Foundational Issues in Cognitive Neuroscience: Introduction

1

Gualtiero Piccinini

Contents

Abstract

Cognitive neuroscience raises several foundational issues. A first issue is how to account for our feeling that we are in control of our actions. A second issue is how to account for the relation between the mind and the nervous system. A third issue is how cognitive neuroscience supports its conclusions. A fourth issue is how cognitive neuroscience explains phenomena. A fifth issue is how to account for the notions of neural representation and neural computation.

Is the Mind Causally Efficacious in the Physical World?

To many people, the mind appears to be very different in nature from the physical stuff of which nervous systems are made. The mind seems to be made of *mental* stuff, or *experience*, stuff that strikes many people as . . . fundamentally different from physical stuff. Perhaps the mind is a nonphysical substance, or at least a collection of nonphysical properties that attaches to the nervous system. But if

G. Piccinini
Department of Philosophy, University of Missouri – St. Louis, St. Louis, MO, USA
e-mail: piccininig@umsl.edu

J. Clausen, N. Levy (eds.), *Handbook of Neuroethics*,
DOI 10.1007/978-94-007-4707-4_3,
© Springer Science+Business Media Dordrecht 2015

the mind is fundamentally different from physical substances and properties, it becomes unclear how mental states could possibly be caused by physical stimuli (as in perception) or cause physical responses (as in behavior). This is the problem of mental causation, which is discussed in ▶ Chap. 5, “Mental Causation” by Holly Anderson.

One view is that mental stuff, no matter how fundamentally different it is from physical stuff, can causally interact with physical stuff. This is called *interactionist dualism*. Interactionist dualism has the advantage of vindicating our feeling that our mind interacts with the physical world, but it has the remarkable disadvantage of making it mysterious how this occurs. No one has ever explained how something nonphysical could possibly interact with something physical. In addition, interactionist dualism implies that cognitive neuroscience cannot explain cognition and behavior—the true explanation lies in the nonphysical (mental) stuff, which presumably is not accessible to standard neuroscientific methods.

Another view is that mental stuff is causally inert after all—or at least it does not causally interact with the *physical* stuff of which our nervous systems are made, or if it is *caused by* physical events, it *causes no* physical events of its own. Maybe mental states interact with each other, but they never interact with anything physical. This view—called epiphenomenalism—frees cognitive neuroscience to explain our behavior in terms of neural mechanisms and processes, and some cognitive neuroscientists have endorsed it. Unfortunately, epiphenomenalism also implies that our mind has nothing to do with our behavior, which many people find hard to swallow.

The last possibility is that the mind is *physical*—presumably, the mind is some aspect of the nervous system and its activity. Those who think the mind is fundamentally different from physical stuff are just wrong. This view is called *physicalism*. Like interactionist dualism, physicalism vindicates our feeling that our mind interacts with the physical world. Like epiphenomenalism, physicalism implies that cognitive neuroscience can explain cognition and behavior. But physicalism raises new questions as well.

How Does the (Physical) Mind Relate to the Nervous System?

Even if the mind is physical in a broad sense (physicalism), there remains the question of how, more precisely, it is metaphysically related to the nervous system. This issue is discussed in the entry by Aizawa and Gillett.

One possibility is that mental properties are just physical properties of the nervous system. This view is called *reductionism*. All that cognitive neuroscience has to do in order to explain cognition and behavior is to find the physical properties of the nervous system that are identical to the mental properties. Reductionism has an appealing simplicity, but—at least in its strongest forms—it has a couple of disadvantages as well. First, (strong) reductionism does not sit well with the prevalent explanatory strategy within cognitive neuroscience, which is mechanistic (see below); this is because fitting a phenomenon within a multilevel mechanistic explanation involves both showing how the phenomenon is produced by a series of organized components, which sounds reductionist, *and* showing how the phenomenon contributes to a higher-level mechanism, which sounds antireductionist. Second, (strong) reductionism does not sit easily with the plausible view that cognition may be physically realizable by systems that possess no nervous system—for instance, silicon-based robots.

Another possibility is that mental properties are *realized* by physical properties of the nervous system without being identical to them. According to this view, cognitive neuroscience may still explain cognition and behavior by finding the physical properties of the nervous system that realize mental properties, but this leaves open the possibility that other physical systems—say, silicon-based robots—also realize the same mental properties in some physically different way. This view also seems to fit well the mechanistic explanatory style that prevails in neuroscience.

A third possibility is that mental properties are higher-level physical properties that emerge in nervous systems in addition to their lower-level physical properties—emergent properties are novel and irreducible to lower-level properties. This is called (strong) emergentism. Emergentism vindicates the common feeling that there is something special about our mind, but it introduces a mystery about how the emergent properties emerge. In addition, the emergent properties appear ontologically redundant: once the lower-level properties have done their causal work, there does not seem to be any causal work left over for the emergent properties to do. If this causal exclusion argument is correct, emergentism collapses back into epiphenomenalism—the view that denies the mind any influence on the physical world.

How Does Cognitive Neuroscience Support Its Conclusions?

Cognitive neuroscientists intervene on the nervous system, collect data, and draw conclusions. What this means and the epistemic risks involved are discussed in the entry by Jacqueline Sullivan.

A cognitive neuroscience experiment requires subjects to engage in a cognitive task. After that, behavior and neural activity are recorded using a variety of techniques that include both neuroimaging methods and neurophysiological recording. Computational models may also be used to understand how subjects may be able to process solve the given task. A prediction is made—for instance, about whether a certain area of the brain is involved in a cognitive task. In the best-case scenario, the data collected will adjudicate between these three competing hypotheses and point to the one that is best supported by the data.

The data from an experiment will serve their function only to the extent that they are reliable and valid. Data are *reliable* just in case they are unlikely to support false conclusions. Data are *valid* just in case they support the conclusions that they were intended to support. The intended conclusion may be about human cognition in the wild, whereas the data may be collected from rats or fruit flies operating in a constrained laboratory environment. The extrapolation from the latter to the former carries inductive risk.

Imaging techniques face a specific set of epistemic challenges, which must be handled appropriately. First, the experimental paradigms used in conjunction with imaging technology may or may not be sufficient to individuate the cognitive function under investigation. Second, typical neuroimaging techniques do not measure neural activity directly, but some correlate of it, which raises the question of what the exact correlation is between the two. Third, neuroimaging data must be processed before they are usable to draw any conclusion, and this data processing may introduce mistakes and biases that affect the reliability and validity of the conclusions. In summary, experiments in cognitive neuroscience are fraught with epistemic risks, which we should consider when evaluating the conclusions that are drawn from them.

How Does Cognitive Neuroscience Explain?

There is an old idea about scientific explanation: it consists of deriving a phenomenon from the laws of nature together with initial conditions. The entry by Kaplan discusses why this old idea is poorly suited for explanation in cognitive neuroscience and proposes a replacement. For starters, cognitive neuroscience rarely discovers or invokes anything resembling laws of nature. A better account is that cognitive neuroscience explains by providing multilevel mechanisms.

A mechanism is an organized collection of components, each of which plays a role in producing a phenomenon. A mechanistic explanation explains a phenomenon by showing how the roles played by each component, when the components are organized in the appropriate way, produce the phenomenon. Each role of each component of a mechanism is a phenomenon of its own, which may be explained mechanistically by going down one level and looking at the component's sub-components, the roles they play, and the way they are organized. By the same token, a phenomenon produced by a mechanism may play a role in a larger mechanism, whose behavior may be explained mechanistically by going up one level and looking at what a mechanism contributes to the larger system that contains it as a component.

Cognitive neuroscience explains cognition and behavior in terms of multilevel neural mechanisms. If we start from cognition and behavior, those are phenomena to be explained in terms of neural systems (cortical areas, cerebellum, brainstem, etc.), their roles, and their organization. The roles played by neural systems, in turn, are phenomena to be explained in terms of their components (cortical columns, nuclei, etc.), their roles, and their organization. Going down to even lower levels, we find neural networks, neurons, and their components. If and when we understand cognition and behavior at all these levels, we will have a multilevel mechanistic explanation of cognition and behavior.

The above picture assimilates explanation of cognition and behavior to explanation in other mechanistic sciences, such as other areas of physiology and engineering. Is there anything that distinguishes neurocognitive explanation from other mechanistic explanations?

How Does the Nervous System Think?

An important part of the explanation for how nervous systems manage to think is that they collect information from the organism and the environment, use that information to construct representations, and perform computations on such representations. Nervous systems have the distinctive function of performing computations over representations in order to control the organism—that sets them apart from most other mechanisms. The entry by Maley and Piccinini elaborates on this point.

To begin with, nervous systems possess receptors that respond to a wide variety of physical stimuli: light, sounds, odors, pressure, and more. These receptors transduce these physical stimuli into spike trains—sequences of neuronal signals that encode information about the physical stimuli. Neural signals are transmitted to the central nervous system, where they are processed. Such processing—called neural computations—extracts information that is implicit in the signals and combines such information with internal information about the organism's needs, beliefs, and desires—all of which are also encoded as states of the central nervous system. The outputs of neural computations are updated internal states as well as motor outputs—the behaviors of the organism.

This is how cognitive neuroscience explains cognition and behavior—as the outcome of computations over representations performed by multilevel neural mechanisms. Or at least, this is how current cognitive neuroscience explains some aspects of cognition and behavior.

Conclusion and Future Directions

Cognitive neuroscience is an exciting field in the middle of a growth spurt. It collects evidence at multiple levels of mechanistic organization by performing sophisticated experiments using innovative techniques. It integrates such evidence from multiple levels into multilevel mechanistic explanations. It explains cognitive phenomena mechanistically in terms of neural computations operating on neural representations.

There remains room for debate and further work on how to understand cognitive neuroscience, the explanations it gives, their relation to psychological explanations, and the role of the mind in physical world.

Cross-References

- ▶ Experimentation in Cognitive Neuroscience and Cognitive Neurobiology
- ▶ Explanation and Levels in Cognitive Neuroscience
- ▶ Mental Causation
- ▶ Neural Representation and Computation
- ▶ Realization, Reduction, and Emergence: How Things Like Minds Relate to Things Like Brains