users. On her view, "wheelchairs occupy an intimate space: simultaneously an expression of our own body-forms and a distinct, independent entity—a vital, vibrant being" (7). That is, the wheelchair is not just dead weight pushed around by some wheeler; rather, the wheelchair takes on the wheeler's (expressive) movements in some moments, resists movement in others, and sometimes seems to have a character of its own. For Belser, then, wheelers and wheelchairs can form an "*intricate partnership*," or a "mutuality between wheeler and chair [that] demands an understanding of how hand and tire interrelate, how weight affects wheel. Motion and presence are unified; the self is expressed in and through wheels" (12, emphasis added). Given the structure of the relationship between wheelers and their wheelchairs, Belser insists it makes more sense to say that the wheeler and their chair act together than it does to say the wheeler pushes their chair.

Belser, however, more often reaches for examples of competent wheelers who contend with and relish the quirks of their wheelchairs. By acknowledging the animacy people with disabilities confer on their wheelchairs, she elevates the status of both people with disabilities and the tools they rely on above that of dead weight. This move preemptively counteracts prevailing ableist attitudes that would characterize people with disabilities as with significantly diminished agency or without agency at all. By attending carefully to how disabled people express their agency in partnership, we are able to understand more clearly how thoroughly relational and socially embedded agency is writ large. In this way, Belser's project lines up with Bierria's: both problematize oppressive misinterpretations of agency by showcasing oppressed peoples' powers as agents. I want, however, to recognize those moments when we use our technology in a clumsy or awkward way so that we can also attend to how we become competent or graceful through an iterative process of cultivating the skills necessary (i.e. for gaining agency competence).

DBS and their users, I contend, occupy a similar shared space in moments of struggle. We can see this most clearly in the moments where users try to figure out what role DBS will play their lives and what role they want DBS to play. Goering et al. (2017) argue that we should think of these systems as relational others in just this way—as systems designed to supplement or support an individual's agency the way a friend or family member would. Just as friends and family enforce the reality constraint on our identity–forming narratives (Baylis 2011), related others also hold us in identities during our performances of them (Lindemann 2014). Related others hold a great deal of power on this view: the right other, in the right position, can often notice changes in me that I cannot see, and they can (perhaps) influence me in ways I am unable to perceive. As we have seen, DBS can influence users similarly: producing changes at a neurological level that are sometimes easy to notice (e.g., when it treats tremor) and sometimes are not (e.g., when it induces mood changes). And just as it is often necessary to set boundaries with controlling friends and family members, it might be necessary for the DBS user to set goals for and boundaries on how their stimulator impacts their agency. To make sense of this boundary-setting, we need to take seriously the possibility of relational forms of agency where "an individual's agency (connecting her motives and actions) can be supported through and entwined with a variety of external others" (Goering et al. 2017, 66)—even if those external others are not so external to us, even when they are, in fact, devices implanted in our bodies.

In order to understand users' lived-experiences of coming to grips with the realities of stimulation, we must attend to the "mutuality" between neurostimators and the stimulated. Many features of life with DBS are salient here: the subcutaneous space the leads and stimulator unit occupy, the maintenance required to prolong battery life or ensure the system's efficacy, the physiological and psychological effects (intended or otherwise), the need to negotiate tradeoffs or "workarounds" to problems (should they arise), how family and friends think about stimulation, the option to change stimulator settings, one's reflections on life with such a system, and so on (Brown et al. 2016). User experiences will vary, of course, but we can say at least this much: many of these experiences require "getting used to," while others will require sustained negotiation to figure out how to live and work with DBS. Dubiel's negotiations were iterative and ongoing as he learned how his stimulator's voltage levels match up with his symptoms and social context. A user navigating "agency ambiguities" has work to do as well: learning what role their stimulator plays in their behaviors and determining what role their stimulator will play in future behaviors.

Earlier, we asked what role DBS plays in users' performances of identity: are they props or fellow agents? Our best answer, for the moment, is that they are somewhere in between. Users confer animacy on their stimulators through their attempts to express themselves with the help of stimulation. When difficulties arise, users must negotiate and collaborate with their stimulators to overcome them. In the future, however, we may have grounds to call neuromodulators full-fledged agents. Currently, neuromodulation systems are configured manually by technicians, neurologists, and (in a growing number of cases) users. Several groups, however, are testing artificially intelligent DBS systems capable of adjusting themselves. Such systems would record and analyze data recorded from sensors implanted in contact with a user's brain in order to determine when to (or when not to) apply stimulation (Hebb et al. 2014; Klein and Nam 2016; Goering et al. 2017). Such a system could present users with even more difficulties. Imagine, for instance, if Dubiel's stimulator were able to detect when he wanted to speak or when he wanted to walk. Such a system could have been convenient—insofar as Dubiel would no longer need to use explicit commands to control his stimulator. It might have also have been more difficult—insofar as Dubiel might have had trouble predicting when the system would switch states, determining what the current state of the system is, and so on. These systems, we might say, act on their own. As such, we might be inclined to say that such systems are not just conferred animacy but, instead, are much more agents in their own right (Klein et al. 2016; Kellmeyer et al. 2016; Goering et al. 2017; Yuste et al. 2017). How often will DBS users need to "carry the performance" for their artificially-intelligent neurostimulators, or will they be "held" too tightly by them?

In summary, I've suggested that DBS is capable of more than just bolstering or threatening user agency. On my view, users are capable of conferring animacy on their DBS systems, and DBS systems are capable of holding their users to certain performances of agency. As such, users form an intricate relationship with DBS that can be challenging and require skill to navigate in difficult circumstances. But what skills do users require to do well in these relationships, and how do they cultivate these skills? In the next section, I explore this question and identify skills users might train as they build relationships with DBS.

5. Agentic skills for relational agency

Earlier on, I suggested that we should think not only about DBS's impact on autonomous agency but also its impact on agency broadly. I continued to say that agency itself a kind of struggle—a struggle made more complicated by DBS and the relationship formed with it. Both Baylis (2011) and Mackenzie (2015) remind us that our autonomy competencies are cultivated and exercised relationally. But, if my proposal above is correct, the DBS user's answer to the question "what do I truly want?" is, in some cases, the product of the collaborative partnership between then and their stimulators. Below, I'll gesture toward three skills that we might be able to identify as users navigate their partnerships with DBS: *introspective vigilance*, *improvisation*, and *relational resourcefulness*.

Some of the difficulties with DBS are a result of the user not being able to determine where their agency ends and the DBS's intervention begins. The case of the woman who puzzled over why she could not cry at a funeral is a clear example of this dynamic. Her worries themselves, however, are an example of an *introspective vigilance* where she is "attuned to signs of frustration and dissatisfaction, attentive to baffling subjective anomalies, and willing to puzzle out gaps in one's self-understanding" (Meyers 2000, 167)—this could be a useful strategy for making sense of her relationship to DBS. That is, her worries about her behavior, her behavior's possible causes, and what she ought to do in response, are apt and crucial to taking control of her experiences. Being aware of the change in her behavior and knowing that it could be related to stimulation is an important first step. This kind of vigilance motivated Dubiel (2009) to insist that his DBS was causing many of his cognitive and speech issues, and he was eventually able to gain control of his stimulator enough to confirm his suspicions and tune his stimulation to match his needs.

In moments when DBS performs in a clumsy way[4] and when the user's agentic plan does not pan out, *improvisation skills* can help smooth the performance over or devise solutions to problems. Recall Dubiel's attempts to conceal his PD symptoms. In moments when his stimulator induced a speech impediment, he chose words he could say despite his impediment in order to perform agency in ways that are more in line with his intentions.[5] Further, the woman who could not cry at a funeral might need to improvise a way to express the emotions she wants to express and to mourn the way she wants to mourn.

Finally, cultivating a kind of *relational resourcefulness* can help the user situate DBS alongside the other resources they can use to support their expressions of agency. Recall the problems DBS users' friends and family members face as they grow accustomed to how DBS changes the user's moods and behaviors. One study of people using DBS for PD found evidence suggesting that friends and family notice personality changes more readily than the users themselves (Pham et al. 2015). This suggests that users may need to rely on friends and family to inform them of these changes. Some users may not want this kind of feedback—for example,

---

[4] I borrow this turn of phrase from Lindemann (2009), where she describes how people with dementia hold their friends and family—dementia makes holding difficult, and so people with it hold their loved ones to the best of their ability, clumsily.

[5] This could, however, be a converse (but also problematic) form of Scully (2010) performing deafness for the sake of prompting others for assistance. The implications of Dubiel's (2009) attempts to "pass" as able-bodied, in the broader context of ableist attitudes about and stigmas against disability, deserve their own investigation.

when the user has family members who think adjusting the stimulator is a solution to all of the users' problems (Klein et al. 2016). Other users, however, may be able to make use this feedback to improve their relationship with DBS—that is, friends and family can make themselves useful by learning how DBS works, by noticing side-effects, by corroborating the user's worries, by advocating for the user, or just by being understanding when the user's agentic plan falls flat.[6]

6. Conclusions

My central claim has been that DBS systems are relational partners insofar as users cultivate relationships with these devices as they perform agency through their use of them. I considered several views that attempt to make sense of the ways that DBS systems change their users' lived-experiences, and identified several problems within them. Chief among these problems was the tacit assumption that, in extreme cases of personality change, DBS users lose their identity or agency as a matter of course. To address this problem, I proposed that our view of agency should capture the struggles people have with their stimulator devices, and their attempts to "make do" with them.

After considering two key examples of the relationships users have with their neurostimulators, I offered a synthesis of several views to frame and characterize these struggles. On my view, DBS systems are capable of holding us—helpfully or disruptively, depending on the context—to certain performances of agency. As such, some users form (or will have to form) intricate partnerships with these systems in order to make use of them. Through these

---

[6] Fins (2018), relatedly, proposes a *mosaic approach to decisionmaking* for people using DBS to counteract injury-induced minimally conscious state (MCS). As the injured person's agency reemerges, Fins argues, they will need the assistance of several other agents—a surrogate, medical professional, and an advocate—that form a mosaic of agents that make decisions together.

relationships, users may confer animacy on these technologies to the point where their exercises of agency are supported, encumbered, and otherwise "intertwined with" them.

I, finally, identified a number of skills that users can and already do use to navigate challenges with and alongside DBS. In the end, we are left a view that—I hope—promotes DBS from the status of an inert device that either enables or confounds its user's agency, to the status of an interactive system with which users collaborate. Further, and more importantly, we are left with a view that promotes the DBS user from a vulnerable target for potential threats, to an (at least partially) empowered agent.

References

Agid, Y, M Schüpbach, M Gargiulo, et al. 2006. "Neurosurgery in Parkinson's disease: the doctor is happy, the patient less so?" *Journal of Neural Transmission* 70: 409–14.

Alonso, Pino, Daniel Cuadras, Loes Gabriëls, et al. 2015. "Deep Brain Stimulation for Obsessive-Compulsive Disorder: A Meta-Analysis of Treatment Outcome and Predictors of Response." *PLoS ONE* 10 (7): e0133591.

Baylis, Françoise. 2011. "'I Am Who I Am': On the Perceived Threats to Personal Identity from Deep Brain Stimulation." *Neuroethics* 6 (3): 513–26.

Baylis, Françoise. 2012. "The Self in Situ: A Relational Account of Personal Identity." In *Being Relational*, ed. Jocelyn Downie and Jennifer J. Llewellyn, 109-31. Vancouver: UBC Press.

Belser, Julia Watts. 2016. "Vital Wheels: Disability, Relationality, and the Queer Animacy of Vibrant Things." *Hypatia* 31 (1): 5–21.

Bierria, Alisa. 2014. "Missing in Action: Violence, Power, and Discerning Agency." *Hypatia* 29 (1): 129–45.

Brown, Timothy, Margaret C. Thompson, Jeffrey Herron, et al. 2016. "Controlling Our Brains a Case Study on the Implications of Brain-Computer Interface-Triggered Deep Brain Stimulation for Essential Tremor." *Brain-Computer Interfaces* 3 (4): 165-70.

Clark, Andy. 2008. Supersizing the Mind: Embodiment, Action, and Cognitive Extension. Oxford: Oxford University Press.

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.

Dos Santos, Joao Flores Alves, Sophie Tezenas du Montcel, et al. 2017. “Tackling Psychosocial Maladjustment in Parkinson’s Disease Patients Following Subthalamic Deep-Brain Stimulation: A Randomised Clinical Trial.” *PLoS ONE* 12 (4): e0174512.

Dubiel, Helmut. 2009. Deep within the Brain: Living with Parkinson’s Disease. New York: Penguin.

Fins, Joseph J. 2018. “Mosaic Decisionmaking and Reemergent Agency after Severe Brain Injury.” *Cambridge Quarterly of Healthcare Ethics* 27 (1): 163-74.

Gallagher, Shaun. 2000. “Philosophical Conceptions of the Self: Implications for Cognitive Science.” *Trends in Cognitive Sciences* 4 (1): 14–21.

Gallagher, Shaun. 2018. “Deep Brain Stimulation, Self and Relational Autonomy.” *Neuroethics*, March: 1–13.

Glannon, W. 2009. “Stimulating Brains, Altering Minds.” *Journal of Medical Ethics* 35 (5): 289–92.

Goddard, Eliza. 2017. “Deep Brain Stimulation through the ‘Lens of Agency’: Clarifying Threats to Personal Identity from Neurological Intervention.” *Neuroethics* 10 (3): 325–35.

Goering, Sara, Eran Klein, Darin D Dougherty, et al. 2017. “Staying in the Loop: Relational Agency and Identity in Next-Generation DBS for Psychiatry.” *AJOB Neuroscience* 8 (2): 59–70.

Haan, Sanneke de, Erik Rietveld, Martin Stokhof, et al. 2013. “The Phenomenology of Deep Brain Stimulation-Induced Changes in OCD: An Enactive Affordance-Based Model.” *Frontiers in Human Neuroscience* 7 (October): 653.

Haan, Sanneke de, Erik Rietveld, Martin Stokhof, et al. 2017. “Becoming More Oneself? Changes in Personality Following DBS Treatment for Psychiatric Disorders: Experiences of OCD Patients and General Considerations.” *PLoS ONE* 12 (4): e0175748.

Hariz, Marwan I., Stig Rehncrona, Niall P Quinn, et al. 2008. “Multicenter Study on Deep Brain Stimulation in Parkinson’s disease: An Independent Assessment of Reported Adverse Events at 4 Years.” *Movement Disorders* 23 (3): 416–21.

Hebb, Adam O., Jun Jason Zhang, Mohammad H Mahoor, et al. 2014. “Creating the Feedback Loop.” *Neurosurgery Clinics of North America* 25 (1): 187–204.

Kellmeyer, Philipp, Thomas Cochrane, Oliver 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): 623–33.

Klaming, Laura, and Pim Haselager. 2013. “Did My Brain Implant Make Me Do It? Questions Raised by DBS Regarding Psychological Continuity, Responsibility for Action and Mental Competence.” *Neuroethics* 6 (3): 527–39.

Klein, Eran, and C. S. Nam. 2016. “Neuroethics and Brain-Computer Interfaces (BCIs).” *Brain-Computer Interfaces* 3 (3): 123–25.

Klein, Eran, Sara Goering, Josh Gagne, et al. 2016. “Brain-Computer Interface-Based Control of Closed-Loop Brain Stimulation: Attitudes and Ethical Considerations.” *Brain-Computer Interfaces* 3 (3): 140–48.

Kraemer, Felicitas. 2011a. “Authenticity Anyone? The Enhancement of Emotions via Neuro-Psychopharmacology.” *Neuroethics* 4 (1): 51–64.

Kraemer, Felicitas. 2011b. “Me, Myself and My Brain Implant: Deep Brain Stimulation Raises Questions of Personal Authenticity and Alienation.” *Neuroethics* 6 (3): 483–97.

Kraemer, Felicitas. 2013. “Authenticity or Autonomy? When Deep Brain Stimulation Causes a Dilemma.” *Journal of Medical Ethics* 39 (12): 757–60.

Kuncel, Alexis M., Scott E. Cooper, Barbara R. Wolgamuth, et al. 2006. “Clinical Response to Varying the Stimulus Parameters in Deep Brain Stimulation for Essential Tremor.” *Movement Disorders* 21 (11): 1920–28.

Lindemann, Hilde. 2009. “Holding One Another (Well, Wrongly, Clumsily) in a Time of Dementia.” *Metaphilosophy* 40 (3-4): 416–24.

Lindemann, Hilde. 2014. Holding and Letting Go: The Social Practice of Personal Identities. Oxford: Oxford University Press.

Lipsman, Nir, and Walter Glannon. 2013. “Brain, Mind and Machine: What Are the Implications of Deep Brain Stimulation for Perceptions of Personal Identity, Agency and Free Will?” *Bioethics* 27 (9): 465–70.

Mackenzie, Catriona, and Mary Walker. 2015. “Neurotechnologies, Personal Identity, and the Ethics of Authenticity.” In *Handbook of Neuroethics*, ed. Jens Clausen and Neil Levy, 373–92. Dordrecht: Springer Netherlands.

Mayberg, Helen S. 2008. “Deep Brain Stimulation for Treatment-Resistant Depression.” *Journal of Affective Disorders* 107 (March): S23.

Mayberg, Helen S, Andres M Lozano, Valerie Voon, et al. 2005. “Deep Brain Stimulation for Treatment-Resistant Depression.” *Neuron* 45 (5): 651–60.

Mayberg, Helen S., Patricio Riva-Posse, and Andrea L. Crowell. 2016. “Deep Brain Stimulation for Depression: Keeping an Eye on a Moving Target.” *JAMA Psychiatry* 73 (5): 439–40.

Mecacci, Giulio, and W. F. G. Pim Haselager. 2014. "Stimulating the Self: The Influence of Conceptual Frameworks on Reactions to Deep Brain Stimulation." *AJOB Neuroscience* 5 (4): 30–39.

Meyers, Diana T. 1989. *Self, Society, and Personal Choice*. New York: Columbia University Press.

Meyers, Diana T. 2000. "Intersectional Identity and the Authentic Self?" In *Relational Autonomy*, edited by Catriona Mackenzie and Natalie Stoljar, 151-180. Oxford: Oxford University Press.

Pham, Uyen, Anne-Kristin Solbakk, Inger-Marie Skogseid, et al. 2015. "Personality Changes after Deep Brain Stimulation in Parkinson's Disease." *Parkinson's Disease* 2015 (13): 1–7.

Schechtman, Marya. 2010. "Philosophical Reflections on Narrative and Deep Brain Stimulation." *The Journal of Clinical Ethics* 21 (2): 133–39.

Schlosser, Markus. 2015. "Agency." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta.

Schüpbach, M, M Gargiulo, M L Welter, and L Mallet. 2006. "Neurosurgery in Parkinson disease A distressed mind in a repaired body?" *Neurology* 66 (12): 1811–16.

Scully, Jackie Leach. 2010. "Hidden Labor: Disabled/Nondisabled Encounters, Agency, and Autonomy." *IJFAB* 3 (2): 25–42.

Sharp, Daniel, and David Wasserman. 2016. "Deep Brain Stimulation, Historicism, and Moral Responsibility." *Neuroethics* 9 (2): 173–85.

Smeding, H. M. M., A. E. Goudriaan, E. M. J. Foncke, et al. 2007. "Pathological Gambling after Bilateral Subthalamic Nucleus Stimulation in Parkinson Disease." *Journal of Neurology, Neurosurgery & Psychiatry* 78 (5): 517–19.

Specker Sullivan, Laura. 2019. "Insight and the No-Self in Deep Brain Stimulation." *Bioethics* 33 (4): 487-94.

Voon, Valerie, Cynthia Kubu, Paul Krack, et al. 2006. "Deep Brain Stimulation: Neuropsychological and Neuropsychiatric Issues." *Movement Disorders* 21 (S14): S305–27.

Wendell, Susan. 1989. "Toward a feminist theory of disability." *Hypatia* 4 (2): 104–24.

Widge, Alik S., Amanda R. Arulpragasam, Thilo Deckersbach, et al. 2015. *Deep Brain Stimulation for Psychiatric Disorders*. Vol. 57. An Interdisciplinary, Searchable, and Linkable Resource. Hoboken, NJ, USA: American Cancer Society.

Yuste, Rafael, Sara Goering, Blaise Agüera y Arcas, et al. 2017. “Four Ethical Priorities for Neurotechnologies and AI.” *Nature News* 551 (7679): 159–63.

Jumping Through Loops: Closed-Loop Control of DBS and Its Impact on User Agency

Abstract

Deep–Brain Stimulators (DBS) has become an effective means of treating a variety of neurological conditions, and an investigational device for treating psychiatric conditions. Most DBS systems work in an *open-loop* way, applying constant stimulation over the lifetime of the treatment. The next generation of DBS systems, however, are *closed-loop*, recording the user's neural activity in real time, and adjusting the level of the stimulator in order to fit their needs or demands. However, depending on what neural activity is recorded, what features are extracted from those recordings, and how those features are used to modulate stimulation parameters, different modes of controlling DBS will likely produce different experiences for their users—with different practical (and moral) outcomes. I suggest that the constituent components of closed-loop system (neural sensing, feature extraction, and stimulation decisions) change the topology of users' decisional loops in ways that call the trustworthiness of these devices (and the institutions that create them) into question.

1. Introduction

Deep–Brain Stimulators (DBS) have become an effective means of treating a variety of neurological conditions: Essential Tremor (ET), Parkinson's disease (PD), dystonia, and refractory epilepsy. Several trials are investigating DBS for a wide variety of treatment-resistant neurological and psychiatric conditions. The most common systems work in an *open-loop* way: stimulation electrodes implanted in the brain are connected to a fully-implanted signal generator. The signal generator supplies an electrical signal to the brain through the electrodes. A technician or trained clinician can turn the system off (and on), and change the signal's parameters (voltage, frequency, etc).

Open-loop stimulation, however, has several drawbacks. Stimulation can cause side-effects, depending on where the stimulation electrodes are placed, the stimulation parameters used, and the user's neurophysiology. Some users, even, report neuropsychiatric problems post-surgery: for example, a small subset of people with Parkinson's develop impulsivity and risk-taking behaviors (e.g., gambling) (Smeding et al. 2007). Further, these systems use a battery contained in the stimulator, and once the battery is depleted—between two and five years—the entire stimulator unit must be replaced through a minor surgery (Fakhar et al. 2013). Finally, some DBS may become less effective over time—many users require increased voltages to receive the same levels of effectiveness (Shih et al. 2013).

The next generation of Deep–Brain Stimulation systems, however, offer a possible solution to many of the aforementioned problems with current-generation systems: persistent side-effects, battery depletion, and decreasing effectiveness. Where the current generation of DBS applies stimulation constantly, next-generation, *closed-loop* DBS (CL-DBS) systems would record the user's neural activity in real time and use it to adjust the level of the stimulator in order to fit their needs or demands (Arlotti et al. 2018; Houston et al. 2019). One such system could record and analyze neural signals related to the user's condition to determine what kinds of stimulation they need. Several groups have termed this *adaptive* DBS, or aDBS (Weiss and Massano 2018; Habets et al. 2018). Another could allow the user to modulate their own neural activity in order to control their stimulator's level (Thompson and Herron 2016). We can call these systems *brain-computer interface–triggered* DBS (or BCI-DBS), given their similarity to systems that use neural signals to control computer interfaces, robots, and similar systems (Klein and Nam 2016; Klein et al. 2016; Brown et al. 2016; Thompson and Herron 2016).

These *adaptive* or *brain-computer interface (BCI)* driven DBS systems could potentially reduce the number of battery-replacement surgeries required over the lifetime of the treatment: by applying stimulation less often and only when needed, power is conserved and battery life is extended. They could also help users avoid the potential side-effects of stimulation: that is, if the system applies stimulation less often, the user will likely experience fewer physiological side-effects, or the user might be able to control when the side-effects occur through control of the stimulator. Further, there is the possibility that applying stimulation conservatively will lower the chance that the user will build a tolerance to stimulation, and raise the chance that comparatively low voltages will remain clinically effective for a long period of time. Finally, there is the possibility that novel and more effective therapies will require the use of closed-loop techniques.

However, depending on what neural activity is recorded, what features are extracted from those recordings, and how those features are used to modulate stimulation parameters, different modes of controlling DBS will likely produce different experiences for their users—with different practical (and moral) outcomes. If the system is automated, then the user will need to learn when the system will adjust its parameters and figure out how to prepare for the consequences of those adjustments. If, however, the system is user-controlled—that is, if the user can control the system by modulating their own neural signals—the user will need to learn how to control the system reliably. They will also need to learn or decide for themselves when it is appropriate to change their stimulator's settings. These modes of control, I've argued elsewhere, seem to present users with challenges to their competence as agents as they manage DBS' tradeoffs in the context of their already complicated lives (Brown, forthcoming). Further, different disorders may call for different modes of control. For example, a person with depression might not trust themselves to manage their own stimulation levels (Klein et al. 2016;

Goering et al. 2017), where a person with Parkinson's might welcome the possibility (Dubiel 2009).

Several authors (Kellmeyer et al. 2016; Goering et al. 2017) take advantage of the turn of phrase "closed-loop," arguing that the CL-DBS user's position inside or outside the system's control loop will have implications for how they exercise agency and how responsible they are for actions undertaken with DBS assistance. But what is the topology of this loop? What are its components? What are its pressure points? And how will users' experiences of agency change as a result of being kept in one part of the loop and excluded from others? To ask the question another way, where will the user (and other stakeholders) be in the loop, and how will that impact the user's agency? In most conversations about closed-loop systems so far, users are kept out of the loop, given that they are unable to exercise decisional control over stimulation parameters. There is, however, more ground to cover.

In this paper, I suggest that many of the most pressing moral issues with CL-DBS concern the design and use of specific components of the closed-loop system: its neural sensing platform, its feature extraction algorithms, and (of course) its stimulation capabilities. I'll maintain that the design of each of these components of a closed-loop system can change the user's decisional loop in ways that reflect the assumptions, biases, and decisions made by a variety of stakeholders and institutions. The way each component is designed, I argue, comes with deep implications for whether or not users can *trust* these systems to support their agency. I close with a series of recommendations for investigators, device manufacturers, medical professionals to follow as they design trustworthy CL-DBS systems.

2. The closed-loop approach to DBS

Let’s begin by taking a closer look at the closed-loop approach to DBS and how it’s different from open-loop approaches. By the end, we should be able to see the components of the closed-loop system plainly. We’ll also find that the two approaches overlap in terms of how users interact with them and the potential for these systems to bolster or undermine users’ expressions of agency.

Recall that open-loop systems are comprised of a signal generator that applies stimulation at either a constant rate or on a predetermined schedule. The loop is “open” insofar as the system has no means of receiving feedback about the system that its stimulating. A typical system that treats ET, for example, would apply stimulation at a frequency and amplitude selected and tuned by a physician—the system itself is not able to determine what settings are needed on its own.

Many open-loop systems for motor disorders, however, admit of some forms of user control. DBS users are typically given a “patient programmer,” a small handheld device that allows the user to switch the device off and on. These allow the user to stop their stimulator in cases of malfunction, or at their own discretion to conserve battery life or to avoid side-effects. Some users’ patient programmers give them even finer control over stimulation parameters (usually voltage and frequency). Some users with PD welcome the ability to match their stimulation parameters with their everyday activities (Dubiel 2009).

Closed-loop DBS systems, on the other hand, record neural signals, analyze the recordings for salient features, and apply stimulation when certain conditions are met. That is, a closed-loop system uses some form of feedback to determine when to stimulate—the feedback “closes the loop.” A closed-loop system that treats ET, for example, may record areas of the brain associated with upper limb function using electrocorticography (ECoG) electrodes, detect the different forms of limb movement that trigger tremor, and apply enough stimulation so that the user can move their limb smoothly (Herron et al. 2016; Thompson and Herron 2016; Houston et al. 2019). Another system that treats epilepsy, like the Neuropace Responsive Neurostimulation (RNS) system, uses an ECoG array to detect pre-seizure activity, and applies stimulation to preemptively disrupt oncoming seizures (Geller et al. 2017; Jobst et al. 2017).

An essential component of a CL-DBS system, then, is some means of tracking *biomarkers* that stand in as surrogates for clinician assessments of symptoms of disease. These biomarkers, then, give the system feedback. Some biomarkers are used to track the symptoms of the condition treat it—as with the Neuropace RNS. Biomarkers can also, however, be used to track activity adjacent to the condition which could serve as useful control signals—as a CL-DBS system for ET would use biomarkers for upper limb movement. Further, while we might imagine that all systems will track *internal* biomarkers like neuronal electrical activity or biochemical changes, it may also make sense to rely on *external* biomarkers. That is, instead of tracking motor cortical activity (Herron et al. 2016; Thompson and Herron 2016; Houston et al. 2019), a system treating ET or PD could track accelerometer and gyroscope data from a worn smartwatch and apply stimulation when it detects tremor (Herron and Chizeck 2014).

Ideal biomarkers, according to Hoang et al's (2017, 2019) comprehensive review of biomarker use in CL-DBS, have several characteristics. Four are relevant for us now. First, biomarkers should correlate to symptoms of the disease treated. We've already seen, however, that other biomarkers could track normal activities that are impeded by the disease treated. Second, biomarker measurements should change dynamically as the system applies stimulation, or as the circuits measure change in real time—the biomarker must give the system access to outcomes. Third, it must be possible to distinguish between the salient biomarker and other, potentially obfuscating, signals: background noise, irrelevant adjacent activity, and so on. Finally, clinicians and technicians should be able to fine-tune measurements of the biomarker to account for differences between users.

The success of a CL-DBS system, then, depends on being able to find biomarkers that correlate either with disease symptoms or activities that are impeded by symptoms. It may be difficult, in many use cases, to find biomarkers of symptoms of certain conditions—depending on how much we know about the condition. For example, we can identify corticography sites for detecting limb movement and tremor. There are established methods for using corticography to track seizure activity. It is not clear, however, what biomarkers we can use to track psychiatric disorders like major depressive disorder (MDD), anxiety, and obsessive-compulsive disorder (OCD) in real time. While some groups have identified potential targets for electrophysiological recording to track the state of MDD, OCD, PTSD and anxiety (Neumann et al. 2014; Bina et al., 2018), some conditions may require the use of sensing technologies not yet available in consumer medical neuro-devices—e.g., biochemical sensors (Hoang et al. 2017).

Biomarker recordings, further, must be *classified* before they can be used to determine when the system will apply stimulation. A CL-DBS system for ET, for example, must extract features like tremor or limb movement from a recording, often in real time, and often while conserving battery power. The analogous system for epilepsy must be able to distinguish between signals related to a potential seizure, everyday neural activity, background noise, and so on. Rao (2019) propose using machine learning techniques—artificial neural networks (ANNs) and deep-learning—to both *classify* data and *make decisions* based on it. They term these systems *neural co-processors*. ANNs are algorithms that emulate neurological systems. ANNs work by feeding a series of simple inputs into a series of interconnected "neurons." These neurons are “weighted” to fire and produce outputs in response to the right series of inputs. We can call the complete set of weights a *model*.

ANNs are powerful precisely because we needn’t set each of their parameters manually: we can automate this process so that the ANN will form a model to fit the data we feed it, classifying it into the categories without much human intervention. We call this process of optimizing the network’s parameters *training* the model. There are several ways to train a model—that is, several ways for the system to learn to classify data. *Supervised learning* requires that a person label data used to train the system using the categories we want the model to classify. For example, to train a CL-DBS for ET to recognize when a user is moving their arm, we would train it with data we’ve manually classified as moving or not moving. The system would attempt to classify the data on its own, check its accuracy against the training data, adjust its parameters, and repeat the process until it reaches a high enough accuracy. With *unsupervised learning*, the system uses a variety of methods to arrive at categories on its own, without relying on labeled training data. Closed-loop DBS systems that make use of machine learning techniques could be used to make use of several sources of data at once, classify more complicated patterns of electrophysiological behavior, and orchestrate complicated stimulation patterns across multiple electrodes (Hoang et al. 2017).

3. Being kept in (or out of) the loop

We might think of ordinary (motor) action as a kind of “closed-loop.” When a person moves their arm deliberately, their intention to move their arm is (at least part of) the cause of their limb’s movement. ET and PD disrupt this loop by introducing an obstacle to limb movement: the person wants and intends to move their arm, but an involuntary tremor impedes them from moving their arm the way they want to. But where ordinary motor function is disrupted by a motor disorder, OL-DBS is meant to restore that control—or make the loop work the way it originally did. OL-DBS, however, introduces more actors into “the loop,” widening it to include clinicians, technicians, and caretakers. Once the system is implanted, a clinician must tune its parameters, the user (or their caretakers) must learn to operate the system (e.g., charge the battery, or power the system off), a technician may need to service the device if it malfunctions, and a surgeon will need to replace the battery when it depletes or degrades.

OL-DBS’ loop might be so open that it feels *loose*. That is, depending on the user’s relationship with the device or their relationship with the people who help them use the device (Brown, forthcoming), the user may not feel like they have much of a say over how the device operates, or how their body operates as a result. In the previous section, we saw that CL-DBS’s loop is “closed” insofar as the system is able to record and analyze biomarkers that help it determine when to initiate stimulation and which stimulation parameters to use. We also saw that at least some of these systems could make use of biomarkers that users can modulate deliberately, giving them BCI control of their DBS.

These systems, in some ways, emulate neural structures in the brain and spinal cord by reading from some neural signals as inputs and stimulating other neural circuits in response. An *adaptive* closed-loop system could modulate stimulation parameters without the help of a clinician, and without needing to use an external device. A *BCI-controlled* system could allow the user to modulate stimulation parameters directly, in real time. Perhaps such systems could “tighten the loop” for some users: matching stimulation to the user’s needs, or giving the user direct control over stimulation.

Closed-loop systems, however, may present their own set of new problems. The choice between an adaptive and a BCI-controlled system will also be the choice between having more or less control over the system. In the former, the system decides when to stimulate based on non-volitional states detected within recordings. In the latter, the system detects changes in neural activity that users can modulate volitionally. The decision to use one type of system over the other, then, is the decision to either allow the system to keep the user out of or include the user in the decisional loop.

Kellmeyer et al. (2016) ask how this decision can strengthen or weaken the user's capacity for autonomous agency. On one hand, being kept out of the loop could relieve the user of the burden of learning to control, navigate, or moderate the system (depending on how the system is controlled). This, however, means that the user must rely on the system to make the right decisions for their needs in a specific context. If the system makes the wrong decisions, the system could impede the user's agency by applying stimulation in the wrong ways, in the wrong contexts, or even failing to apply stimulation when the user is relying on it. In some ways, we can think of the device itself (or, more modestly, the device's designers and maintainers) as responsible for the successful or failed attempts to apply stimulation. After all, the system itself takes control of a process that we could imagine the user having control of.[1]

---

[1] There are several questions we could ask. Do we (or can we) praise or blame our limbs for reflex activity? Do we (or can we) praise or blame areas of our brains for generating certain kinds of activity—e.g., dreams, emotions—without our explicit say-so? Some people colloquially talk about different parts of their bodies as though they have their own agency. A runner, for example, might say, "my legs won't do what I tell them to do today," if they feel fatigued or uncoordinated. A person who is prone to nightmares, to use another example, might say, "my brain does strange things at night" to explain the imagery within and frequency of their nightmares. We can imagine a person speaking similarly about their CL-DBS, if they know that the system is complex and dynamic. That is, a person with PD might say, "my implant does what it wants to do at times" when explaining why their stimulator causes them side-effects unpredictably. It's not clear, however, if this way of speaking is apt. A deeper exploration is warranted, but such an exploration is outside of the scope of this paper. See, however, ("Rethinking Neuroethics in the Light of the Extended Mind Thesis" 2007) and Buller (2011) for the beginnings of such an exploration in the context of Andy Clark's (2008) extended mind hypothesis.

To say that the device itself is an agent, and can be responsible for some action, seems far-fetched at first. However, the more automated systems rely on machine learning techniques to both classify neural recordings and make treatment decisions, the more it seems apt to say that DBS systems exercise agency, or that users *share agency* with their implants (Yuste et al. 2017; Goering et al. 2017; Brown, n.d.). Putting aside the question of whether we can call an artificially–intelligent medical device an agent, some users may welcome the assistance of a more robust device that can anticipate their needs. Some users, however, may feel even more *out of the loop* simply because such a system sets stimulation without their input. Users may find themselves fighting against the system for decisional control (Brown et al. 2016), or even feeling unclear about the source of their actions (Klein et al. 2016).

On the other hand, if the user is kept in the decisional loop, the user will be able to retain autonomy over the decision to stimulate, and (by extension) the results of the stimulation. Some users may welcome being closer to the center of DBS' decisional loop. This could, however, render users responsible in ways that complicate their agency. That is, since users will need to learn to operate the system (by learning to modulate their neural activity), and will be tasked with operating it regularly, BCI-DBS could prove burdensome to some users (Brown et al. 2016).

Where Kellmeyer et al consider CL-DBS' effects on autonomous agency insofar as it renders the user more or less responsible for actions, Goering et al. (2017) attend to how closed-loop systems could support user agency. Their move is similar to François Baylis' (2011), who maintains that some commenters' (Glannon 2009; Kraemer 2011; Schermer 2011) worries that OL-DBS causing personal identity changes are the result of a misunderstanding of how personal identity works. Personal identity, on her view, is dynamic, narrative, and relational. While OL-

DBS may pose a threat to users when it “undermines agency to such an extent that the person is no longer able to meaningfully contribute to the authoring of her own life” (525), most of the changes users undergo as a result of using DBS are on par with the changes they’d undergo after any other significant event in their lives.

Goering et al. (2017) extend a similar view to CL-DBS. They argue that these worries about CL-DBS undermining and bolstering agency underestimate how much of agency itself is relational. That is, most agents rely on other agents and assistive tools to support their agency. On their view, CL-DBS users will rely on their implanted systems to “support [them] in achieving [their] aims,” but this isn’t much different from how we “rely on other people and things (including devices like cell phones, computers, etc.) to support our activities” (Goering et al. 2017, 67). Goering et al. (2017), taking cues from Barclay (2000), suggest that agency itself is relational, that “an individual’s agency (connecting her motives and actions) can be supported through and entwined with a variety of external others” (66). When we recognize just how relational agency is more broadly, it is easier to see how neurotechnologies like CL-DBS fit into the networks of relationships we already participate in.

CL-DBS users, then, rely on their implanted systems in a kind of relationship, but this relationship is not so different from other relationships they enter and rely on for support. Further, relationships between CL-DBS systems and users exist within the space of the other relationships users rely on. If a CL-DBS system behaves in a way that makes a user feel “out of the loop”—if it malfunctions, or makes the user behave in undesirable ways—the user will have to rely on others to support them. Friends and family might show the user patience as they adjust to the system. Or they might help them see how their relationship with DBS has turned difficult or unproductive. Or, finally, they might support the user as they “reclaim their agency” from CL-

DBS (Goering 2014). The user's team of clinicians can adjust the system to better suit the user's needs, teach the user how to get the system to behave the way they want it to, or affirm that the system is working as it's intended. This, on Goering et al. (2017)'s view, is a point of solace: CL-DBS is less foreign than it seems, and users will have most of the tools they need to handle its challenges.

Let's retrace our steps. Kellmeyer et al. (2016) argue that being in the loop (or being left out of it) can complicate agency and responsibility. Goering et al. (2017) argue that these loops are similar to loops we're all familiar with. We might say that one benefit of using a closed-loop system instead of an open-loop system is to tighten the user's decisional loop. It turns out, however, a closed-loop system may *loosen* the decisional loop to include more participants instead—CL-DBS may even make the decisional loop *porous*. With an OL-DBS system, the stimulator often runs constantly unless the user or clinician decides to turn it off. With a CL-DBS system, however, the system makes a stream of decisions on the user's behalf. As we saw in §2, users will have multiple points of contact with these closed-loop systems.

This, however, means that users can be kept in or out of the decisional-loop in a wide variety of ways. Users, for example, may feel in the loop about how data is classified, only to feel out of the loop about how or when it administers stimulation. If Kellmeyer et al. (2016) are right to say that CL-DBS systems could shift responsibility from the users to devices and their manufacturers, then we must consider carefully what those responsibilities are and decide if they ought to take them. If Goering et al. (2017) are right to say that CL-DBS will play a familiar role in users' expressions of agency, we must ask a familiar question: can users *trust* CL-DBS to act on their behalf and for their best interests? In the next section, we'll explore this question.

4. Can CL-DBS be trusted in the loop?

By now, we should have a better grasp of what constitutes a closed-loop system: a platform that records neural signals, a means of extracting features from those neural signals, and algorithms determine when (and how) stimulation is applied. We should also have a better grasp of how these systems could impinge on or support user agency. But we also saw that "being in the loop" is not as simple as either being "in control" or "out of control" of the system or its effects. Being in the loop is also a function of being able to get perspective on, make sense of, negotiate with, or push back against CL-DBS. Even this, however, is more complicated than it seems prima facie.

CL-DBS' decisional loop is a system whereby neural recordings are classified and used to make the decision to stimulate. Users can be brought into different parts of this loop with different effects. For example, we've already seen that for some use-cases there are several possible biomarkers available. For ET and PD, it remains to be seen if CL-DBS systems will make use of external wearable sensors or internal neural recordings to track either disease symptoms (e.g., tremor). The choice of biomarker itself determines *how* users are brought into the loop, and what's at stake. In this section, I'll demonstrate how CL-DBS systems can bring users into the loop through several entrypoints—in how/what they record, how they classify those recordings, and what decisions they make from them—and each of these entrypoints have the potential to complicate user agency by changing the topology of the loop. These complications to agency raise the question: can users trust CL-DBS in their decisional loops?

First, however, let's say more about what we mean by "trust" here. Carolyn McLeod, in her *Self-trust and Reproductive Autonomy* (2002), offers a useful framework for understanding trust that should also extend to CL-DBS. She argues against views that characterize trust as an optimism about someone's (or something's) competence and goodwill. When a person trusts their doctor, it's not only because they believe the doctor has an expanse of medical knowledge or because they believe the doctor has advanced skill as a medical practitioner—though we hope that the doctors we trust have these attributes. Even the most medically competent doctors can fail their patients: by failing to respect their autonomous choices, by not recognizing their cultures and backgrounds, or by allocating medical resources unjustly. Nor is it enough for a doctor to have goodwill—or "kindly good feelings" (21)—toward patients. Further, if a doctor is only respectful and just as a result of their goodwill, it is crucial that we ask if the doctor will remain respectful and just if and when their goodwill fades. Trustworthy doctors are respectful and just toward their patients even if they don't have good feelings toward them.

McLeod, then, defends views on which we trust one another to not only act competently but also with moral integrity. We want the people we trust to have an "enduring commitment to acting in a morally respectful way toward us and we want their actions to accord with that commitment" (23). On her view, then, trusting another means that we know what they stand for and that they will stand up for it. At first glance, it seems like McLeod's view pertains to people and not devices. If Goering et al. (2017) are right, however, people will rely on their CL-DBS systems the way they might rely on friends and family who support their agency. This reliance would seem to require a kind of trust: that the system will act in ways that do not betray the user's values. One question we can ask, then, is if users will be able to count on CL-DBS to eschew systems of injustice or oppression. This is a difficult task indeed, given the technologies underlying CL-DBS and given how these technologies change the topology of users' decisional loops. In what remains of this section, we will take a closer look at each component of the closed-loop system to see how users are rendered vulnerable to possible violations of their trust.

Biomarkers, Recordings, and Big Data

Up until now, we've talked about biomarker recordings in a neutral way. We've taken for granted that the users would be comfortable allowing a system to record their biological data continuously, just as long as the system is able to improve their quality of life. We should not, however, accept this assumption for too long. CL-DBS systems will likely record data that many of us will count as sensitive or private. How comfortable will CL-DBS users with ET, for example, be with the knowledge that their system collects motor activity data during their most intimate moments?

It is not clear what kinds of data we should permit closed-loop systems to collect, how these data ought to be stored, who should be allowed access to these data, and what uses of these data are off-limits. It is important to remember, then, that medical device companies (e.g., Medtronic, Boston Scientific, and Neuropace) are designing CL-DBS systems within a broader context where larger tech companies (e.g., Google, Microsoft, Amazon, and IBM) have made it possible to analyze data on a massive scale using cloud-based storage and computing platforms. Several of these companies are encouraging device manufacturers to link their products to an entire ecosystem of services so that users can receive deeply personalized healthcare. Implanted medical devices—most commonly cardiac monitors and cardioverter-defibrillators—that send their data to electronic medical record systems are becoming increasingly common. CL-DBS systems are likely next targets for integration.

Kellmeyer (2018) warns against potential moves to privatize or commodify biometric data. This warning is apt amidst parallel efforts to make human neural tissue images (Peng et al. 2015) and neural recordings (Faraut et al. 2018; Miller 2019) available in open datasets, as well as to consolidate neural tissue samples and neural recordings into centralized biobanks (Nichols et al. 2014; Freund et al. 2018). Centralizing and analyzing neural data at scale will likely make exciting new innovations possible, but it's also possible that the organizations behind these efforts will use data in ways that go against the wishes of both individuals and communities. We need only look toward Arizona State University's misuse of biological samples and genetic data collected from the Havasupai tribe—where researchers analyzed data for features that would contradict the tribes' narratives about their origins—to see the potential for harm (Mello and Wolf 2010; Sterling 2011). CL-DBS has the potential to draw from and contribute to the growing pool of user data. More robust forms of CL-DBS could make use of multiple biomarkers, electronic health records, data collected from smartwatches and health-oriented wearables (e.g., Fitbit), GPS location data, and so on. CL-DBS will also collect neural data, produce analyses of those data, and also generate their own diagnostic data.

Users may, sure, want control over what signals are recorded, how they're stored, and who they're shared with. They may also want to know what institutions will do with their data, and whether or not they can trust these institutions to have their best interests in mind if (or when) they use the data they collect. CL-DBS designers, however, will need to be careful about which sources of data their systems collect and analyze. Further, they will need to decide if and how the data they collect contributes to the growing pool of data. A CL-DBS designer has a choice to restrict or prevent third-party access to neural recordings or more detailed diagnostic information—either by anonymizing it or discarding it after on-board analysis (Bonaci et al. 2014; Bonaci, Calo, and Chizeck 2014). Measures like these could make CL-DBS more trustworthy by ensuring that they do not share data without the user's explicit consent.

Classifier Algorithms and the Risk of Implicit Bias

Earlier (in §2), we explored the possibility of using ML algorithms to classify signals used for closed-loop control (Rao 2019). Artificially intelligent systems, specifically ANNs, are notoriously vulnerable to pernicious social biases—depending on the methods used to train the model, the data used to train it, what measures put in place to validate the model (if any), and how representative early adopters of the technology are of the broader population of users. Examples of these biases are abundant. Joy Buolamwini (2018) found that three commercial facial recognition algorithms that also perform gender detection (created by Microsoft, IBM, and Face++) were less able to classify the gender of darker-skinned faces.[2] Wilson, Hoffman, and Morgenstern (2019) found that models used for object detection in early autonomous vehicles are less able to detect dark-skinned pedestrians.

People with disabilities, in particular, are at risk of being subjected to broader societal biases through the technologies they're expected to use every day (Guo et al. 2019). Assistive technologies, like object recognition systems for the visually impaired (MacLeod et al. 2017), may operate in ways better suited for the unimpaired or nondisabled. These systems may even fail to recognize them entirely. These crucial missteps raise a number of issues. The very act of creating an algorithm to classify data seems to fix an ontology in place. For example, returning to Buolamwini (2018), to create an algorithm that classifies pictures of faces by gender, we must presume that: there are stable gender categories, we know what they are, and we can detect them using facial features alone. These assumptions reinforce gender norms and exclude gender identities or expressions that do not conform to these norms.

These issues transfer to how we classify neural data. Our methods of classifying data could make use of the right ontologies, but it's also possible that they will perpetuate what Dembroff (2018) calls *ontological oppression*—or "when […] social kinds (or the lack thereof) unjustly constrain (or enable) persons' behaviors, concepts, or affect due to their group

---

[2] Buolamwini also found that these algorithm performed poorly on famous photographs of influential black women—misclassifying or even failing to classify Oprah Winfrey, Michelle Obama, Shirly Chisolm, Sojourner Truth, and so on.

membership" (5). CL-DBS systems are well-positioned to implement or (more likely) perpetuate these kinds of oppressions. CL-DBS classification algorithms are tasked with detecting neural states that correlate with categories of disease. These systems may not "see" us the way we want to be seen.

Imagine, for example, a CL-DBS system for depression that is unable to detect depression states in people of color, or a system for epilepsy that produces high rates of false positives in women. Keeping users in the loop could make these problematic outcomes less likely by making sure that models are trained against as many different kinds of people as possible. But it's not clear if this will be enough. In each patient population, device engineers must be careful about what ontological categories they identify and make sure they do not reinforce pre-existing injustices. It is possible, however, that any ontology CL-DBS employs will be exclusive or marginalizing—but this only means designers will need be even more careful about determining which categories their analyses recognize (and which they ignore).

Stimulation and Feelings of Agency

CL-DBS' ability to determine when to stimulate presents its own set of challenges for user agency. Klein et al. (2016) conducted a series of interviews with people using DBS for psychiatric conditions (MDD and OCD) to hear their experiences with OL-DBS and gauge their concerns about CL-DBS. Participants expressed worries about both adaptive and BCI-triggered systems. Several users reported losing a sense of their own agency at times with OL-DBS—or what what Klein et al. (2016) termed "ambiguous agency" (5). These users sometimes find it difficult to determine if their actions are truly theirs. In one participant's words, "I've begun to wonder what's me and what's the depression, and what's the stimulator" (5). Participants predicted that these feelings of ambiguity would be "amplified in patients who receive closed-loop systems, if system maintained them in a relatively constant state of subjective well-being, and prevented them from experiencing a 'normal' range of emotion (e.g. feeling sadness at a funeral)" (4). These predictions seem apt. After all, OL-DBS users know (or should know) that their stimulator is running constantly unless they turn it off manually, so at the very least they know when the system is stimulating and when it isn't. An adaptive DBS user may not have the same information available to them: the system's algorithm determines when it applies stimulation, and when that happens may not be clear to the user. This may cause problems if users are having trouble determining if their behaviors are the result of stimulator activity or their condition. Several participants also worried that someone might be able to "game" their adaptive CL-DBS to trigger stimulation manually. If, for example, a system can detect sadness and stimulate in response, the user might deliberately make themselves sad to trigger stimulation.

We saw (in §2) that in some ways we are always "kept in the loop" to some degree insofar as users have the power to push back against the system. Even in loops where CL-DBS makes decisions without much volitional control, the user can still opt out of treatment, ask their clinicians to adjust stimulation parameters, or seek support from related others who have their best interests in mind. But even still, we should ask how much control users should have over calibrating the system—initially and over the lifetime of the treatment. That is, CL-DBS is tasked with making decisions on the user's behalf, it still remains to be seen what role users will play in training, tuning, and correcting algorithms to behave the ways they want them to behave. Further, the problems that arise in other parts of the closed-loop—problems with signal recording and classification—are likely *cumulative*. That is, decisions made about what to record and how to interpret these recordings will necessarily determine how and when stimulation is applied. If classification algorithms make use of ontologies that don't always map onto a person's experiences, then those algorithms will likely decide to apply stimulation in ways that do not align with their goals.

We should be able to see, now, that CL-DBS users face choices about what role larger social institutions will (or get to) play in their decisional loops. Conversely, they will also face choices about how their actions (with regard to CL-DBS) contribute to these institutions. Each component of the closed-loop system—sensing, analysis, and stimulation—comes with its own tensions and implications for agency. This conclusion should not feel foreign or non sequitur. After all, these tensions lie in the background of how we use technology more broadly, and many of us are reminded of them when we agree to End-User License Agreements (EULAs) for online services, post on social media platforms, receive targeted advertisements, and so on. If I agree to give a smartphone app access to data from my social media account, for example, it's possible that the app's developers will collect and analyze these data, sharing insights about it with third parties. CL-DBS users will likely be confronted with similar challenges for their agency, and this means they're rendered vulnerable to a variety of harms and disenfranchisements. Disabled and neurodivergent people who will rely on CL-DBS to backup and support their agency—the kind of backup and support able and neurotypical people can take for granted—may have difficulty trusting CL-DBS.

Maria Lugones gives us a useful framework for understanding these vulnerabilities and their contexts. In her *Peregrinajes* (2003), she distinguishes between two epistemic roles with respect to social and political action: the *strategist* and the *tactician*. "Theoreticians of society and politics," she explains, "have often conceived of themselves as perched up high, looking at or making up the social from a disengaged position" (207). Such theoreticians take themselves to be able to strategize from afar without joining the fray. These strategists contrast themselves with "subjected subjects assumed to negotiate daily survival myopically from within the concreteness of body-to-body engagement" (207)—or, the oppressed who resist oppression. These "subjected subjects" are not able to strategize from above. They are only able to employ tactics on the ground. They operate from what Lugones calls the streetwalker's perspective.

Lugones insists, however, that the strategist's supposed independence from the concrete is an illusion. A myriad of institutions and practices backup and support their agency in ways that are near-invisible to the strategist, but all-too-apparent to the tactician. After all, the tactician cannot count on the institutions to support them: these institutions are set against them. CL-DBS users will likely be more like the tactician than the strategist. CL-DBS systems will, in some ways, open their decisional loops to so many others: individuals, communities, institutions. Many of these parties, these supposed "stakeholders," will be nondisabled and neurotypical, benevolent but detached.[3] Stakeholders who aren't end-users, perhaps, are the strategists here: creating and modifying devices, building and shifting institutions to support these devices, and operating under presumptions generated without on-the-ground knowledge.

We started this section by asking, "Can users trust CL-DBS to be in the loop?" Now we can ask a similar question, "Can users trust CL-DBS, and the institutions that make CL-DBS possible, to back them up?" When we talk in terms of CL-DBS "being in the loop," we foreground issues that arise when CL-DBS takes control from the user or gives them more control. When we think about CL-DBS "backing up" their users, we foreground the issues that could arise in virtue of how CL-DBS is linked to interpersonal, academic, medical, corporate, and broader societal structures. In the next section, I'll suggest that users can trust CL-DBS systems if device designers put several safeguards in place.

---

[3] Admittedly, I have cherry-picked the most optimistic scenario.

5. Recommendations

Below, I propose three possible safeguards device manufacturers ought to place on CL-DBS systems. This is not an exhaustive list by any means, but it should serve as a useful starting point. First, device manufacturers, engineers, and clinicians must ensure that the CL-DBS operates transparently. This certainly means that study participants and patients must receive enough education to consent to research or therapy, as they are already required to. But the landscape of these technologies—what data they collect, what insights they arrive at, and the treatments they administer—is shifting constantly. Informed consent, then, must be ongoing, relational, and exploratory (Klein and Ojemann 2016). These goals may not be attainable without adding interfaces that give users the means to access information about what data their closed-loop system collects, how it analyzes that data, and how it determines when to stimulate. These goals, further, will likely require clinicians and device manufacturers to be even more deliberate about patient education than they might have been in the past.

Second, technicians and clinicians will likely need to give end-users the option of playing a deeper role in the setup and ongoing maintenance of their implant. This deeper role may not come in the form of giving users more control over how their stimulator operates on a day to day level, as too much control could be become burdensome. Instead, technicians could build systems that give users easier methods of opting out of data collection/sharing, correcting classification algorithms, and correcting stimulation misapplications. Determining how to design interfaces like these in ways that are unobtrusive, easy to use, and effective, should be (and likely will be) the subject of future research in the field (Wolkenstein, Jox, and Friedrich 2018).

Finally, device manufacturers must ensure that they, and the devices they manufacture, are trustworthy (Kerasidou 2016). Their efforts must go beyond lowering the risk of adverse effects or malfunctions. Their efforts must also go beyond increasing the usability of their devices and decreasing compliance with treatment. Their efforts must go beyond reading and heeding ethics guidelines (Johnsson et al. 2014). That is, manufacturers must take care to ensure that their on-device algorithms and external services account for and counteract possible injustices. How this looks will certainly depend on the algorithms used for classification and actuating stimulation. In the case of ANNs, these measures might take the form of using carefully-selected training data to counteract cultural biases and building out a rigorous process of auditing models. In the case of any algorithm, however, deep engagement with the communities of users is vital. In academic research on CL-DBS, a community-based participatory research (CBPR) model is a good start (Israel et al. 1998; Shore et al. 2008). Even further, however, diversifying teams of researchers and engineers—intellectually and culturally—could ensure that these teams are able to anticipate and address possible injustices that would normally go ignored.

6. Conclusion

We have seen that while CL-DBS may prove to be a useful intervention for a variety of conditions, the switch from open-loop to closed-loop may impact user agency in complicated ways. I argued that users are brought into or excluded from the decisional loops created by CL-DBS—through what the system records, how those recordings are analyzed, and how the system applies stimulation. Each of these components of the closed-loop system have the potential to complicate user agency by changing the topology of the user's decisional loop to include an entire cohort of potentially supportive others—systems, organizations, institutions, and so on.

This raised the question: can users trust CL-DBS to act on their behalf? Crucially, I warned about the possibility that each of these components, if designed carelessly, could subject users to injustice. I argued that device manufacturers must evaluate how each of the closed-loop system's components is implemented, determine how users are bought into or left out of the loop in virtue of these components, and determine if they influence, support, or stifle user's agency in morally salient ways. Finally, I gave a series of recommendations for how researchers and manufacturers can prevent these issues. Moving forward, each stakeholder in the creation of CL-DBS must ensure that each component of CL-DBS will back users up. If they fail to recognize how CL-DBS can fail its users, they risk bringing forces of oppression into their users' bodies.

7. References

Arlotti, Mattia, Sara Marceglia, Guglielmo Foffani, Jens Volkmann, Andres M Lozano, Elena Moro, Filippo Cogiamanian, et al. 2018. "Eight-hours adaptive deep brain stimulation in patients with Parkinson disease." *Neurology* 90 (11): e971–e976.

Barclay, Linda. 2000. "Autonomy and the social self." In *Relational Autonomy*, edited by Catriona Mackenzie and Natalie Stoljar. Oxford: Oxford University Press.

Baylis, Françoise. 2011. ""I Am Who I Am": On the Perceived Threats to Personal Identity from Deep Brain Stimulation." *Neuroethics* 6 (3): 513–26.

Bina, Robert W, and Jean-Phillipe Langevin. 2018. "Closed Loop Deep Brain Stimulation for PTSD, Addiction, and Disorders of Affective Facial Interpretation: Review and Discussion of Potential Biomarkers and Stimulation Paradigms." *Frontiers in Neuroscience* 12 (May): 305.

Bonaci, T, R Calo, and H J Chizeck. 2014. "App stores for the brain: Privacy & Security in Brain-Computer Interfaces." *2014 IEEE International Symposium on Ethics in Engineering, Science, and Technology (ETHICS)*, 1–7.

Bonaci, T, J Herron, C Matlack, and H J Chizeck. 2014. "Securing the exocortex: A twenty-first century cybernetics challenge." *2014 IEEE Conference on Norbert Wiener in the 21st Century (21CW)*, 1–8.

Brown, Timothy Emmanuel. Forthcoming. "Building Intricate Partnerships with Neurotechnology: Deep-Brain Stimulation and Relational Agency." *International Journal of Feminist Approaches to Bioethics*.

Brown, Timothy, Margaret C Thompson, Jeffrey Herron, Andrew Ko, Howard Chizeck, and Sara Goering. 2016. "Controlling our brains a case study on the implications of brain-computer interface-triggered deep brain stimulation for essential tremor." *Brain-Computer Interfaces*, September, 1–6.

Buller, Tom. 2011. “Neurotechnology, Invasiveness and the Extended Mind.” *Neuroethics* 6 (3): 593–605.

Buolamwini, J. 2018. “Gender shades: Intersectional accuracy disparities in commercial gender classification.” In *Conferenceon Fairness, Accountability, and Transparency*, 1–15.

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

Dembroff, R. 2020. “Real Talk on the Metaphysics of Gender.” *Philosophical Topics*.

Dubiel, Helmut. 2009. *Deep Within the Brain*. Living with Parkinson’s Disease. Penguin.

Fakhar, Kaihan, Erin Hastings, Christopher R Butson, Kelly D Foote, Pam Zeilman, and Michael S Okun. 2013. “Management of Deep Brain Stimulator Battery Failure: Battery Estimators, Charge Density, and Importance of Clinical Symptoms.” *PLoS ONE* 8 (3): e58665.

Faraut, Mailys C M, April A Carlson, Shannon Sullivan, Oana Tudusciuc, Ian Ross, Chrystal M Reed, Jeffrey M Chung, Adam N Mamelak, and Ueli Rutishauser. 2018. “Dataset of human medial temporal lobe single neuron activity during declarative memory encoding and recognition.” *Scientific Data* 5 (1): 279.

Freund, Michelle, Anna Taylor, Cathy Ng, and A Roger Little. 2018. “The NIH NeuroBioBank: creating opportunities for human brain research.” *Handbook of Clinical Neurology* 150 (January): 41–48.

Geller, Eric B, Tara L Skarpaas, Robert E Gross, Robert R Goodman, Gregory L Barkley, Carl W Bazil, Michael J Berg, et al. 2017. “Brain-responsive neurostimulation in patients with medically intractable mesial temporal lobe epilepsy.” *Epilepsia* 58 (6): 994–1004.

Glannon, W. 2009. “Stimulating brains, altering minds.” *Journal of Medical Ethics* 35 (5): 289–92.

Goering, Sara. 2014. “Is it Still Me? DBS, Agency, and the Extended, Relational Me.” *AJOB Neuroscience* 5 (4): 50–51.

Goering, Sara, Eran Klein, Darin D Dougherty, and Alik S Widge. 2017. “Staying in the Loop: Relational Agency and Identity in Next-Generation DBS for Psychiatry.” *AJOB Neuroscience* 8 (2): 59–70.

Guo, Anhong, Ece Kamar, Jennifer Wortman Vaughan, Hanna Wallach, and Meredith Ringel Morris. 2019. “Toward Fairness in AI for People with Disabilities: A Research Roadmap.” *arXiv.org*, July.

Habets, Jeroen G V, Margot Heijmans, Mark L Kuijf, Marcus L F Janssen, Yasin Temel, and Pieter L Kubben. 2018. “An update on adaptive deep brain stimulation in Parkinson’s disease.” *Movement Disorders* 33 (12): 1834–43.

Herron, J, and H J Chizeck. 2014. "Prototype closed-loop deep brain stimulation systems inspired by Norbert Wiener." In *2014 Ieee Conference on Norbert Wiener in the 21st Century (21CW)*.

Herron, Jeffrey A, Margaret C Thompson, Timothy Brown, Howard J Chizeck, Jeffrey G Ojemann, and Andrew L Ko. 2016. "Chronic electrocorticography for sensing movement intention and closed-loop deep brain stimulation with wearable sensors in an essential tremor patient." *Journal of Neurosurgery*, November, 1–8.

Hoang, Kimberly B, Isaac R Cassar, Warren M Grill, and Dennis A Turner. 2017. "Biomarkers and Stimulation Algorithms for Adaptive Brain Stimulation." *Frontiers in Neuroscience* 11 (October): 117.

Hoang, Kimberly B, and Dennis A Turner. 2019. "The Emerging Role of Biomarkers in Adaptive Modulation of Clinical Brain Stimulation." *Neurosurgery* 6 (3): 15.

Houston, Brady, Margaret Thompson, Andrew Ko, and Howard Chizeck. 2019. "A machine-learning approach to volitional control of a closed-loop deep brain stimulation system." *Journal of Neural Engineering* 16 (1): 016004.

Israel, Barbara A, Amy J Schulz, Edith A Parker, and Adam B Becker. 1998. "Review of Community-Based Research: Assessing Partnership Approaches to Improve Public Health." *Annual Review of Public Health* 19 (1): 173–202.

Jobst, Barbara C, Ritu Kapur, Gregory L Barkley, Carl W Bazil, Michel J Berg, Gregory K Bergey, Jane G Boggs, et al. 2017. "Brain-responsive neurostimulation in patients with medically intractable seizures arising from eloquent and other neocortical areas." *Epilepsia* 58 (6): 1005–14.

Johnsson, Linus, Stefan Eriksson, Gert Helgesson, and Mats G Hansson. 2014. "Making researchers moral: Why trustworthiness requires more than ethics guidelines and review." *Research Ethics* 10 (1): 29–46.

Kellmeyer, Philipp. 2018. "Big Brain Data: On the Responsible Use of Brain Data from Clinical and Consumer-Directed Neurotechnological Devices." *Neuroethics* 90 (May): 1–16.

Kellmeyer, Philipp, Thomas Cochrane, Oliver Müller, Christine Mitchell, Tonio Ball, Joseph J Fins, and Nikola Biller-Andorno. 2016. "The Effects of Closed-Loop Medical Devices on the Autonomy and Accountability of Persons and Systems." *Cambridge Quarterly of Healthcare Ethics* 25 (4): 623–33.

Kerasidou, Angeliki. 2016. "Trust me, I'm a researcher!: The role of trust in biomedical research." *Medicine, Health Care and Philosophy* 20 (1): 43–50.

Klein, Eran, Sara Goering, Josh Gagne, Conor V Shea, Rachel Franklin, Samuel Zorowitz, Darin D Dougherty, and Alik S Widge. 2016. "Brain-computer interface-based control of closed-loop brain stimulation: attitudes and ethical considerations." *Brain-Computer Interfaces* 3 (3): 140–48.

Klein, Eran, and C S Nam. 2016. “Neuroethics and brain-computer interfaces (BCIs).” *Brain-Computer Interfaces* 3 (3): 123–25.

Klein, Eran, and Jeffrey Ojemann. 2016. “Informed consent in implantable BCI research: identification of research risks and recommendations for development of best practices.” *Journal of Neural Engineering* 13 (4): 043001.

Kraemer, Felicitas. 2011. “Me, Myself and My Brain Implant: Deep Brain Stimulation Raises Questions of Personal Authenticity and Alienation.” *Neuroethics* 6 (3): 483–97.

Lugones, María. 2003. *Pilgrimages / Peregrinajes*. Theorizing Coalition Against Multiple Oppressions. Rowman & Littlefield.

MacLeod, Haley, Cynthia L Bennett, Meredith Ringel Morris, and Edward Cutrell. 2017. *Understanding Blind People’s Experiences with Computer-Generated Captions of Social Media Images*. New York, New York, USA: ACM.

McLeod, Carolyn. 2002. *Self-Trust and Reproductive Autonomy*. MIT Press.

Mello, Michelle M, and Leslie E Wolf. 2010. “The Havasupai Indian tribe case–lessons for research involving stored biologic samples.” *The New England Journal of Medicine* 363 (3): 204–7.

Miller, Kai J. 2019. “A library of human electrocorticographic data and analyses.” *Nature Human Behaviour* 3 (11): 1225–35.

Neumann, W-J, J Huebl, C Brücke, L Gabriëls, M Bajbouj, A Merkl, G H Schneider, B Nuttin, P Brown, and A A Kühn. 2014. “Different patterns of local field potentials from limbic DBS targets in patients with major depressive and obsessive compulsive disorder.” *Molecular Psychiatry* 19 (11): 1186–92.

Nichols, Lisa, Michelle Freund, Cathy Ng, Alice Kau, Melissa Parisi, Anna Taylor, David Armstrong, et al. 2014. “The National Institutes of Health Neurobiobank: A Federated National Network of Human Brain and Tissue Repositories.” *Biological Psychiatry* 75 (12): e21–e22.

Peng, Hanchuan, Michael Hawrylycz, Jane Roskams, Sean Hill, Nelson Spruston, Erik Meijering, and Giorgio A Ascoli. 2015. “BigNeuron: Large-Scale 3D Neuron Reconstruction from Optical Microscopy Images.” *Neuron* 87 (2): 252–56.

Rao, Rajesh PN. 2019. “Towards neural co-processors for the brain: combining decoding and encoding in braincomputer interfaces.” *Current Opinion in Neurobiology* 55 (April): 142–51.

“Rethinking Neuroethics in the Light of the Extended Mind Thesis.” 2007. *The American Journal of Bioethics* 7 (9): 3–11

Schermer, Maartje. 2011. "Ethical Issues in Deep Brain Stimulation." *Frontiers in Integrative Neuroscience* 5 (May).

Shih, Ludy C, Kathrin LaFaver, Chen Lim, Efstathios Papavassiliou, and Daniel Tarsy. 2013. "Loss of benefit in VIM thalamic deep brain stimulation (DBS) for essential tremor (ET): How prevalent is it?" *Parkinsonism & Related Disorders* 19 (7): 676–79.

Shore, Nancy, Kristine A Wong, Sarena D Seifer, Jessica Grignon, and Vanessa Northington Gamble. 2008. "Introduction to Special Issue: Advancing the Ethics of Community-Based Participatory Research." *Journal of Empirical Research on Human Research Ethics* 3 (2): 1–4.

Smeding, H M M, A E Goudriaan, E M J Foncke, P R Schuurman, J D Speelman, and B Schmand. 2007. "Pathological gambling after bilateral subthalamic nucleus stimulation in Parkinson disease." *Journal of Neurology, Neurosurgery & Psychiatry* 78 (5): 517–19.

Sterling, Robyn L. 2011. "Genetic Research among the Havasupai: A Cautionary Tale." *AMA Journal of Ethics* 13 (2): 113–17

Thompson, Margaret C, and J Herron. 2016. "Demonstration of a Chronic Brain-Computer Interface using a Deep Brain Stimulator." In *Th International Brain Computer Interface Bci Meeting*. Asilomar.

Weiss, Daniel, and João Massano. 2018. "Approaching adaptive control in neurostimulation for Parkinson disease: Autopilot on." *Neurology* 90 (11): 497–98.

Wilson, Benjamin, Judy Hoffman, and Jamie Morgenstern. 2019. "Predictive Inequity in Object Detection." *arXiv.org*, February.

Wolkenstein, Andreas, Ralf J Jox, and Orsolya Friedrich. 2018. "BrainComputer Interfaces: Lessons to Be Learned from the Ethics of Algorithms." *Cambridge Quarterly of Healthcare Ethics* 27 (4): 635–46.

Yuste, Rafael, Sara Goering, Blaise Agüera y Arcas, Guoqiang Bi, Jose M Carmena, Adrian Carter, Joseph J Fins, et al. 2017. "Four ethical priorities for neurotechnologies and AI." *Nature News* 551 (7679): 159–63.

Deep Brain Stimulators and the Aesthetics of Agency

Abstract

Some users of deep-brain stimulation (DBS) report unwelcomed changes to their sense of self and their interpersonal relationships post-operatively. These reports have been the catalyst for arguments over the nature of these changes, their moral salience, and their broader significance for theories of self. But where commentators employ various theories of self to account for users experiences of DBS, I argue that the *aesthetic dimensions* of these users' experiences are left underexplored. In this piece, I propose a framework—drawing from everyday aesthetics and the aesthetics of games—for understanding users' testimonies in terms of DBS's impact on the creative act of exercising agency. I close with recommendations for how empirical research and device design can better track and account for the aesthetics of exercising agency.

1. Introduction

Devices that apply electrical stimulation to the central nervous system are becoming a more common part of medical practice. One such technology, Deep–Brain Stimulation (DBS), is already an FDA-approved treatment for the symptoms of Parkinson's Disease (PD), Essential Tremor (ET), and Epilepsy—with experimental use for psychiatric conditions like Major Depressive Disorder (MDD), Obsessive-Compulsive Disorder (OCD). These systems consist of a signal generator, often implanted in the user's chest, that sends electrical impulses to electrodes implanted in carefully-selected areas of the user's brain.

Stimulation sometimes comes with a number of psychological side-effects, with some DBS users reporting unwelcomed changes to their sense of self and their interpersonal relationships after the operation (Agid et al. 2006; Schüpbach et al. 2006). These experiences are often described as psychosocial maladaptations (Schüpbach et al. 2006; Gilbert 2012; Mecacci and Haselager 2014; Dos Santos et al. 2017). Several commenters argue that these changes constitute a threat to users insofar as it makes users feel alienated (Kraemer 2011), changes their identities (Glannon 2009, 2014), or diminishes their competencies for (autonomous) agency (Baylis 2011; Mackenzie and Walker 2015). To the latter-most worry, some commenters attend to the potential for DBS (and adjacent neurotechnologies) to bring about changes to the user's sense of agency (Haselager 2012; Lipsman and Glannon 2013; Gallagher 2018)—or, the sense that they are the source of their intentions and/or actions that result from those intentions. Many of these worries are driven by the intuition that DBS can change fundamental features of the user's selfhood—their behaviors, beliefs, capacities, and lived-experiences—in ways that render the user vulnerable to harm. That is, DBS may change the user fundamentally, and in ways that are difficult to recognize, understand, accept, or reverse.[1] Many authors note that several of the changes brought about by DBS have direct implications for informed consent (Giordano 2015; Klein and Ojemann 2016; Witt 2017), continuity of care (Rossi, Okun, and Giordano 2014), or personal responsibility (Klaming and Haselager 2013; Sharp and Wasserman 2016; Brown, Moore, et al. 2016)—e.g., when people with PD report pathological gambling (Smeding et al. 2007) or mania (Leentjens et al. 2004) while using DBS.

Some changes, however, are not as easy to characterize in terms of some direct encroachment on the user's autonomous agency or identity. Some users may feel, for example, like they've changed in ways they still have yet to grow accustomed to (Haan et al. 2015), or that it's not clear if they're the source of their own actions (Klein et al. 2016), or that it's difficult

---

[1] It is not clear that the effects of DBS, and other implanted neurotechnologies, are reversible. The very act of implanting electrodes alongside neural tissue causes a "micro lesioning effect," where tissue is damaged and (at least in the case of motor disorders) can bring about temporary relief from symptoms (Cersosimo et al. 2009). It is not clear if micro lesioning caused by implantation (or explantation) play a role in long-term, irreversible neurological changes.

imagine life without symptoms of their condition (Gilbert 2012), or (again) that their actions feel artificial (Schüpbach et al. 2006). In many of these cases, DBS worked the way it was supposed to, bringing about changes that the user wanted (or needed), but the user's everyday experiences have changed in ways that are disconcerting.

How, then, do we characterize the experiences of a person who feels alienated or artificial (Kraemer 2013), but hasn't undergone a change in selfhood that undermines their autonomous agency directly? These experiences, even if they seem less salient, can easily become deep difficulties for users who are coming to terms with life with their DBS. After all, if a person's actions feel inauthentic after using DBS, this experience of inauthenticity may complicate how users experience and express their agency. Experiences like these, I propose, are better described in aesthetic terms—that is, in terms of how DBS impacts the creative act of being an agent, and how the agent's appreciation and appraisal of the experience of being an agent.

I'll defend my proposal as follows. I'll begin by attending DBS users' reports of feeling alienated, and Felicitas Kraemer's interpretation of these cases—on which these experiences are the result of DBS changing the user in ways that make them feel distant from or closer to their authentic self. This will serve as a springboard for thinking about user testimonies in terms of the aesthetics of being an agent. That is, I'll argue that what a person means when they say they feel inauthentic is easier to understand when we make sense of how DBS changes the aesthetic features of being an agent. I'll, then, give an account of what we could possibly mean by "aesthetic" and how we can think of agency in terms of the aesthetic. Finally, I'll conclude by exploring the implications of my view for future research on DBS.

2. DBS Users and Felt-Inauthenticity

Let's begin by taking stock of users' reports of their experiences with DBS. We'll find that several users report changes to experiences of their agency, and several of these are disruptive in ways that make agency more difficult to exercise. Agid et al. (2006) conducted unstructured, open interviews with 29 participants using DBS for PD. They report that participants experienced a variety of changes to their self-perceptions after surgery. One participant remarked that he felt "unable to recognize [himself] without the problems he had before" (411). Another remarked that their symptoms gave them a purpose in life, and that DBS took that purpose away and left him feeling aimless: "before stimulation, every day was a struggle. Now, I miss the time when I used to fight. Nowadays, I'm like a soldier when the war is over, there's no longer anything to fight against. My life seems empty" (411).

Agid et al also observed what they call "negative anticipation," or the suspicion that their symptoms will return. This negative anticipation, one participant describes, is a kind of habit that they can't leave behind: "Even though I know my motor state is perfectly normal, I just can't get rid of the thoughts I used to have, when I was ill. I still have the same reflexes to initiate walking and the rituals before taking … medicine" (411). Some participants described feeling artificial or subject to external influence. One participant remarked feeling like "an electronic doll," "'Robocop'," or "under remote control" (411).

Further, several participants noticed changes in their interpersonal relationships—some feel unrecognized, overprotected, or rejected by their friends, family members, or caretakers. "During all these years, I let myself be treated like a child because I didn't have the strength to fight back," as one participant explained poignantly, "That period is over now. I'm going to claim my old place back again. Now I want to have control over my life again, recover the life I had before PD" (412). Speculating on the possibility that some of these difficulties are the result of DBS "modifying the patient's personality" directly, Agid et al. conclude that this is likely not the case. Instead, they argue it is more likely that these users simply had trouble adjusting to life without many of the symptoms of PD, symptoms they had little power to change or had grown accustomed to.[2]

De Haan et al also conducted a set of semi-structured interviews with 18 participants using DBS for OCD. Their goal was to determine if DBS, at least when used to treat OCD, causes the kinds of personality changes or feelings of inauthenticity noted in the Agid et al. (2006) study. Their participants also experienced self-perception changes post-op. But it is important to note that several participants did not feel they had changed as a person at all. Instead, some of these participants preferred to talk in terms of how their actions have changed. As one participant (P3) recounts, "I did not become a different person, no. […] Just that the traits that I had, have become a little stronger, so to say. That kind of thing."

Several of them explained that they were "more themselves" as a result of being compulsion free (or even just better able to control their compulsions). As one participant, Lea (P20), explained, "I do feel that I am indeed more Lea. I still have my compulsions, but the compulsions used to rule my life. The compulsions were in charge, ruled over me. And now I am a bit in charge again over myself" (11). For another (P5), DBS was the means to become who they used to be: "I think I am coming back to my own self. If you are in a fixed pattern for decades and through such an operation you get back to your own self more, then that is a good thing" (11).

---

[2] Gilbert (2012) calls this phenomenon the "burden of normality." I avoid this language because of possible ableist readings of the words "burden" and "normal" in this context.

Others, however, noted several more direct changes. In the case of side-effects (increased libido in particular) some were comfortable with these changes, while others were dissatisfied enough to ask for adjustments to their stimulation settings. Further, several participants felt they became more assertive after surgery, and a subset of them took this to be a sign that they were able to express themselves more now without being limited by OCD. Some, however, were uncomfortable with how assertive they became. As one such participant notes, "I used to be so insecure that I did not dare to say anything, and now I do say it. Sometimes a bit too much" (13).

Several participants felt they had yet to *get used to* themselves or *find out who they were* as DBS users—or, perhaps more saliently, as people who've found a way to manage their compulsions. One participant describes this as a process of self-(re)discovery, "You are actually starting to find out who you really are, right? Because it was gone, yes. […] Because of the compulsions and also because of the bullying. Because if you hear every day that you are a loser, an asshole. Five years long. Yes, then you will start to believe in that, right?" (16) Finally, some participants described feeling ambiguous about which behaviors represent them best or belong to them more, with some participants relying on friends and family to remember what they were like before OCD. As one participant (P13) remarks, "my oldest sister she noticed that too: she saw me at a few birthday parties, and then at some point she says, because she knows me the longest, she is six years older than I am […] so she has experienced me from an early age. And she said at a certain point: I see the old Peter from the old days again."

Finally, Klein et al (Klein et al. 2016) conducted a series of interviews with participants using DBS for Major Depressive Disorder (MDD) and OCD to gather their impressions of next-generation, automated versions of DBS. They identified many of the same themes found in Agid et al and De Haan et al's respective studies. Several of them explained that DBS allowed them to return become who they were when OCD and MDD weren’t an issue. Again, one participant notes, “I still like the same things. I don’t have, like, different values or anything. I just enjoy things more. I’m me without depression” (5). Several participants, however, were inclined to talk about DBS in terms of returning to a more “authentic” version of themselves, or feeling more like themselves. One participant (KI3) claimed DBS allowed them to “access parts of [themselves] that gave [them] enjoyment, that [they weren’t] able to when [they didn’t] have the surgery” (5).

Some, however, describe more a disconcerting experience: not being able to distinguish between their own behaviors and behaviors influenced by DBS. Klein et al call this phenomenon *ambiguous agency*. As one participant explains: “There are parts of this where you just wonder how much is YOU any more, and you wonder kind of, ‘How much of it is my thought pattern? How would I deal with this if I didn’t have the stimulation system?’ You kind of feel artificial” (5). Other participants noticed that family members would talk about their behaviors differently. “Your feelings don’t get attributed to *you*,” one participant explains, “They’re all due to the device.” Another participant describes, “there’s people in my family that sometimes question, you know, how much of it is me any more and how much of me is, you know, bein’ programmed. So that’s been a, that’s a hard thing to deal with sometimes.”

DBS, we can see, has the potential to deeply complicate users’ experiences. In some cases, DBS complicated identity assertion or formation (“I’ll need to get used to myself”). In other cases, DBS complicated the user’s actions and deliberations (“Am I the one acting, or is it my neurostimulator?”). In all of these cases, however, it seems that DBS played a perceptible role in changing the experience of exercising agency—that is, the experience of deliberating over choices, forming intentions, acting on them, and evaluating the consequences of their actions.

Whether the user is overwhelmed (or excited) by their increased libedo, or they're unable to determine if they're the source of their own actions, these cases demonstrate just how much the experience of agency matters. In the next section, we'll focus on people who report that they feel inauthentic or isolated after using DBS. We'll find that within the debate over how to characterize these feelings, there are tools we can use to make sense of how DBS changes users' experiences of agency.

3. Feelings of authenticity, experiences of agency

Agid et al. (2006) use the language of maladaptation or maladjustment to describe users difficult experiences using DBS. Felicitas Kraemer (2011), however, argues that this language "fails to capture an important dimension of the patients' experience, namely, the subjective experience of authenticity and alienation" (488). To overlook these experiences would be to overlook a means of gauging and understanding the successes and failures of DBS therapy. Authenticity, on her view, is the felt mental state indicated by language like, "I feel (more) like myself." Inauthenticity, and alienation, are just the opposite of authenticity: the felt mental state indicated by the language, "I don't feel like myself." Kraemer notes (following Lionel Trilling and Charles Taylor) that authenticity and alienation have a normative dimension: we ought to strive for authenticity and avoid alienation. And so the experience of using DBS, in cases like the ones we've considered above, is an empowering or difficult process of coming to grips with how authentic or alienated they feel while using DBS. And so "when evaluating the ethical … implications of behavior changes that result from DBS, the subjective state of felt-authenticity and felt-alienation should be taken into consideration" (Kraemer 2011, 496) so that we can better attend to how each patient's experience using DBS may be a process of "recognizing, exploring, and enacting what they regarded as their 'true selves'" (Kraemer 2011, 496).

But while the language of authenticity may be useful for capturing salient features of DBS users' testimonies, several critics worry that describing users' experiences in terms of their felt-authenticity or felt-alienation is to commit to an impoverished view of self-identity that makes it difficult to track a variety of important features of identity: the role relationships play in identity-formation, how identities shift from one social context to the next, and how identities inevitably change over time in response to life events. Francois Baylis (2011)—following Marya Schechtman (2010) and followed by Catriona Mackenzie & Mary Walker (2015)—employs a narrative account of identity to argue along these lines. DBS implantation, she argues, is not unlike other events (e.g., severe injury, relationship loss, etc.) that have an impact on our self-narratives.

After all, if therapy is successful, a DBS user will incorporate their new experiences into their narratives, for better or worse. Friends and family will, in turn, corroborate or refute their narratives—applying what Baylis calls an "equilibrium constraint" on them. In worse cases, a DBS user may feel uncertain or upset about the self-narratives available to them, unsatisfied with the narratives their friends and family make on their behalf, or uncomfortable "performing" any of these self-narratives (Baylis 2011, 2012). But these problems are not unique to DBS; they are a normal part of identity formation, expression, and (to the point) revision. Perhaps, then, we needn't talk in terms of "authentic" self-hood at all—instead we should talk in terms of unfolding dynamic, relationally-constructed, narratives.

Haan et al. (2017), however, worry that Baylis is too quick to dismiss authenticity-speak. After all, even if we should think of receiving DBS therapy as a life event among other life events, "DBS not just any life event" (5)—it can change the way a person expresses themselves, the way they perceive their own actions, and the way others interact with them. That is to say, DBS is significant enough to be salient in a way not captured by Baylis' version of the narrative identity view. That's not to say that these features are better captured by positing the existence of an authentic self. De Haan et. al. simply maintain that "given the prominent role of the brain in who we are, it makes sense to pay careful attention to what happens when it receives electrical stimulation" (5)—these experience are significant, even if not for the reasons Kraemer provides.

Haan et al. (2015) also note that "in general [their study's] participants did not change their attitudes and beliefs towards things, but rather their way of acting and interacting" (18). For several of the participants in the studies we've considered, it was difficult to determine if they (or the DBS) were the source of their own actions, if their actions really had changed, and/or if they were comfortable with how their behaviors changed. These difficulties seem to stem from how DBS impacts the experience of being an agent. Baylis might be right to say identity is more dynamic than an authenticity view would allow, but (at the very least) authenticity language captures some feature of users' experience of DBS that may not have been captured otherwise. It will be helpful, then, take a closer look at DBS' impact on self-hood in terms of how users' *experience of agency* change post-op.

Lipsman and Glannon (2013), to start, describe agency as "the capacity to initiate and execute plans of action" (468)—but they remind us exercising agency gives rise to lived-experiences of deliberating, choosing, and acting. The quality of these experiences can impact the outcome of expressions of agency: if a condition "impairs" a person's "affective and motivational states," it also "impairs [their] capacity to form and carry out intentions to act" (468). During a depressive episode, for example, a person may not have the desire to deliberate, may find some of their options unappealing, or may not feel like they can take control of themselves enough to act or take on a plan. Agency, at least in this case, seems bound up in the details of their experience of being an agent. And so Lipsman and Glannon argue that DBS users can view stimulation in a number of ways depending on the outcomes of therapy and their effect on their experiences of agency. They can view it as "a threat to [their] experience of being in control," "an enabling tool that has become integrated into [their] brain," or "something that enhances [their existence] through diminished affective and motivational capacities and ensures that [the user] is the genuine source of [their] actions" (469). For each possibility, however, DBS plays some role in modulating the user's experience of agency.

And so even if authenticity language leads us to a much too static view of self-hood—instead of the more dynamic views Schechtman, Baylis, Mackenzie & Walker defend—DBS users seem to reach for language related to felt-authenticity to describe their experiences with DBS (and the conditions it treats) too often to ignore. I am not too interested in finding or developing an alternative account of authentic selfhood. Instead, I want to suggest that we ought to attend to, and take seriously, how DBS users describe their lived-experiences of selfhood. If they find terms like authenticity, alienation, and ambiguity useful for describing their lived-experiences, then our theories should try to track why those terms are useful. We ought not dismiss them out of hand because of the metaphysical commitments they imply. After all, it may not matter in the clinical, interpersonal context if we can give a philosophical account of personhood that tracks their language.[3] As (Haan et al. 2015) write:

> For the clinical DBS practice, the important question is not the theoretical question of which definition of personal identity to adopt and the ensuing discussion whether these changes amount to changes in patients' identity. The important question is rather one that patients themselves struggle with, namely how to assess whether these changes are expressive of themselves or not. Patients' own assessment is of course central. (23)

---

[3] This also implies that we should stay open to other language users employ to describe their experience, even if their language diverges from our preferred accounts of selfhood.

I propose, then, that we shift the conversation so that we start from the presumption that authenticity-talk is a means of describing what it's like to be an agent, rather than an expression of a shift in identity simpliciter. This, perhaps, would give us a clearer picture of how DBS changes the way the user acts in ways that give rise to disconcerting experiences—including, but certainly not limited to, inauthenticity and alienation.

Erik Parens, in his *Shaping Our Selves*, makes use of this intuition as he evaluates the debates over using medical technologies for enhancement purposes. These debates—he notes, following Charles Taylor in his *Ethics of Authenticity*—are over what role technologies ought to play in how we live up to the *moral ideal of authenticity*, where by "it is my job as a human being to find my way of flourishing, of being true to myself" (48). He reminds us that "whether someone uses a scalpel to shape her body or a drug to shape her brain, she does so to shape her experience—to alter what it is like for her to be her" [31]. Restorative or enhancement technologies, on his characterization, "change what it feels like to be in the world," "affect how one is treated by others," and "improve [one's] experience of [one's] self" (33). And since these technologies are so personal, insofar as they have the potential to shape these intimate details of our experiences, our debates over the moral implications of using them feel as though we are "arguing for the value of our own existence, or justifying our way of being in the world" (47).

This is only to be expected. If these technologies are capable of shaping who we are—as moral, practical, and (as I'll argue) aesthetic agents—then argument over (im)permissible uses of these technologies will likely feel like an attack for some and an affirmation for others. These debates, then, have an urgency to them that pushes us to take a side, look at the issue through a single conceptual *lens*, approach decision-making from a single *stance*, and defend our chosen side against the opposing side. Parens encourages us to resist this tendency toward *monocularity*.

Instead, he insists that we should challenge ourselves to look at both sides as far as we can—that is, he urges us to adopt *binocularity*.

In the case of human enhancement, Parens distinguishes between those who embrace the *gratitude stance* and those who embrace the *creativity stance*. People in the gratitude stance, on one hand, accept, appreciate, and show gratitude for what we are given. By embracing gratitude, we "recognize that we have been thrown into being by forces we don't yet understand and that we didn't create" and this "feels worthy of celebration" (61). From this stance, we evaluate one another as objects "whose behaviors are determined by a staggeringly complex array of biological and social forces" (5). People in the creativity stance, on the other hand, embrace our collective ability to shape our environments and experiences—our ability to, in some ways, create ourselves. Even further, our creative and productive capacities give us "an ethical responsibility to creatively transform our selves and the world" (61). From this stance, we evaluate one another as subjects "who have the experience of choosing freely" (5).

Parens, again, insists that we ought not take either stance monocularly. Instead, we should try to look at our enhancement uses of technology through both lenses. This binocularity, he argues (crucially), is not just helpful in how we approach the debate, but in how we evaluate the people who use technologies for enhancement purposes. When we are able to evaluate one another first as objects (or thrown into the world with traits) and then as subjects (who are able shape themselves), we are able to make better sense of how people describe their experiences of using medical technologies for restorative and enhancement purposes. We can begin to see, however, that to take either stance is to appreciate and embrace some set of features of our agency, make moral or practical decisions as a function of the features we are congenial with, and evaluate how those decisions shape our lived-experiences. To take a stance is to recognize, appreciate, even *celebrate* that we've been "thrown into being" or that we are able to "creatively transform ourselves." But to use any of this language—of appreciation and celebration—implies that we bring *aesthetic* values into our moral (and practical) decision-making, evaluations, and debates about enhancement.

For some, my appeal to the aesthetic is a sign that I've gone off the rails. But surely some notion of the aesthetic has been lurking beneath the surface of Parens' discussion of authenticity (and Kraemer's reading of DBS users' reports). We can find evidence of this in Charles Taylor's *Ethics of Authenticity* (1992)—a text both Kraemer and Parens make use of. Taylor argues that the pursuit of self-fulfillment is not just a kind of amoral (or even immoral) self-indulgence. Instead, it is an attempt to live up to the moral ideal of authenticity. The very idea that we have original selves that we ought to strive to be, Taylor notes, implies that we must try to uncover or discover who we are—but this self-discovery only happens when we strive to convey ourselves to others. "This suggests right away," Taylor continues, "a close analogy, even a connection, between self-discovery and artistic creation" (61) insofar as "self-discovery requires poiesis, [or] making" (62). And while some argue that this is a kind of detached self-indulgence, Taylor gives us a picture of authenticity that tasks agents with both uncovering and conveying themselves ("truth to self") as well as navigating externally imposed rules ("intersubjective justice"). The moral ideal of authenticity is a project of both self-discovery and self-articulation—of both making sense of ourselves and allowing others to make sense of us.

In a way, Taylor's view is a bridge between Kraemer and Baylis' respective views. Kraemer asks us to attend to DBS impacts users' states of felt-authenticity, along with the role DBS plays in how users strive for authenticity broadly. Baylis reminds us that DBS therapies contribute to our identities in the same way other important experiences do: in their narrative and relational contexts. Taylor, however, reminds us to attend to both the experience of striving toward authenticity and to the external constraints on that striving.

I, however, want to draw our attention to the role DBS can play in our experiences of agency broadly. Authenticity and alienation, after all, are but two of the possible lived-experiences of agency (or, more broadly, selfhood) as it is modulated by DBS. The cases we considered in the previous section give us example of experiences that are more complicated than the language of felt-authenticity can capture. When users with OCD find that they need to "get used to themselves" or that they're uncertain if they've become too assertive, despite seeming satisfied with the outcomes of DBS therapy, their authenticity-language gives us at the very least a clue as to what their experiences are like. But if we want an even clearer picture of what those experiences are like, and if we want to make sense of other experiences that are not as well-captured through authenticity-language—if, perhaps, they'd prefer to talk about DBS the same way they would talk about using a tool (Brown, Thompson, et al. 2016)—we need to attend to the exploratory, productive, *aesthetic* features of the experience being an agent.

On my view, then, *aesthetic* frameworks will help us make sense of what neurotechnology users (and, likely, medical technology users broadly) mean when they use authenticity-language, ambiguity-language, striving-language, or other (as of yet unexplored) languages to describe how devices impact their experiences of being agents. DBS users are not just striving to discover themselves or to rediscover some prior form of themselves. They are also striving to produce and live with themselves after starting DBS therapy. In the next section, I give an overview of a relatively recent area of aesthetics that should give us tools to understand these experiences: the aesthetics of everyday life and the aesthetics of games.

4. The aesthetic dimension of agency

Above, I recounted a debate between Felicitas Kraemer and Francois Baylis over the legitimacy of using authenticity-language to describe the outcomes of using DBS. I argued that we should take authenticity-speak more seriously on the grounds that we ought to take users' descriptions of their experience seriously—provided that we have the tools to make sense of them. I also noted that a hint for how to make sense of user experiences sits in the background of the debates over enhancement technologies and their (potential) threats and benefits to authenticity: self-discovery and self-articulation are creative, artistic, *aesthetic* projects. I argued we would find the tools to make sense of users' experiences through aesthetic frameworks. In this section, we'll work through several views on the aesthetics of ordinary experience to locate these tools.

In common parlance, the term "aesthetic" often refers to the things we do and experience when we encounter artworks. When I view a painting in a gallery space, for example, I often operate under the (tacit or explicit) assumption that the purpose of observing art is to analyze, take in, appreciate, or otherwise allow myself to be moved—to pleasure or discomfort—by the painting. It is, however, difficult to give a completely satisfying account of what happens when we experience art, and so it's hard to give a set of necessary and sufficient conditions for what we can call 'aesthetic.' For the time being, however, we might be able to settle for an (almost) ostensive definition of the aesthetic in artistic contexts. Let's, then, think of the aesthetic as a set of responses given, attitudes held, judgments given, values considered, and emotions felt in certain contexts—paradigmatically, when we observe art.

But it seems we also respond, hold attitudes, judge, evaluate, and emote in similar ways outside of artistic contexts as well. We can say that a walk in the park, for example, can feel a certain way (harried, serene, and so on) in addition to our appreciation of the objects within that experience (the trees, walkway, the park's other visitors). Experiences like these have an "aesthetic dimension" that we can appreciate.[4] My suggestion, here, follows recent literature on "everyday aesthetics," attempts to account for how we participate in the aesthetic in ordinary contexts. We can think of John Dewey's view in *Art As Experience* (Dewey 2005) as an earlier example of everyday aesthetics.

Dewey begins with the assertion that if we want to understand how we experience art, or how the aesthetic operates, we need to understand how we experience life broadly. On his description, experience "occurs continuously" in such a way that "the experience had is inchoate" (35). But we are also, he argues, "live creatures" that are situated in and struggle against a dynamic environment. And through our struggles, we produce "rhythms" where "moments of fulfillment punctuate experience with rhythmically enjoyed intervals" (Dewey 1934, 18). These moments occur when "the material experienced runs its course to fulfilment" (35). The result is "*an* experience"—or an "integral experience"—where we undergo an experience profound enough to distinguish it from other experiences. *An* experience builds up tension, and ends in resolution, harmony, fulfillment, completion, or satisfaction. It has a *telos* that it reaches, that "is so rounded out that its close is a consummation" (35).

And so, on Dewey's view, our intentional actions and activities become the material of aesthetic experiences. As he writes: "a piece of work is finished in a way that is satisfactory; a problem receives its solution; a game is played through." And when the integral experience has

---

[4] When terms take the form "aesthetic x"—e.g., aesthetic experience, aesthetic appreciation, aesthetic property—it seems as though we've defined them in a circular or question-begging way. Francis Sparshott lampooned this phenomenon in an appendix of his *The Theory of the Arts* entitled, "Aesthetic This, Aesthetic That" (Sparshott 1982, 467–85). Below, I try my best to characterize the "aesthetic dimension" of agency without relying too much on these terms.

run its course, it is significant enough to remember. When we say “that meal *was* an experience,” for example, we might go on to remember it as “an enduring memorial of what food might be” (36). As such, Dewey recognizes how everyday life can rise to a level where it moves us—to the point where we can call the experience aesthetic—and he recognizes that our everyday activities can produce these aesthetic experiences.

But it seems that even experiences that don’t rise to the level of *an* experience also have an aesthetic dimension. Not all experiences flow from genesis to terminus; some salient experiences begin without our awareness, like when we don’t notice how tired we are after a long day, or when a long-time friendship has grown into a romance under our noses. Further, not all experiences resolve satisfactorily, if they resolve at all. Some problems go without a solution, and some tensions never release. Take, for example, the case of unrequited love: it is possible to love someone in a way that is unreciprocated and ongoing in a way that never resolves but is an appreciable source of tension. Think, also, of people using DBS for OCD who say they’re “getting used to” who they are when they use DBS. Will these users ever truly get used to their new way of being? Will this tension ever resolve in a way that qualifies as *an experience*? It may, but it also may not. But in either case, it seems possible to, or inexorable that we, appreciate these features of experience—even if they are or remain incomplete. I say *possible* because we appreciate integral experiences as we experience them (e.g., the satisfactory meal is not just satisfying in retrospect, it’s satisfying *as I eat it*). If something interrupts the integral experience, it might ruin the experience, but the experience happened none the less. I say *inexorable* because it seems, at least to me, that a large number of our experiences are of this sort. Sometimes we are in the middle of some experience, or we’ve pushed some ongoing drama to the side, or we know that whatever tensions exist may never resolve.

I want to suggest that experiences of being an agent can have an aesthetic dimension, and Dewey's view gives many of the tools for doing just that. Dewey, however, sets the standard for *an* experience so high—or gives criteria so narrow—that his view doesn't quite capture the aesthetic dimensions of experiences like the ones that accompany DBS use. My concern here tracks an on-going debate over the possibility of appreciating the aesthetics of ordinary things.

Tom Leddy (2010, 2012) argues, in a way that extends Dewey, that we appreciate the aesthetic qualities of ordinary things insofar as they strike us as (or become) extraordinary. To others [Irvin (2008); Saito (2008); Forsey (2013); and so on], however, it is possible (or even common) to evaluate the aesthetics of ordinary things on their own terms. On views like these, the content of everyday life is saturated with the aesthetic, even in ordinary moments. After all, we do attend to, evaluate, and take pleasure in a myriad of ordinary objects of experience: the arrangement of dishes on the dishrack, the process of folding clothes, the sound of butter spreading on toast. I side with this latter position, wherein we can say that ordinary life is aesthetic, even when it is ordinary. But if we accept this much, what stops us from saying that we respond to our own intentions and actions in ways that are characteristically aesthetic? Our experience of our own agency is a part of everyday, ordinary life. If the ordinary can be experienced aesthetically, then agency itself may have an aesthetic dimension. And if we can accept the possibility of this aesthetics of agency—and perhaps that's a big "if"—we can also accept the possibility that DBS' interventions on behavior have an aesthetic dimension as well.

My move here is in many ways an extension of Yuriko Saito's attempt—in her Everyday Aesthetics (2008)—"to expand and diversify the domain of the aesthetic" (10) to include experiences that neither form some special species of experience nor require a special mode of appreciation. She argues, further, that we respond to our everyday experiences in this way (aesthetically) so often that they “perhaps constitute the core of our aesthetic life” and we respond this way so readily that our responses are “almost automatic” (10). We may not always acknowledge the aesthetic in our everyday lives, but its features are still common.

Even further, on Saito’s view, our aesthetic responses play a key role in our practical lives. After all, a variety of our “decisions and actions are primarily, if not exclusively, guided by aesthetic considerations” (46-47), and this is evidence that our aesthetic responses “propel us toward everyday decision and actions” (11). She cites a variety of examples—“choosing what to wear and what sort of haircut to get, cleaning and ironing clothes, […] deciding whether or not to dye our hair” (46)—but most of us can imagine countless examples of how our aesthetic sensibilities guide our everyday activities. In some ways, it would be difficult to imagine what motivates most of our everyday decisions if not our aesthetic sensibilities.

Imagine deciding, for example, what to eat for dinner without thinking about what you’re “in the mood for.” The decision to eat one food over the next is very often more than just a pragmatic decision. It is, instead, guided by the person’s sensibilities, preferences, and expectations for the experience to come. It seems inevitable, on my extension of Saito’s view, that aesthetic responses to experience also have an impact on possibilities imagined, decisions made, and actions taken (i.e., on agency).

But if we accept Saito’s framework, it seems like we can also say that we can we respond to or appraise the aesthetic features of exercising our agency. Not only do we appreciate the aesthetic features of everyday activities, we also appreciate the aesthetic features of action itself—of being presented with a choice, deliberating over a set of options, making a decision, acting on it, and living with the consequences. Think of, for example, what it’s like to develop romantic feelings for a close friend. For many of us, these feelings strike us as a challenge, and so we feel confronted with a choice: conceal our feelings (and keep the relationship unchanged) or confess (and change that relationship—perhaps irrevocably). Surely there are parts of the experience that seem to have very little to do with agency—that is, they seem to be out of our control prima facie. To be infatuated, after all, is to undergo strong physiological responses when we're around the person we've fallen for: a racing heart, sweating palms, fluttering feelings in the stomach. Responses like these are often a component of rich emotional episodes: we feel anxious when we're around our would-be lovers, lonesome when they're away, and (perhaps) jealous when they're with someone else.

Even still, these physiological and affective features of experience can color, aide, or confound our attempts to exercise agency. So many of us who have fallen for someone can say what it's like to keep our feelings secret, to be uncomfortable around the person we feel so strongly for, to fantasize about what could (or what couldn't possibly) be, to dread the threat of rejection, to finally buckle under the pressure of our feelings, to confess those feelings awkwardly (or, perhaps, suavely), to live with the consequence (be it a rejection or a happily-ever-after or the vast range in between). These are the kinds of aesthetic experiences that inspire music and poetry, and they have their root in our first-person experiences of making choices, acting on them, and living with consequences. Insofar as we appreciate these features of our experiences as they happen, the aesthetic is, to extend Saito, often a crucial catalyst and guide for action. I don't want to simply belabor the point that our decisions (in matters of romance or otherwise) are guided by aesthetic sensibilities and preferences, nor do I mean to imply that we treat agency as some state of affairs we observe from the outside.[5] Instead, I am suggesting that these aesthetic features are an important part of agency itself.

After all, it is possible for these features to color the entire experience of acting. Take, for example, Helmut Dubiel's now ubiquitous account of having, and concealing, Parkinson's Disease. After learning that he had PD, he tried to conceal it, but his symptoms grew so severe that he could no longer keep them hidden. His attempts to conceal the symptoms of his PD from others were part of a larger pattern of keeping private matters private. "I wasted a large portion of my time and vital energy shielding my secrets from discovery," he explains, "this activity was so central to me that it was not limited to individual aspects of keeping the respective secrets but consolidated into a *habitus* where I would act cautiously, like an agent behind enemy lines" (110, Dubiel's emphasis). These aren't just practical realities for Dubiel—he doesn't simply keep secrets for personal benefit. His affective states are bound up in his habit of keeping secrets. Dubiel's experience is saturated with the aesthetic insofar as it *feels a certain way* to be "an agent behind enemy lines," and those feelings mattered enough for him to reflect on the experience and describe what it was like to go through it.

To this point, Thi Nguyen (2019) argues not only that the lived-experience of agency is an avenue for rich aesthetic experiences, he ventures that agency can become an artistic medium unto itself through gameplay. Nguyen, building on Bernard Suits' account of gameplay, notes that many games require that we take on *temporary ends*. When we play hide-and-seek, for example, we volunteer to hide ourselves or to find our hidden friends. These are ends that we

---

[5] Many aestheticians take the view that aesthetic appreciation relies on "psychical distance" between the spectator and the spectacle. We cannot, on this view, appreciate a performance of Romeo and Juliet if we believe that Romeo and Juliet's respective actors really are at risk of being killed. Conversely, we cannot appreciate the performance if we have no interest at all in the dealings between the Capulets and Montagues. See Bullough, 1912 for one of the first formulations of this view.

only take on for the sake of playing the game, and not for the sake of accomplishing some practical end in an efficient way. To have these experiences, however, we must take on the game's temporary ends genuinely. We must immerse ourselves in the game enough that "winning … temporarily [occupies] our mind as if it were a final end. [We] must submerge [ourselves] in an alternate agency" (12). Just like we take on temporary ends, we also inhabit a *temporary agency* to achieve them.

But why do we play games at all if their ends are unnecessary and their agencies are temporary? On Nguyen's view, "we can take on these ends temporarily for the sake of the experience of pursuing them"—the thrill of almost being found, the satisfaction of finding someone well-hidden. To play a game is to undergo an aesthetic experience. "These are the experiences of acting, deciding, solving, and doing," Nguyen insists, "not only of appreciating the movement […] but of originating it" (7).

When we play games, we don't just enjoy the experience from an outside perspective (as a spectator would), we enjoy the lived-experience of participating in the game. Nguyen argues that this "is a paradigmatic form of the aesthetic experience of one's own activity. Once we've seen it in games, we can find aesthetic experiences with this character in all sorts of practical tasks" (8). After all, if we take on temporary ends and agencies for the sake of aesthetic experiences, it doesn't seem to be at all a stretch to say that we can find similar experiences within our more long-standing, overarching ends and agencies.

Nguyen makes his account of games compatible with Dewey and Saito in one stroke: "Dewey suggests that artworld practices are […] crystallizations of experiential qualities of everyday experiences. […] If you believe anything of this sort, then you might think that games are the artistic crystallizations of everyday practical reasoning and practical action—of the experience of choosing, deciding, and strategizing." Nguyen, however, inherits some of Dewey's hesitance to embrace the aesthetic within everyday life. We play games, he insists, because they are artfully sculpted scenarios through which we have novel, rich, memorable aesthetic experiences. In ordinary life, on the other hand, many of these experiences are uncommon and hard-won—most are not as rich and not as memorable.

Crucially, however, if we take on temporary agencies within games for the sake of the experiences we have as a result, it also makes sense that we reconfigure our agencies for the sake of our everyday experiences. Outside of a game, however, the stakes are higher: we cannot just choose to play another game, take on another agency entirely. Living a life worth living often means reconfiguring our environments, sensibilities, abilities, desires, goals, plans, intentions, and decisions so that our experiences are rich, rewarding, meaningful.

Let's re-trace our steps. Dewey gives us an account of how agential features comprise "an experience"—but his view seems to ignore a variety of ordinary, *incomplete* experiences. Saito's view makes room for the appreciation of the aesthetic in everyday life, but she focuses on the objects of experience rather than the aesthetic dimensions of agency itself. Nguyen's view, finally, locates the aesthetic within the experience of acting itself, and games demonstrate how agency itself is aesthetic. In the previous section, we traced a thread in the debate over the use of authenticity-language as a way of describing experiences of using DBS. We found that the attempt to be authentic is a kind of production, a creative act that participates in the aesthetic. As such, I suggested that we could turn to aesthetic frameworks to make sense of how DBS users use authenticity-language. Between the views we've considered in this section, we have enough tools to make sense of the lived experience of DBS. Recall, how Parens notes that restorative and augmentative technologies "change what it feels like to be in the world" (33). On my view, a crucial feature of what it feels like to be in the world is the *aesthetic dimension* of being an agent—of which feelings of alienation and authenticity, isolation and closeness, ambiguity and certainty, strife and ease are an important part.

DBS has the profound potential to change how it feels to express agency insofar as it removes or introduces experiences that impinge on, extend, or modulate agency. When therapy is successful, DBS treats symptoms of conditions that either constrain or color the user's ongoing, ordinary, lived-experiences of agency. In these moments, DBS reconfigures agency in ways that produce experiences—of being more authentic, of being free to act, of having new desires, of self-(re)discovery—the user can appreciate as they unfold. When it is unsuccessful, DBS fails to change the features of the user's experience of agency they wanted (sometimes desperately) to change. And, finally, when DBS introduces side-effects—in the form of behavioral, physiological, or affective changes—their agencies are reconfigured in ways that produce puzzling experiences. Elevated libido, manic episodes, and feelings of ambiguity are difficult because of the aesthetic experience that arises from each respective configuration of the experience of agency.

When users use authenticity-language to describe these experiences of agency, it's possible that the user has a metaphysics of self in mind that they tacitly accept, and I accept that it can be fruitful to try to make sense of that metaphysics. An overlooked but salient way to understand device user's authenticity-language is to see it as an expression of their struggle to shape and articulate their selves in ways that fit with their aesthetic self-making project (that is, with their aesthetic practice of agency). Seeing it this way not only allows us to better understand how DBS modulates feelings of authenticity, but it also allows us to anticipate how DBS will impact users' (aesthetic) experiences of agency more broadly: insofar as it shapes how users are able to express agency, DBS shapes how it feels to be an agent.

5. A Demarcation Problem

One might object, however, that my view runs into the following demarcation problem. It is one thing to say that some of our actions are worth appreciating or evaluating—when the ordinary becomes extraordinary, or when an artist crystallizes experiences into artworks—but it is another to say that agency itself is *saturated* with the aesthetic. Surely, the objection might go, not all of our actions are saturated this way. Some actions are just actions, and so they have nothing at all to do with the aesthetic. Alternatively, there are times when we don't pay attention to the aesthetic, there are times when a person just wants to complete a task.

Perhaps, then, my view either (in the best case) provides no way of demarcating the aesthetic from the non-aesthetic or (in the worst case) views all exercises of agency through aesthetic lenses. In the former case, my use of the term "aesthetic" is ineffective; in the latter, the term "aesthetic" becomes vacuous. I believe my view embraces the latter implication, but that doesn't render the view vacuous. To the contrary, we can think of this implication as a strength of the view. After all, to say that all exercises of agency have an "aesthetic dimension" is just to say that agency is bound by the aesthetic.

Suppose that I posed the following (mostly) analogous view instead: "Our experiences have a temporal dimension, and so we ought to make sense of the temporal dimension of being an agent. Once we do, we will be better positioned to understand how DBS effects agency through the lifetime of the treatment." If someone objects that my use of the word "temporal" is vacuous here—because that means all agency is temporal—I would be comfortable with this observation and wonder why it amounts to an objection to the view. After all, agency is temporal insofar as there is a crucial difference between deliberating *before* we act and deliberating *after* we act, and this is the case even if we do not constantly attend to the difference. (But, then again, don't we always at least act with the difference well within reach?) This is not a perfect analogy, of course. There is no binding law of physics that imbues human agency with aesthetic features the way agency is constrained by spacetime. But if we are even able to call one moment ordinary and the other extraordinary, then we can also say that each moment is comprehensible in terms of the quality of the experience (even if that quality is dull, or insignificant).

Further, as Sheri Irvin remarks on the possibility that single components of experience can also be the focus of aesthetic evaluation: "Once we see the particular phenomena as embedded within a more complex structure of experience, however, the reluctance to apply aesthetic concepts should diminish. Though one might resist speaking of a color or simple auditory tone, perceived on its own, as aesthetic, it is clear that such simple elements contribute to the aesthetic character of the more complex structure of which they are a part" (Irvin 2008, 30). We might think of exercises of agency similarly: even when we can make a decision unimpeded by the aesthetics of the experience, even if that experience is unimportant on its own, each individual experience is embedded in a larger structure of experience. That is, all of these experiences are still products of the complex structures that have an aesthetic dimension and give rise to our ongoing feelings of agency. We can (and often do) think of each individual experience as boring, unimportant, and forgettable when they aren't fun, important, and memorable. But even the seemingly unimportant experiences influence how we appraise how it feels to act more generally—when we are able to appreciate how those experiences fit together.

6. Recommendations

With last section's demarcation problem considered, let's turn to some of the possible implications of applying the above frameworks to neurotechnology design. Recall that Nguyen's view of games and gameplay not only give us a means of understanding the significance of the lived-experience of playing games, it also gives us a means of understanding how powerful game designers are. If playing a game means taking on a temporary agency, designing a game means structuring the environment, rules, and goals players accept as part of the game. The game designer, as such, creates a world, identity, and way-of-life for the player to adopt temporarily. But if we accept my extension of Nguyen's view, device-makers have a similar power. They have the power to shape user experiences of agency in ways that are empowering or discouraging, fulfilling or alienating.

Many features of how DBS systems work can impact user experiences: the way the signal generator is shaped along with where it is meant to be implanted, the (in)ability to recharge the system, the kinds of adjustments that users can make to the system, the system through which users are able to control them, the algorithms that control how fast (or slow) adjustments are applied. Design decisions about these systems may (or likely will) change how users experience and exercise agency. If, for example, people are using DBS to treat drug addiction, having any control of the system could feel like a temptation to circumvent the treatment. Again, the stakes may be higher for DBS users than game-players: the DBS user doesn't just take on a temporary agency when they use DBS, they take on therapy to reconfigure their everyday practical agency, in part so that their experiences of it can improve. These design decisions aren't salient just because of the practical outcomes—where the user is either able or not able to use the device.

Device-maker's decisions are salient also because of how they structure the user's experiences of agency broadly.

In many ways, then, my view might require that designers go beyond a User Centered Design (UCD) model—where products adhere to standards of effectiveness, efficiency, and satisfaction (Kübler et al. 2014; Sullivan et al. 2017)—to consider the role DBS plays in reconfiguring the user's agency more broadly and how they experience those reconfigurations. Understanding the pathway from design-of to experiences-of to testimony-about how DBS changes experiences of agency can help device-makers better anticipate how their devices will impact users downstream. Further, understanding how disability structures experiences of agency—beyond the mere preconception that the goal of neurotechnology is to "fix" disability—should help device-designers imagine where DBS fits within users' lives.

As such, neuroethical research can elucidate these features of using DBS (and neurotechnology more broadly) in two ways. First, recall how many of the interviews with DBS users we've surveyed here were attempts to ascertain the nature of "psychosocial maladaptations" to DBS—and many were attempts to track the nature of changes to users' sense of identity. In future interviews with DBS users, neuroethics researchers could pose questions prompting DBS users to elaborate on the aesthetic features of using DBS directly—that is, if using DBS ever changes how they act in ways that are satisfying or dissatisfying, extraordinary or mundane, puzzling or illuminating, motivating or demotivating, and so on.

Second, neuroethical research can try to gauge how often device designers and engineers attend to the aesthetic features of end-user's experiences. While there exist several studies and investigations of researcher perspectives on the goals of neural technology research (Sullivan et al. 2017; Kübler et al. 2014), few of these studies elucidate the role of the aesthetic in the design of these technologies. Future studies could ascertain how often device-makers ask users about the aesthetics of their experiences using devices, or how the aesthetics of users' experiences figure into device designs.

7. Conclusion

We have seen that Deep–Brain Stimulators have the potential to impact their users' in ways that are helpful or harmful. We've also seen that some commenters worry that DBS will—in the worst cases—make users feel inauthentic to themselves, confused about the source of their actions, or otherwise less able to exercise agency. I suggested that we must attend to the mechanics of users' experiences if we want to understand how these experiences come about, how they impact users, and how we might support users in the ways they need. I proposed that an aesthetic framework would give us a finer-grained understanding of these mechanics and help us understand DBS users' reports of their experiences. DBS (and, perhaps, neurotechnology more broadly) use, on my view, has an impact on the aesthetic features of being an agent. That is, DBS' treatment outcomes, side-effects, and mode of interaction can (and in some cases, like with psychiatric DBS, is meant to) change the way users act and, crucially, *what it's like to act*. That is to say, using DBS can be liberating, frustrating, confusing, illuminating, or—yes—alienating. We saw that where DBS users talk about their experiences in terms of felt-authenticity, we can take them at their word and dig even deeper. But also, knowing that felt-authenticity and alienation are just a pair of the many ways that users experience their agency, we can also be receptive to the many possible experiences users may (or will) have while using DBS. Finally, we saw that understanding DBS design in terms of agency and its aesthetics gives us a potentially powerful tool for making sense of the power device engineers and clinicians have over users' experiences.

References

Agid, Y, M Schüpbach, M Gargiulo, and L Mallet. 2006. "Neurosurgery in Parkinson's disease: the doctor is happy, the patient less so?" *Journal of Neural Transmission* 70: 409–14.

Baylis, F. 2012. "The self in situ: A relational account of personal identity." In. Vancouver: UBC Press.

Baylis, Françoise. 2011. ""I Am Who I Am": On the Perceived Threats to Personal Identity from Deep Brain Stimulation." *Neuroethics* 6 (3): 513–26.

Brown, Timothy, Patrick Moore, Jeffrey Herron, Margaret Thompson, Tamara Bonaci, Howard Chizeck, and Sara Goering. 2016. "Personal responsibility in the age of user-controlled neuroprosthetics." *2016 IEEE International Symposium on Ethics in Engineering, Science and Technology*, 1–12..

Brown, Timothy, Margaret C Thompson, Jeffrey Herron, Andrew Ko, Howard Chizeck, and Sara Goering. 2016. "Controlling our brains a case study on the implications of brain-computer interface-triggered deep brain stimulation for essential tremor." *Brain-Computer Interfaces*, September, 1–6.

Cersosimo, Maria G, Gabriela B Raina, Eduardo E Benarroch, Fabián Piedimonte, Gabriela González Alemán, and Federico E Micheli. 2009. "Micro lesion effect of the globus pallidus internus and outcome with deep brain stimulation in patients with Parkinson disease and dystonia." *Movement Disorders* 24 (10): 1488–93.

Dewey, John. 2005. *Art as experience*. Perigee.

Dos Santos, Joao Flores Alves, Sophie Tezenas du Montcel, Marcella Gargiulo, Cecile Behar, Sébastien Montel, Thierry Hergueta, Soledad Navarro, et al. 2017. "Tackling psychosocial maladjustment in Parkinson's disease patients following subthalamic deep-brain stimulation: A randomised clinical trial." *PLoS ONE* 12 (4): e0174512.

Forsey, J. 2013. "Appraising the OrdinaryTension in Everyday Aesthetics." In *Proceedings of the European Society for Aesthetics*.

Gallagher, Shaun. 2018. "Deep Brain Stimulation, Self and Relational Autonomy." *Neuroethics*, March, 1–13.

Gilbert, F. 2012. "The burden of normality: from 'chronically ill' to 'symptom free'. New ethical challenges for deep brain stimulation postoperative treatment." *Journal of Medical Ethics* 38 (7): 408–12.

Giordano, James. 2015. "Conditions for Consent to the Use of Neurotechnology: A Preparatory Neuroethical Approach to Risk Assessment and Reduction." *AJOB Neuroscience* 6 (4): 12–14.

Glannon, W. 2014. "Neuromodulation, agency and autonomy." *Brain Topography*.

———. 2009. "Stimulating brains, altering minds." *Journal of Medical Ethics* 35 (5): 289–92.

Haan, Sanneke de, Erik Rietveld, Martin Stokhof, and Damiaan Denys. 2015. "Effects of Deep Brain Stimulation on the Lived Experience of Obsessive-Compulsive Disorder Patients: In-Depth Interviews with 18 Patients." *PLoS ONE*.

Haselager, Pim. 2012. "Did I Do That? BrainComputer Interfacing and the Sense of Agency." *Minds and Machines* 23 (3): 405–18.

Irvin, Sherri. 2008. "Scratching an Itch." *The Journal of Aesthetics and Art Criticism*.

Klaming, Laura, and Pim Haselager. 2013. "Did My Brain Implant Make Me Do It? Questions Raised by DBS Regarding Psychological Continuity, Responsibility for Action and Mental Competence." *Neuroethics* 6 (3): 527–39.

Klein, Eran, Sara Goering, Josh Gagne, Conor V Shea, Rachel Franklin, Samuel Zorowitz, Darin D Dougherty, and Alik S Widge. 2016. "Brain-computer interface-based control of closed-loop brain stimulation: attitudes and ethical considerations." *Brain-Computer Interfaces* 3 (3): 140–48.

Klein, Eran, and Jeffrey Ojemann. 2016. "Informed consent in implantable BCI research: identification of research risks and recommendations for development of best practices." *Journal of Neural Engineering* 13 (4): 043001.

Kraemer, Felicitas. 2011. "Me, Myself and My Brain Implant: Deep Brain Stimulation Raises Questions of Personal Authenticity and Alienation." *Neuroethics* 6 (3): 483–97.

———. 2013. "Authenticity or autonomy? When deep brain stimulation causes a dilemma." *Journal of Medical Ethics* 39 (12): 757–60.

Kübler, Andrea, Elisa M Holz, Angela Riccio, Claudia Zickler, Tobias Kaufmann, Sonja C Kleih, Pit Staiger-Sälzer, Lorenzo Desideri, Evert-Jan Hoogerwerf, and Donatella Mattia. 2014. "The User-Centered Design as Novel Perspective for Evaluating the Usability of BCI-Controlled Applications." *PLoS ONE* 9 (12): e112392.

Leddy, Thomas. 2010. "The Nature of Everyday Aesthetics." In *The Aesthetics of Everyday Life*, edited by Andrew Light and Jonathan M Smith, 3–22.

———. 2012. *The Extraordinary in the Ordinary*. The Aesthetics of Everyday Life. Broadview Press..

Leentjens, A F G, V Visser-Vandewalle, Temel Y, and Verhey F R J. 2004. "Manipuleerbare wilsbekwaamheid: een ethisch probleem bij elektrostimulatie van de nucleus subthalamicus voor ernstige ziekte van parkinson." *Nederlands Tijdschrift Voor Geneeskunde* 148 (28): 1394–8.

Lipsman, Nir, and Walter Glannon. 2013. "Brain, Mind and Machine: What Are the Implications of Deep Brain Stimulation for Perceptions of Personal Identity, Agency and Free Will?" *Bioethics* 27 (9): 465–70.

Mackenzie, Catriona, and Mary Walker. 2015. "Neurotechnologies, Personal Identity, and the Ethics of Authenticity." In *Handbook of Neuroethics*, 373–92. Dordrecht: Springer Netherlands.

Mecacci, Giulio, and W F G Pim Haselager. 2014. "Stimulating the Self: The Influence of Conceptual Frameworks on Reactions to Deep Brain Stimulation." *AJOB Neuroscience* 5 (4): 30–39.

Nguyen, C Thi. 2019. "Games and the art of agency." *Philosophical Review* 128 (4): 423–62.

Rossi, P Justin, Michael Okun, and James Giordano. 2014. "Translational Imperatives in Deep Brain Stimulation Research: Addressing Neuroethical Issues of Consequences and Continuity of Clinical Care." *AJOB Neuroscience* 5 (1): 46–48.

Saito, Yuriko. 2008. *Everyday Aesthetics*. Oxford University Press.

Schechtman, Marya. 2010. "Philosophical reflections on narrative and deep brain stimulation." *The Journal of Clinical Ethics* 21 (2): 133–39.

Schüpbach, M, M Gargiulo, M L Welter, and L Mallet. 2006. "Neurosurgery in Parkinson disease A distressed mind in a repaired body?" *Neurology* 66: 1811–6.

Sharp, Daniel, and David Wasserman. 2016. "Deep Brain Stimulation, Historicism, and Moral Responsibility." *Neuroethics* 9 (2): 173–85.

Smeding, H M M, A E Goudriaan, E M J Foncke, P R Schuurman, J D Speelman, and B Schmand. 2007. "Pathological gambling after bilateral subthalamic nucleus stimulation in Parkinson disease." *Journal of Neurology, Neurosurgery & Psychiatry* 78 (5): 517–19.

Sparshott, Francis Edward. 1982. *The Theory of the Arts*. Princeton University Press.

Sullivan, Laura Specker, Eran Klein, Timothy Brown, Matthew Sample, Michelle Pham, Paul Tubig, Raney Folland, Anjali Truitt, and Sara Goering. 2017. "Keeping Disability in Mind: A Case Study in Implantable BrainComputer Interface Research." *Science and Engineering Ethics* 398 (5): 297.

Taylor, Charles. 1992. *The Ethics of Authenticity*. Harvard University Press.

Witt, Karsten. 2017. "Identity change and informed consent." *Journal of Medical Ethics* 43 (6): 384–90.

Closing Remarks

In each of the above papers, I've tried to make sense of users' experiences of using DBS. In *Building Intricate Partnerships*, I suggested that DBS users are able to form intricate partnerships with their implants in order to express agency in the ways they want to. These partnerships, I argued, are made easier if the user has "agentic skills": introspective vigilance about how DBS influences them, the ability to improvise when DBS doesn't perform as they want it to, and the ability to make use of other resources if necessary. In *Deep–Brain Stimulation and the Aesthetics of Agency*, I proposed that we use aesthetic frameworks to make sense of how DBS impacts users' creative expressions and experiences of agency. We found that device designers have a great deal of power over the aesthetic features of users' day-to-day experiences: their design decisions can shape users' experiences of agency. Finally, in *Jumping Through Loops*, I examined the components of CL-DBS—biomarker recording, classification, and stimulation—in order to evaluate if CL-DBS systems will be trustworthy. We also found that each component of the loop could weave broader systems of injustice into the inner workings of their implanted devices.

Together, these papers have implications for how users, neuroethics researchers, neural engineers, and device manufacturers approach DBS. First, the conceptual tools I've provided serve as a robust framework for attending to, collecting, and addressing users' testimonies about (CL-)DBS. For example, when researchers have conversations with DBS users, they ought to ascertain what kinds of relationships they've formed with their implants and how those relationships shape experiences of agency (for better and/or worse). Future, more empirical, work will uncover more modes of interaction between DBS and users as well as more descriptions of how DBS impacts the aesthetic dimensions of agency. Second, each of the papers above describe how the design of (CL-)DBS can have a profound, morally salient impact on user experiences. By working backward from users' testimonies about their experiences with (CL-)DBS, device designers can figure out what about their designs works, what doesn't, and—potentially—how to change their devices to address user concerns. Knowing what about DBS makes users feel alienated, for example, is as important as knowing *that* they feel alienated. As neurostimulation devices play an increasingly central role in the treatment of both motor and psychiatric disorders—and as these devices become more complicated and robust—it is incumbent on us that we attend to the details of users' experiences.