

Foreword

Patricia S. Churchland

A celebration in honor of Owen Flanagan and his storied career in philosophy stirs up memories of his tenacious curiosity about all things, along with his willingness to size up problems without being hemmed in by philosophical fashions.¹ It is also an occasion to take a long look at the academic discipline of philosophy over the last forty years or so and its sometimes quite strict orthodoxies. The glance backward is especially fitting when the connecting thread is Flanagan and the way he went about things or, perhaps more accurately, did *not* go about things.

Especially in the second half of the twentieth century, many philosophers tended to be deferential to the “mainstream” and thus to the luminaries who set the philosophical agenda, but not Flanagan. He could be respectful, but he did not shape his inquiries to fit trendy theories and methods. For example, when others were pushing “necessary truths” rooted in their convictions about what words *do* mean and *must* mean, Flanagan preferred to take a good look at what the data revealed about the phenomena the words described. When others were assembling natty metaphors to “explain” consciousness, Flanagan was the head-scratching skeptic asking what explanatory value the metaphors cornered and how they measured up against scientific data. When moral philosophers lined up to proclaim themselves either Kantians or Utilitarians and do battle accordingly, Flanagan wisely saw that there was a far deeper story that both factions totally missed.

Fads and fashions seem to find a welcome home in philosophy, which is why the Cambridge group, Beyond the Fringe, and their later Oxfordian counterpart, Monty Python’s Flying Circus, found such ready material. The satire staged in 1977 by John Cleese and Jonathan Miller of a common room conversation between two philosophers concerning what

philosophers do remains surprisingly current. At one point, Cleese responds to some banal point of Miller's by saying "yes, yes" and Miller sincerely queries, "Do you mean 'yes' in the affirmative sense," to which Cleese, after a thoughtful pause, says, "no."² The entire caricature is stunningly silly, and yet the echoes of actual philosophical practice are all too audible. Of course, fads and foolishness can be found in science too, but perhaps any discipline is especially vulnerable to fads when little is known concerning its chosen topics, thereby rewarding debating swagger and sometimes plain silliness. For example, this was true concerning the nature of the heart until William Harvey discovered in the early seventeenth century that it was a meat pump.

What Flanagan realized early on, and what motivated his first book, *The Science of the Mind*,³ is that if we want to understand phenomena such as consciousness, decision-making, and reasoning, we will not get very far unless we have evidence, as opposed to mere speculations, metaphors, and dogmatic edicts, about what certain words mean. Thus, he felt at home with naturalism, the very general view that understanding natural phenomena requires, among other things, paying attention to the facts regarding the phenomena. Or, to paraphrase Quine in his defense of naturalism, there is no principled division between science and philosophy.

Although Flanagan's heresy on how to make progress might have been a break with trends in philosophy circa 1980, it was not a break with the grand philosophical tradition of Aristotle, René Descartes, David Hume, Adam Smith, William James, and Hermann von Helmholtz. Rather, it was an extension of that same tradition. At this point in the twenty-first century, Flanagan's unfashionable regard for the empirical world of experiments, data, and how science makes progress has paid off—handsomely. By contrast, some of the hotly contested ground of the 1990s now looks a bit forlorn, more abandoned than productively cultivated.

As Flanagan well appreciated, neuroscience, ethology, anthropology, and experimental psychology have made discoveries that, if we permit, have an impact on traditional philosophical questions, such as the nature of decision-making, consciousness, and the origins of moral values.

As we look to the future, the nature of knowledge, learning, and memory is a broad topic where the impact of neuroscience on philosophical theories is likely to be substantial, if not outright revolutionary. Any epistemologist worth her salt will be intrigued by the discovery that nonverbal animals clearly have spatial knowledge, such as where home is and where the food

is, as well as nonspatial knowledge, such as who one's mate is. Like us, non-verbal animals learn from experience, and what is learned is put to good use in their behavior. Neuroscientists have inquired into the mechanisms that support this knowledge, and impressive data relevant to these mechanisms have been discovered. In 2014, the Nobel Prize in Physiology or Medicine went to Edvard Moser, May-Britt Moser, and John O'Keefe "for their discovery of cells that constitute a positioning system in the brain."

It turns out that the hippocampus and connected structures (e.g., medial entorhinal cortex, pre-subiculum) are essential for spatial knowledge as well as for episodic memory generally (e.g., what music I listened to at breakfast). There is a demonstrable link between repeatable patterns of activity of the neurons of the hippocampus of a rat and the spatial routes it uses to navigate the world it inhabits. In short, the hippocampus and ancillary structures construct organized, flexible, and usable representations of the rat's spatial environment. Spatial knowledge is only part of what the hippocampus constructs. Other cognitive functions, such as social knowledge of who to trust, also depend on the hippocampus.

The neural mechanisms supporting spatial cognition are complex, though major clues derive from the precise tuning properties of distinct groups of neurons. Some neurons (place cells) respond preferentially whenever the rat enters a particular place in an area such as a room. Other neurons (head direction cells) respond specifically to the direction in which the rat's head is pointing. Yet other neurons respond to a specific distance from the *borders* of the region that the rat is in. Remarkably, in the *medial entorhinal cortex*, which feeds into the hippocampus, a population of neurons collectively tiles the entire space that the rat is in. These neurons are called *grid cells*, and they allocentrically represent a spatial area (as opposed to merely a response to sensory cues in the area). When the rat explores a new room, these neurons then tile that new area. Under those conditions, the hippocampal place cells remap as well. Moreover, the upshot of this complex spatial framework is that the rat knows *where it is* in the larger environment. It knows where to go to get what it seeks. The rat need not see the food or the hiding place to make a beeline for it.⁴

Exactly how the different types of cellular signals are integrated to yield the systematic representations of space and other forms of relational representations (such as remembering the flow of events across time and the ability to draw novel inferences from memory) are not yet well understood.

The rat's spatial knowledge is more abstract than one might have guessed, as is illustrated by recording the activity of neurons. When a rat is resting or merely planning to forage, an observable pattern of activity is seen in the rat hippocampus that replicates the activity seen, for example, when it is actually running that successful route to the food cache. In short, its brain rehearses the winning route even while not running the route, and the time taken is compressed. This is probably a lot like what I do when I rehearse (i.e., imagine) a good route to a chanterelle mushroom patch in the woods. I would say, and many neuroscientists do say, that the rat knows where the food is likely to be found much as I know my route to the chanterelles. When a rat confronts a range of possible routes, its brain will play out (i.e., imagine) a series of routes before making a selection. This is probably a lot like what I do when I confront a range of routes for a portage of my canoe, some of them more hazardous than others. This rehearsal of routes involves representing the future—where I could go or will go—not just the past (what I did last time). In short, there is a lot of smart cognitive business going on here. This is not simple stimulus-response conditioning.

For philosophers to appreciate the epistemological significance of hippocampal discoveries made over the last four decades, they will find it useful to know as much as possible about the physiology of the hippocampus and its role in behaviors such as foraging and escaping predators. Along the way, they may bump into developmental psychologists, such as Jean Mandler, who are testing the hypothesis that spatial primitives (such as “close by” and “far away”) in baby humans are the scaffolding for many categories that babies come to acquire as they develop.⁵ They will perhaps see a link to the vast array of spatial metaphors common in so many languages that scaffold the way we make inferences.⁶ Undoubtedly they will be more productive if they relax the rigid dogma that “knowing how” involves a completely different system from “knowing that,” along with the favored assumption that believing something always requires having a language in which to represent what is believed.

One dividend of absorbing the hippocampal research is likely to be progress on intentionality, otherwise known as the *aboutness* of representations. Put another way, the question is how representations can represent something in the world, including something in the body and brain itself, such as metacognitive representation, where a monkey or a human knows what it does know. Although intentionality is sometimes considered an

exclusively philosophical topic, at this stage the topic appears to be sorely in need of fresh ideas. Consequently, it may be productive to reflect on the neurobiological discoveries showing that representations in neural networks are about spatial and temporal relationships in the world. The allocentric system of grid cells is undoubtedly about some small part of the world. As the mechanisms sustaining a “positioning system” come to be more fully understood, will that help explain what it is for activity in neurons to be about things in the world? A good bet says *yes*.⁷

Should the philosophical contribution consist in divining the deep conceptual truths that limit what the neuroscientists can meaningfully say or discover about the nature of representation—of knowledge and belief? This is not likely to be a rewarding strategy. Frequently in science, we have to make do with an imprecise expression to refer to a phenomenon, at least until the science develops to the point where greater precision is possible. *Gene* is a well-known case in point. To make that term more precise, scientists in the 1940s did not turn to philosophers to reveal the conceptual truths about the concept “gene.” They slowly got more data, whereupon the meaning of the word *gene* became better defined in tandem. Many expressions used in neuroscience are like *gene* circa 1940, because much about how the brain works remains poorly understood. The words *represents* and *believes* are two examples, and their meaning is apt to change profoundly over the coming decades as more is understood about how the brain learns, knows, and represents. Although immersing yourself in the neurobiological data is no guarantee of a theoretical breakthrough, it pretty obviously seems a necessary condition.

Although Flanagan was not especially focused on *represents*, he had long suspected that a lot of heavy philosophical posturing about what words must mean (supposed conceptual truths) did not accurately reflect what was going on in the actual world of language, science, and practical reasoning. All too often, what are shopped as conceptual truths anchored by word meanings are actually empirical speculations floating free of an empirical base.

As I look on the Flanagan shelf in my library, I tend to think my favorite Flanagan book is *The Bodhisattva's Brain: Buddhism Naturalized*. It taught me a lot about Buddhism and about philosophy for the ages, but Flanagan's audacity also made me laugh a lot. The first two paragraphs are quite astonishing, and bear repeating:

Suppose we permitted ourselves this luxury: invite Confucius, Siddhartha Gautama, Mohammed, Joan of Arc, Catherine the Great, Karl Marx, Thomas Jefferson, Sojourner Truth or any other interesting or wise dead person with a view, who is representative of a tradition, into our conversations about our problems—poverty, health care, capitalism, how to be a good person, how to live well, to flourish, to be happy, and listen to what they say. This is anachronistic. Some say anachronism is bad, even that it is not allowed. Allow it.

Next imagine responding to the anachronistic answers of our respected ancestors with our own reflective standards of cogency, wisdom, and breadth and depth, feeling free to judge their answers as helpful or inadequate for our problems in our time. This is ethnocentric. It is temporally different but logically identical to judging the ideas and ways of other contemporary peoples as well suited for us or not suited for us or not good for us in our time. Some say ethnocentrism is bad, even that it is not allowed. Allow it.

Thus, Flanagan launched a deep, thoughtful, and always fresh exploration of an ancient and still pressing question: what is it to live well? Flanagan's approach never appeals to a divine authority, to the edicts of pure reason, or to some shattered algorithm that aims to solve all moral problems. He knows all those pitfalls inside out. Such knowledge motivates his advice that when we are stuck on a moral problem, we need to go into the world; there is nowhere else to go. In *The Bodhisattva's Brain*, into the world is indeed where he takes us. He realizes perfectly well that what he is doing is not "the way things are done in philosophy." Instead, Flanagan is inviting us to give our curiosity space to stretch, take in oxygen, see with new eyes, hear with new ears—and to think for ourselves. Rather like Socrates, he shows by example that we should admit our ignorance, avoid arrogance and self-deception, and all the while resist getting blinded by what is fashionable. This is philosophy worth celebrating.

November 21, 2018

Notes

1. Particular thanks are owed to John Kubie, who taught me a lot about the hippocampus, and to John Kubie and Eddy Nahmias, who helped correct an earlier draft.
2. <https://www.youtube.com/watch?v=qUvf3fOmTTk>.
3. Owen Flanagan, *The Science of the Mind* (Cambridge, MA: MIT Press, 1984; repr. 1991).

4. For a much more complete review article, see Neil Burgess, "The 2014 Nobel Prize in Physiology or Medicine: A Spatial Model for Cognitive Neuroscience," *Neuron* 84, no. 6 (2014): 1120–1125.
5. Jean Mandler, "The Spatial Foundations of the Conceptual System," *Language and Cognition* 2, no. 1 (March 2010): 21–44.
6. George Lakoff, *Women, Fire, and Dangerous Things* (Chicago: University of Chicago Press, 1987).
7. Incidentally, the *Stanford Encyclopedia of Philosophy*, typically very up-to-date, does not yet have an entry for spatial knowledge that even mentions the hippocampus, and no general entry for spatial knowledge at all. This omission suggests a bonanza for an epistemologist with pluck and vision.