

Optogenetics, Pluralism, and Progress

Jacqueline Anne Sullivan*†

Optogenetic techniques are described as “revolutionary” for the unprecedented causal control they allow neuroscientists to exert over neural activity in awake-behaving animals. In this article, I demonstrate by means of a case study that optogenetic techniques will only illuminate causal links between the brain and behavior to the extent that their error characteristics are known and, further, that determining these error characteristics requires (1) comparison of optogenetic techniques with techniques having well-known error characteristics (methodological pluralism) and (2) consideration of the broader neural and behavioral context in which the targets of optogenetic interventions are situated (perspectival pluralism).

1. Introduction. A primary aim of neuroscience is to understand how neural activity gives rise to behavior. Neuroscientists have long supposed that by interfering with the activity of brain circuits, synapses, cells, or molecules and tracking the effects of those interventions on behavior, they can establish which structures are causally implicated in which behaviors. The recent “optogenetic revolution” in neuroscience has served to increase confidence in the merits of this approach. Optogenetic techniques, developed by Karl Deisseroth and colleagues (Boyden et al. 2005), allow neuroscientists to use genetically encoded proteins to make selected populations of neurons sensitive to light. Optical stimulation is then used to noninvasively turn these neurons on and off with millisecond precision in awake-behaving animals and investigate the impact on behavior. Hailed as the “method of the year 2010” in the journal *Nature Methods* (2011), optogenetic techniques were rapidly implemented in research laboratories around the world (see, e.g., Bickle 2016) and

*To contact the author, please write to: 3148 Stevenson Hall, 1151 Richmond St., Western University London, ON N6A 5B8, Canada.

†The author would like to thank Wendy Parker and an anonymous referee for very helpful comments on an earlier draft of this article. Thanks also to Marshall Abrams and members of the audience at the PSA 2016 symposium on Optogenetics and Control over Neural Systems for very helpful discussion.

Philosophy of Science, 85 (December 2018) pp. 1090–1101. 0031-8248/2018/8505-0028\$10.00
Copyright 2018 by the Philosophy of Science Association. All rights reserved.

immediately began to generate exciting results, most notably that they could be used to induce false memories in rats (see, e.g., Ramirez, Tonegawa, and Liu 2014; see also Robins 2016). In fact, insofar as the techniques have set a new standard for “making causal links between elements of neural circuits and behavior” (Adamantidis et al. 2015, 1207; see also Häusser 2014), their incorporation into research studies in combination with other techniques has become a standard for publication in systems and behavioral neuroscience.

In this article, I use a case study to demonstrate that optogenetic techniques will only illuminate causal links between the brain and behavior to the extent that their error characteristics are known, and developing this knowledge requires methodological and perspectival pluralism. Optogenetic techniques have been described as “revolutionary” for the unprecedented causal control they allow neuroscientists to exert over neural activity in awake-behaving animals. Yet, the techniques have been introduced into domains of neuroscience where it is common for investigators to intervene at a single locus in the brain and track a select set of changes in behavior without adequate consideration of neural events downstream of the manipulation or a detailed analysis of behavior (see, e.g., Sullivan 2010; Krakauer et al. 2017). It is argued by appeal to the case study that this “reductionistic perspective” disables scientists from appreciating the errors to which novel tools like optogenetics may be subject when they are used to test causal hypotheses linking the brain to behavior. Using the case study, I show that identifying such errors requires (1) a comparison of optogenetic techniques with techniques having well-known error characteristics (“methodological pluralism”) and (2) a consideration of the broader neural and behavioral context in which the locus of an optogenetic intervention is situated (“perspectival pluralism”; e.g., Giere 2006; Sullivan 2014, 2016).

I begin, in section 2, with a preliminary appraisal of optogenetic techniques, comparing them to more traditional intervention techniques used in systems and behavioral neuroscience. In section 3, I describe a research study that challenges the idea that optogenetics alone is poised to shed light on how neural circuits give rise to behavior. Appealing to a set of conceptual tools from the philosophical literature on scientific experimentation and pluralism in section 4, I evaluate the research study and derive some general lessons about the necessity of methodological and perspectival pluralism for progress in understanding causally complex systems in neuroscience.

2. What Is Revolutionary about Optogenetics? A primary aim of neuroscience is to determine how the brain gives rise to or *causes* behavior. One standard investigative approach is to train nonhuman animals (e.g., rodents) on cognitive tasks in experimental learning paradigms as a means to produce detectable changes in behavior and to use either transient or permanent intervention techniques to alter neural activity before, during, or after training

to assess the impact on behavior. Woodward's (2003) manipulability-based account of causation nicely captures the interventionist approaches operative in such experiments (see also Craver 2007). An investigator ideally aims to intervene in a single variable and establish that the variable is necessary for the behavior under normal conditions. If she observes the relevant behavioral differences between an experimental and a control group, she infers that the manipulated variable is "causally responsible for" or "implicated in the production of" the behavioral changes of interest.

In order to convey why optogenetics is a "revolutionary" intervention technique, it is relevant to briefly describe the traditional intervention techniques that it has to a significant extent come to replace. Lesioning, for example, was at one time considered to be the best method for intervening in brain activity. Put simply, lesioning involves the selected removal of brain tissue in an anesthetized animal before or after training in an experimental paradigm. One disadvantage of this method is that it is temporally crude; animals must go through a recovery period after surgery, which may allow sufficient time for so-called plastic changes in neural networks, synapses, cells, and molecules to take place. Such changes, insofar as they are difficult to detect, serve as confounding variables in the attempt to establish causal links between the brain and behavior. Given the possibility of such changes, it is widely accepted that an animal with a brain lesion should not be regarded simply as a normal animal minus the brain area that has been removed, as another area may subserve recovery of that function (e.g., von Melcher, Pallas, and Sur 2000). Lesions are also spatially crude; it is difficult to ensure that the intended brain area has been removed completely or to tell whether the surrounding tissue also has been damaged. These features can only be verified via postmortem analysis (see Bechtel and Stufflebeam [2001] for further discussion).

Investigators can also knock out genes in laboratory mice by replacing or disrupting an existing gene, thereby inactivating it. The aim is to determine whether the gene plays a causal role in the production of the behavior of interest. One advantage of this technique over lesioning is that it is spatially precise; only the relevant genes are altered, and healthy knockout mice appear to develop "normally." As with lesions, however, one cannot treat knockout mice as normal mice simply lacking a gene. The absence of the gene across the animal's developmental trajectory may affect its overall phenotype, including its morphological, biochemical, physiological, and behavioral properties. Some of these changes may be verified postmortem, but others may go undetected. Gene knockouts thus do not provide clean interventions, which is why causal claims about the roles of genes in behavior are often generically phrased (e.g., "plays some role" in the behavior of interest).

Drugs can also be used to alter the activity of neurons by modifying the activity of transmembrane receptors and ion channels in the whole brain or selected brain areas. Drug infusion via a cannula, a thin tube directly inserted

into the brain, compared to intraperitoneal and intravenous drug injections, affords an investigator the greatest spatial and temporal control of the available options for drug delivery. Drugs can be infused at a precise location at the exact moment during a cognitive task when the investigator wishes to alter activity in that area. Despite these advantages, this method is still spatially and temporally crude. If too little drug is used, it may not diffuse sufficiently so as to facilitate or block neuronal activity in the desired location; if too much is used, it may diffuse beyond the intended area and alter neuronal activity outside of the desired location. While getting the drug concentration and amount of drug just right is critical, determining the right concentration and rate of delivery usually involves trial and error. Additionally, an investigator can only verify during postmortem examination whether the cannula was in the right location during the experiment.

Optogenetics is the newest addition to the stock of intervention strategies available to neuroscientists. First developed and described by Karl Deisseroth and colleagues (Boyden et al. 2005), optogenetic techniques involve two steps. First, an animal that expresses opsin DNA in a selected brain area must be created, either by inserting an engineered opsin DNA-virus vector directly into neurons in that brain area or genetically engineering an animal that expresses opsin DNA in that region. This measure programs these neurons to synthesize light sensitive opsin proteins that embed in the membranes as ion channels. An investigator can then take these animals, run them through experimental paradigms and deliver optical stimulation via a chronic implant directly to these neurons at the desired moment. The channels open in response to this stimulation, resulting in immediate changes in ionic concentrations that either activate (depolarize) or inactivate (hyperpolarize) these neurons.

Optogenetics has been described as a revolutionary intervention technique because it has none of the limitations of the traditional techniques described above and several major advantages (see, e.g., Häusser 2014; Bickle 2016). In contrast to lesions and drug infusions, the technique is relatively noninvasive, only requiring a cannula insertion to allow for delivery of the virus vector or optical stimulation. Moreover, the effects are reversible—neuronal activity in a brain area is understood to return to normal after optical stimulation is shut off. It is also temporally and spatially precise; an investigator can turn a restricted population of neurons on or off with millisecond precision while an animal engages in a cognitive task. So, not only does the intervention purportedly directly hit its intended target, if there are observable changes in the behaviors of interest they can be immediately detected. For all of these reasons optogenetics has been touted as “bring[ing] neuroscience closer to causality” (Adamantidis et al. 2015, 1207) than previous techniques. Furthermore, the techniques have effectively changed standards for publishing research in systems and behavioral neuroscience. Including at least one

optogenetic experiment in a research study aimed at establishing causal links between neural circuits and behavior has become, over the past 10 years, a desired if not required benchmark for publication (see, e.g., Bickle 2016).

No one would deny that the development of optogenetic techniques—the ability to control neurons with light—is a revolutionary and exciting achievement independent of how the techniques are ultimately used. The question of interest to me in this article, however, is whether optogenetic techniques alone can revolutionize those areas of neuroscience into which the technique has come to be widely used. Bickle has described the methodological strategy in these areas of neuroscience as that of “intervene cellularly/molecularly and track behaviorally” (2006, 425). Although Bickle made this claim over 10 years ago, well before the onset of the optogenetic revolution in neuroscience, this “reduction-in-practice” strategy, as Bickle dubbed it, remains dominant. In other words, with few exceptions, research using lesions, gene knockouts, and drug infusions historically has focused on testing causal hypotheses about the role of a single brain area, neural circuit, synapse, cell, or molecule in a single set of changes in behavior. Rarely in such research studies do scientists investigate changes upstream or downstream of the area, circuit, or synapse in which they intervene. In other words, they do not know the impact of these interventions on brain areas intermediate between those areas in which they disrupt activity and the changes in behavior they observe. This does not, however, stop them from causally linking neural activity to behavior on the basis of data produced using such interventions.

It is also uncommon for investigators in these areas of neuroscience to closely analyze or describe changes in behaviors over and above those changes in behavior they are trying to produce. For example, Krakauer and colleagues claim that while “technique-driven neuroscience” claims to be addressing broadscale questions, like “How does the brain generate behavior?” investigators in this area refuse to engage in the “careful dissection of behavior into its component parts and subroutines” (2017, 481). Sullivan (2010, 2014) voices similar concerns, pointing to the fact that investigators deploying the intervention techniques described above are often not concerned with the component cognitive processes that are involved in the production of the behavioral effects under study. She regards this as an impediment to explanatory progress because discovering how the brain gives rise to behavioral functions requires knowing how to parse and classify those behaviors.

In response to these criticisms it may be granted that single research studies require investigators to focus on testing a restricted set of causal hypotheses and that this sometimes means they must constrain their inquiry to a single brain area and a single set of behaviors. Yet, the problem is that the causal claims emanating from neuroscience following the rise of optogenetic techniques have been anything but modest. Michael Häusser (in Adamantidis et al. 2015, 1207) expresses a widespread and popular belief that optogenetic

techniques are “powerful tools for making causal links between elements of neural circuits and behavior—in that we can prove both necessity (by inactivating neuronal populations) and sufficiency (by activating the same neurons).” If anything, optogenetics has increased confidence that criticisms of reductionist methodologies are misguided.

Are the critics wrong? Neuroscientist György Buzsáki claims that optogenetics has gone through an “initial hyper-enthusiastic phase” in which investigators have been making “outrageous claims about the novel method’s power and specificity,” but this has yet to give way to a “maturational stage” in which investigators come to understand “the objective and reliable values of the method” (in Adamantidis et al. 2015, 1208). Although Buzsáki indicates that this maturational stage will require “hard work,” he says nothing about the nature of the kind of labor required. This is where I think the critics get things right—that the strategies that might propel neuroscience into and through this maturational stage require a change in methodology and perspective. As I argue in the remainder of this article, a recent research study nicely exemplifies the kind of multipronged strategy likely to be successful.

3. Optogenetics and “Off-Target Effects.” The research study that I focus on here arose out of an earlier study that it is relevant to briefly describe. In the original study, neuroscientist Bence Ölveczky and his research team at Harvard trained several groups of rodents on a motor skill learning task (Kawai et al. 2015). Before each training session they deprived rats of water and then placed them into a box containing a lever. When the rats pressed the lever twice in short succession (700 ms apart), they received a water reward. Across five to seven training sessions each lasting approximately 1 hour, the rats learned the required motor sequence.

Once Ölveczky and colleagues observed that the paradigm produced robust learning, they took one group of rats and lesioned the primary motor forelimb cortex (PMFC) before training. They took another group and lesioned the PMFC after training. They observed that rats that underwent PMFC removal before training were subsequently unable to learn the motor sequence. In contrast, rats in which the PMFC was removed after training could perform the task as well as they had before PMFC removal. Ölveczky and colleagues interpreted these data as indicating that the PMFC is causally implicated in learning the motor sequence rather than in executing it.

In an effort to determine whether these findings were robust across different techniques for intervening in PMFC activity, Ölveczky and colleagues decided to conduct a comparative study (Otchy et al. 2015). They trained three separate groups of rats in the motor sequence paradigm and conducted three different intervention experiments. In one group of rats they injected 100 nanoliters (nl) of the GABA_A-agonist muscimol via cannula directly into the PMFC opposite the dominant paw—a measure designed to temporarily activate inhib-

itory neurons in this brain area, thus shutting down excitatory activity. This group exhibited “marked drops in performance” on the task and “disrupted paw kinematics” (359), insofar as they could not make the requisite fine motor movements with their paws to execute the required motor sequence. In a second set of experiments, Ölveczky’s team lesioned the PMFC in a group of rats and allowed them to heal for at least 5 days after surgery. When these rats were returned to the task context, their performance was observed to be comparable to their prelesion performance, which was consistent with findings from the first study described above (Kawai et al. 2015). In the third set of experiments, a group of rats expressing the optogenetic activator Chrimson in PMFC were trained on the task. On the first lever press after these rats had reached asymptotic performance, Ölveczky and colleagues used a light stimulus to activate PMFC neurons in “the hemisphere contralateral to the dominant paw” (Otchy et al. 2015, 359). They observed that optogenetic stimulation correlated with an immediate decrease in task performance and a disruption in the ability to move the paw so as to elicit the required motor sequence. Undertaking this comparative study, Ölveczky and colleagues found that, whereas permanent removal of PMFC via lesioning had no impact on motor sequence execution, the two forms of transient disruption of PMFC activity used, including optogenetics, did. Data obtained using drug infusions and optogenetics suggested that the PMFC is causally involved in successful execution of the motor task, whereas data obtained using lesions suggested that it is not.

How might this discrepancy between permanent and transient interventions be interpreted? One obvious response would be to reject the data from the lesioning experiments because optogenetic techniques purportedly offer superior spatial and temporal control and the observed disruptions in behavior were robust across two different intervention techniques (i.e., drug infusion and optogenetic manipulation) rather than observed using only a single technique (i.e., lesioning). Ölveczky and colleagues did not do this, because they believed they had adequately controlled for potential confounds in the lesioning studies. Specifically, they claim, “the discrepancy could not be explained by experience-dependent relearning in lesioned animals because the skills recovered their idiosyncratic pre-lesion form without any intervening practice” (Otchy et al. 2015, 359). Another potential response to the discrepancy would be to question the generalizability of the result. In other words, just because differences between permanent and transient interventions were observed in one neural system in one animal model with respect to a single behavior does not mean that such discrepancies would be observed in other systems.

Ölveczky and colleagues, however, did not stop at a single comparative study. They sought to determine whether transient and permanent inactivations of a brain structure involved in courtship song learning in zebra finches

resulted in a similar discrepancy. In these experiments, after they allowed juvenile birds to learn to produce courtship songs by memorizing and vocally mimicking adult zebra finches, they disrupted excitatory projections to the hyperstriatum ventral pars caudalis (HVC), an essential component of the song control circuit, to determine the impact on learned vocalizations. In one set of experiments they lesioned a sensorimotor nucleus upstream of HVC, the nucleus interfascialis (Nif), and in another they injected muscimol into Nif to transiently inhibit excitatory activity.¹ They found that permanent lesions of Nif did not affect learned birdsong (the birds' songs were similar to prelesion), but transient inactivations did. These data lend support to the more general claim that transient and permanent inactivations of neural activity may yield contradictory information about the causal role those neurons normally play.

Ölveczky and colleagues' explanation for the observed differences between transient manipulations like optogenetics and permanent manipulations like lesioning is that the former have what they dub "off-target effects." Given that brain circuits are massively interconnected, taking a single brain area momentarily off-line likely perturbs the ability of other interconnected circuits to perform their normal functions. Using a single transient intervention thus makes it difficult to attribute the behavioral effects to that brain area.² As Ölveczky and colleagues put the point, "transient circuit manipulations may have their own interpretive difficulties that stem from acute effects on the independent functions of non-targeted circuits" (Otchy et al. 2015, 362).

The underlying aim of Ölveczky and colleagues' research study was to determine whether the impressive causal claims linking neural circuits to behavior that have been made on the basis of optogenetic tools are warranted. Into the vast sea of enthusiasm for these techniques, Ölveczky and colleagues inject a healthy dose of skepticism. More importantly, the study reveals something interesting about the kinds of strategies that may ultimately be successful in propelling neuroscience toward the goal of understanding how nervous systems give rise to behavior.

4. Progress and Pluralism. An experiment is ideally supposed to leave an investigator epistemically better off than she would have been had she never run that experiment. Improving one's epistemic situation by means of an ex-

1. In another study, Roberts and colleagues (2012) demonstrate that optogenetic stimulation of HVC correlates with juvenile zebra finches being unable to learn birdsong from an adult tutor.

2. Ölveczky and colleagues also aim to demonstrate by means of a neural network simulation, which is part of the same research study, that homeostatic regulation of neural activity after lesioning likely contributes to recovery of function in lesioned animals.

periment requires that the experiment be *reliable*, or capable of producing data requisite to adjudicate among competing hypotheses about a phenomenon of interest. One valuable way to understand reliability is in terms of *severity* (Popper 1962, 1979; Mayo 1991, 1996). According to Mayo, a *severe* test of a hypothesis is a test that very likely “would not yield . . . a passing result” for a hypothesis if that hypothesis were indeed false (1991, 529). As Mayo acknowledges, empirical tests are rarely if ever maximally severe. Rather, experiments may be regarded as situated on a continuum between 0 severity and maximal severity. Increasing the severity of a test requires, on her view, the use of error-correcting or error-reducing strategies.

Implementing such strategies, of course, requires that an investigator have an awareness of the kinds of errors to which a test may be subject. Yet, as Ölveczky and colleagues’ study nicely illustrates, developing an awareness of these errors within the confines of a single experiment itself may be difficult, especially when one aims to establish a causal claim about a causally complex system. In such instances, the relevant questions become (1) What is the right approach for getting a handle on the errors if there are any? and (2) How might those errors be minimized? Ölveczky and colleagues are clearly interested in both of these questions, and they use what I regard as a two-pronged pluralistic strategy in order to answer them. On the one hand, this strategy is methodologically pluralistic (see also Südhof 2015; Pratt and Prather 2016) and, on the other, theoretically or conceptually pluralistic. These features set it apart from the vast majority of research studies being undertaken in interventionist neuroscience today. I want to consider each component of this strategy in turn.

Ölveczky and colleagues could have taken the now well-trodden path in interventionist neuroscience and used optogenetic techniques to intervene in neural activity in the PMFC to determine the impact on learned motor behavior. Had they done this, they likely would have concluded on the basis of the rodent experiments that the PMFC was causally implicated in the execution of the motor task. Similarly, if they had only blocked Nif in zebra finches using a pharmacological intervention, they would likely have concluded that this brain area was causally implicated in the elicitation of learned birdsong. Using a combination of methods, old and new, to test the same causal hypothesis in both sets of experiments, they discovered discrepancies in the data that prompted them to carefully scrutinize the methods and assess their severity. Because lesioning techniques have been in use for many years and their error characteristics are more or less known, Ölveczky and colleagues used available methods of error reduction to ensure that the lesioning experiments were (as severe as possible) tests of the causal hypotheses under consideration. In the case of transient manipulations, like optogenetics and pharmacology, Ölveczky and his team used a different strategy. They humbly acknowledged that the error characteristics of these techniques are not well

known and that, given the current state of knowledge about the brain, they were uncertain what measures could be taken to reduce these errors.

In attempting to explain the discrepancy between transient and permanent intervention techniques, Ölveczky and colleagues also do something that “current studies” that “focus on only one circuit and one behavior” (Südhof 2015, 339) do not do. Specifically, they shift their perspective from those brain areas they investigated in their study to a consideration of the broader neural circuitry in which those areas are embedded. In doing so, they generate potentially testable hypotheses as to the downstream consequences of acute circuit manipulations. They also generate hypotheses as to the nature of the homeostatic mechanisms that come online when a brain area is permanently lesioned and use computational approaches to test these hypotheses. What they do, I contend, is complement their methodological pluralism with “perspectival pluralism.”³

William Wimsatt’s understanding of a perspective nicely captures what I have in mind. According to Wimsatt, “perspectives involve a set of variables that are used to characterize systems or to partition objects into parts,” and they inform “the characteristic ways in which those observers” who adopt them “interact causally with [a] system” (2007, 227; see also Giere 2010). Different scientists take different theoretical perspectives on those systems they are interested in understanding, and oftentimes when a scientist works in a single area of neuroscience in which a single perspective is dominant, she views that system from that perspective. As I explained in section 2 above, historically those areas of neuroscience that use intervention techniques to link neural circuits to behavior have been committed to a “reduced” perspective on the brain, behavior, and the relationship between the two. Although there are certainly exceptions—Ölveczky and his collaborators, for example—there is widespread resistance among scientists working in this research area and philosophers who find the “intervene cellularly-molecularly-track-behaviorally” approach unproblematic (e.g., Bickle 2006, 2016) to embrace different theoretical perspectives on the systems being investigated (see Sullivan 2010, 2014; Krakauer 2017). Ölveczky and colleagues’ study, however, nicely illustrates the pitfalls of adhering to a single perspective and why it is antithetical to progress. It is not clear how a scientist who adopts a single theoretical perspective on a complex causal system will be able to ensure that she is subjecting those hypotheses to severe tests, as there might be unknown confounding variables over which she is not imposing control. What is potentially gained by adopting multiple different perspectives on the brain and ner-

3. In fact, it may indeed be more accurate to say that they take a methodologically pluralistic approach because they are already thinking about their objects of study from a plurality of different ontological or conceptual/theoretical vantage points. All that is important for my purposes is that these two kinds of pluralism are mutually informing.

vous system is an awareness of the potential confounding variables as well as clues for how to design more reliable experiments.

Of course, Ölveczky and colleagues' approach may be criticized for failing to be thoroughly pluralistic. Krakauer and colleagues, for example, emphasize the importance of the perspectives of "organism-level thinkers who develop detailed functional analyses of behavior, its developmental trajectory and its evolutionary basis" (2017, 481) for forging links between neural activity and behavior. Sullivan (2010, 2014), in contrast, emphasizes the value of engaging in task analyses for individuating cognitive capacities and determining what areas of the brain are responsible for bringing them about. It is likely that a combination of different approaches will be required in order to make experimental tests of causal hypotheses linking the brain to behavior more severe.

5. Conclusion. Optogenetic techniques have spurred a revolution in neuroscience. Yet, the power of these techniques to establish causal links between the brain and behavior is contingent on understanding how the techniques work and their limitations. Ölveczky and colleagues' 2015 study (Otchy et al. 2015) provides good grounds for thinking that such understanding may only be achieved if "reduction-in-practice" neuroscientists cultivate what Chang aptly refers to as "humility" (2012, 255). Specifically, they need to recognize the value of methodological and perspectival pluralism for identifying the benefits and limitations of novel experimental techniques, no matter how powerful at first blush these techniques appear to be.

REFERENCES

- Adamantidis, Antoine, et al. 2015. "Optogenetics: 10 Years after ChR2 in Neurons—Views from the Community." *Nature Neuroscience* 18 (9): 1202–12.
- Bechtel, William, and Robert Stufflebeam. 2001. "Epistemic Issues in Procuring Evidence about the Brain: The Importance of Research Instruments and Techniques." In *Philosophy and the Neurosciences: A Reader*, ed. William Bechtel, Pete Mandik, Jennifer Mundale, and Robert Stufflebeam, 55–81. Oxford: Blackwell.
- Bickle, John. 2006. "Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience." *Synthese* 151:411–34.
- . 2016. "Revolutions in Neuroscience: Tool Development." *Frontiers in Systems Neuroscience* 10 (24): 1–13.
- Boyden, Edward, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. 2005. "Millisecond-Timescale, Genetically Targeted Optical Control of Neural Activity." *Nature Neuroscience* 8 (9): 1263–68.
- Chang, Hasok. 2012. *Is Water H₂O? Evidence, Realism and Pluralism*. Dordrecht: Springer.
- Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- Giere, Ronald. 2006. "Perspectival Pluralism." In *Scientific Pluralism*, ed. Stephen Kellert, Helen Longino, and C. Kenneth Waters, 26–41. Minnesota Studies in Philosophy of Science 19. Minneapolis: University of Minnesota Press.
- . 2010. *Scientific Perspectivism*. Chicago: University of Chicago Press.

- Häusser, Michael. 2014. "Optogenetics: The Age of Light." *Nature Methods* 11:1012–14.
- Kawai, Risa, Timothy Markman, Rajesh Poddar, Raymond Ko, Antoniu Fantana, Ashesh Dhawale, Adam Kampff, and Bence Ölveczky. 2015. "Motor Cortex Is Required for Learning but Not for Executing a Motor Skill." *Neuron* 86 (3): 800–812.
- Krakauer, John, Asif Ghazanfar, Alex Gomez-Marin, Malcolm MacIver, and David Poeppel. 2017. "Neuroscience Needs Behavior: Correcting a Reductionist Bias." *Neuron* 93 (3): 480–90.
- Mayo, Deborah. 1991. "Novel Evidence and Severe Tests." *Philosophy of Science* 58 (4): 523–52.
- . 1996. *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
- Nature Methods*. 2011. "Method of the Year 2010." *Nature Methods* 8 (1): 1.
- Otchy, Timothy, Steffan Wolff, Juliana Rhee, Cengiz Pehlevan, Risa Kawai, Alexandre Kempf, Sharon Gobes, and Bence Ölveczky. 2015. "Acute Off-Target Effects of Neural Circuit Manipulations." *Nature* 528:358–63.
- Popper, Karl. 1962. *Conjectures and Refutations*. New York: Basic.
- . 1979. *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press.
- Pratt, Kara, and Jonathan Prather. 2016. "Systems Neuroscience: How the Cortex Contributes to Skilled Movements." *Current Biology* 26:R60–R82.
- Ramirez, Steve, Susumu Tonegawa, and Xu Liu. "Identification and Optogenetic Manipulation of Memory Engrams in the Hippocampus." *Frontiers in Behavioral Neuroscience* 7 (226): 1–9.
- Roberts, Todd, Sharon Gobes, Malavika Muragan, Bence Ölveczky, and Richard Mooney. 2012. "Motor Circuits Are Required to Encode a Sensory Model for Imitative Learning." *Nature Neuroscience* 15 (10): 1454–59.
- Robins, Sarah. 2016. "Optogenetics and the Mechanism of False Memory." *Synthese* 193:1561–83.
- Südhof, Thomas. 2015. "Reproducibility: Experimental Mismatch in Neural Circuits." *Nature* 528 (7582): 338–39.
- Sullivan, Jacqueline. 2010. "Reconsidering Spatial Memory and the Morris Water Maze." *Synthese* 177:261–83.
- . 2014. "Is the Next Frontier in Neuroscience a Decade of the Mind?" In *Brain Theory: Essays in Critical Neurophilosophy*, ed. Charles Wolfe, 45–67. New York: Palgrave Macmillan.
- . 2016. "Construct Stabilization and the Unity of the Mind-Brain Sciences." *Philosophy of Science* 83 (5): 662–73.
- Von Melcher, Laurie, Sarah Pallas, and Mriganka Sur. 2000. "Visual Behavior Mediated by Retinal Projections Directed to Auditory Pathway." *Nature* 404:871–76.
- Wimsatt, William. 2007. *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.