occasion. But organisms often confront novel situations that require tailoring their basic mechanisms in new ways. To deal with these situations, control mechanisms must exhibit a degree of flexibility, directing basic mechanisms to operate in novel ways. In Section 10, we will explore ways in which control mechanisms are organized so as to support creating effective responses to novel situations.

6.6 Summary

We have introduced several different perspectives on explanation: mechanistic, dynamic, and topological. Each appeals to different factors and seems applicable to specific phenomena. This suggests a pluralistic perspective, recognizing different types of explanation. It also suggests that different perspectives might be integrated, and we offered dynamic mechanistic explanations as one integrated perspective. Lastly, we noted the importance of control in biological organisms and described how the mechanistic perspective can be extended to characterize control mechanisms.

7 What Are Levels in Neuroscience and Are They Reducible?

The term *level* is widely invoked in neuroscience, and researchers and commentators often debate whether some levels should be reduced to others. Unfortunately, the term level is used in a wide variety of senses. In this section, we differentiate three notions of level that are prominent in discussions about neuroscience and identify the implications of each for reduction.

7.1 Marr’s Levels (Perspectives)

David Marr, a pioneer in the development of computational modeling in neuroscience (Section 3.5), began his book *Vision* (1982) with a critical assessment of what he saw as the current state of the discipline. Neuroscientists were accumulating many findings about how various parts of the nervous system operate using techniques such as those discussed in Section 3. But they were making little progress in providing an understanding of how the brain works. On his analysis, this was due to focusing on just one level, which he termed the *hardware implementation* level. Accounts at this level focus on parts of the brain and how each operates. To make progress in understanding the brain, he argued for the need for two other levels: those of *representation and algorithm* and of *computational theory*. At the representation and algorithm level, he argued that researchers should treat the parts of the brain as representing content and applying rules to manipulate those representations. Much of Marr’s own work was focused on the representation and algorithm level as he attempted to describe visual processing in terms of states that represent specific features of stimuli and algorithms that specify how the brain operates on one representation to generate another (see Section 8). But the most novel, and arguably the most important, of his levels was that of computational theory. At this level he proposed that researchers should address questions such as “What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?” (p. 25).

To address the goal and appropriateness of neural computation, it is not sufficient to look inside the nervous system any more than one can determine the goal and appropriateness of what goes on in a computer by just looking inside it. With the computer, we turn to the users and the tasks for which the computer is used. The comparable move with respect to the nervous system is to look to the environment in which it works – the organism and the physical and social environment in which the organism operates. As Marr’s focus was on vision, the relevant environment is the visual world. His contention is that by analyzing the structure of the visual world, we can understand what the visual system needs to do. Marr notes that James Gibson, a psychologist with whom he mostly disagreed, came closest to understanding what it meant to focus on the computational level. Gibson (1979) argued that visual experience does not consist of independent pixels but is highly structured. An illuminating example is that if an object is approaching you, or you it, it expands in your visual field. If a ball is expanding equally in all directions, then you may intersect it. Whether you do and how soon depends on how fast it expands. Studying the challenges an organism faces in its environment, Marr insisted, is critically important to understanding what a nervous system is doing and whether what it does is appropriate (Shagrir, 2010).

Marr’s levels might better be glossed as perspectives that investigators need to take in understanding the nervous system. One perspective is to focus on parts of the nervous system and how each is operating. A second focuses on the procedures that the brain is using to process sensory information or generate actions. The third focuses on the organism and its environment in order to identify the tasks that the nervous system must perform if the organism is to be successful. Each perspective may inform one of the others (knowing the task that the nervous system must perform can guide the search for representations and algorithms that operate over them). However, one perspective cannot provide the insights provided by another. Rather than reducing one level or perspective to another, Marr argues that researchers need to pursue all three.

7.2 Mechanistic Levels and Reduction

The framework of mechanistic explanation (Section 6.1) brings a different conception of levels: the components of a mechanism can be viewed at a lower level than the mechanism itself. Likewise, mechanisms can be constituents of mechanisms at higher levels. This look down to lower levels and up to higher levels iterates due to the fact that mechanisms are often parts of yet higher-level mechanisms and their parts are themselves mechanisms consisting of parts. The notion of part and whole is fundamental to this notion of level. As a result of the process of decomposing mechanisms into other mechanisms, mechanistic levels are hierarchical.

In decomposing mechanisms, researchers are going down levels. How far should they go? There is a long tradition in science, referred to as *reductionism*, which argues that explanation should appeal to as low a level as possible because lower levels are, in some sense, more basic. For some theorists, the most basic level is that of fundamental physics; for these theorists, the ultimate goal is to explain all happenings in the universe, including those in our brains, in terms of the entities and activities of the most basic physical particles. For now, there seems to be little prospect of explaining biological phenomena in such terms. Nonetheless, some theorists argue that the goal should be to explain behavior and cognition at the lowest level possible. Bickle (2006) defends what he terms *ruthless reduction*: intervene at the lowest level possible (currently the molecular level) and measure effects at the behavioral level. When successful, this provides evidence that entities at the molecular level are causally efficacious and so explain the behavioral phenomenon.

Mechanistic explanations are clearly reductionist in one sense: they appeal to the components of a mechanism in explaining its behavior. And insofar as the components of a mechanism are themselves mechanisms, they support going to yet lower levels. This is conveyed in Figure 21: to explain the behavior of a mouse navigating the Morris Water Maze (Section 5.2), neuroscientists appeal to the hippocampus as the locus of spatial maps, to synapses between neurons in which chemical changes are realized (long-term potentiation), and finally to NMDA (N-methyl-D-aspartate) receptors in the postsynaptic membrane. Yet, it is important to realize that at each lower level, the component is one part of a mechanism: A single synapse does not realize a cognitive map, but only in the context of other neurons in the hippocampus. Indeed the hippocampus does not operate on its own but only in the context of a larger mechanism involving other brain areas such as the entorhinal cortex (Bechtel, 2009). At each of these levels, components are organized into larger mechanisms that do things that the components cannot do. While appealing to lower levels, mechanistic accounts

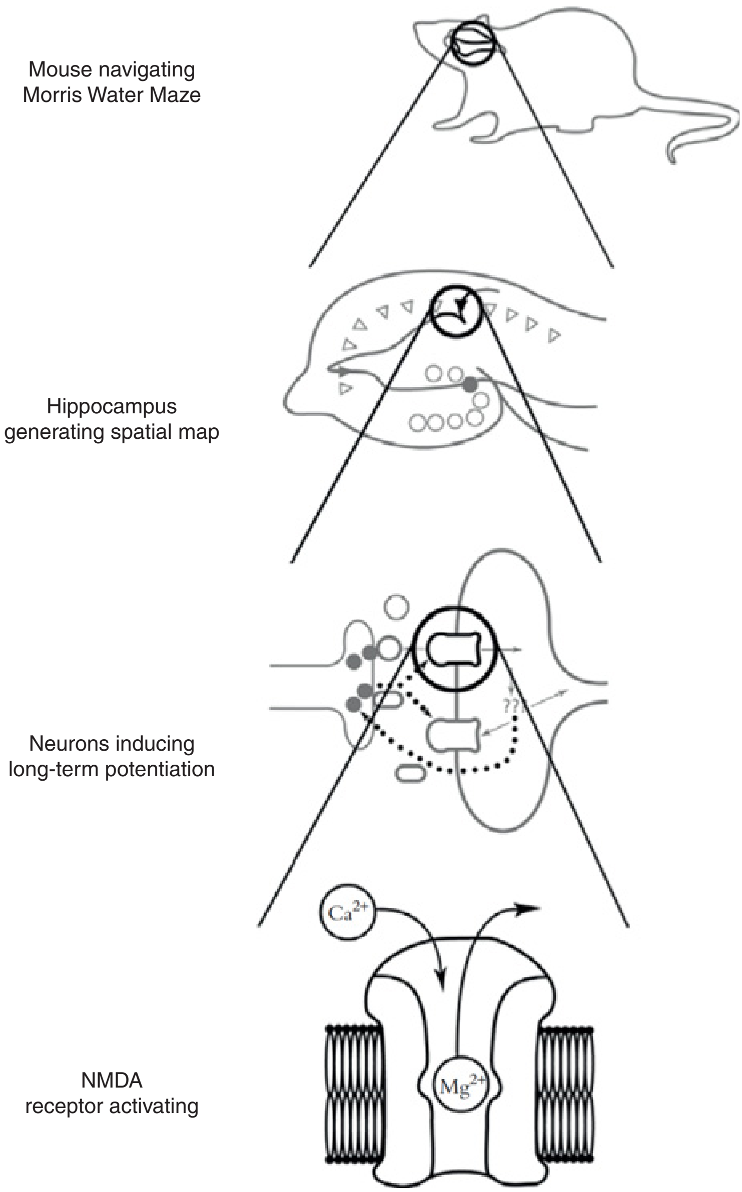

Figure 21 Multiple mechanistic levels invoked in explaining memory. Reprinted from Craver (2007) with permission of Oxford University Press.

do not privilege any one of them. Rather, mechanistic accounts emphasize equally how components are integrated into larger systems. One reason it is important to go up levels is that in many cases, lower-level components perform different operations when they are part of different higher-level wholes.

7.3 Levels of Control

One of the factors that can make a lower-level mechanism behave differently is that control mechanisms (Section 6.5) operate on it. Control mechanisms give rise to their own relation of levels. Insofar as a control mechanism operates on and changes the parts or operations of another mechanism, one can view it as at a higher level. And insofar as another control mechanism operates on it, that control mechanism is at a yet higher level. But this relation differs from the relation between mechanistic levels in two respects. The control mechanism is not a whole containing the controlled mechanism. And although the relation between control mechanisms can be hierarchical, it need not be (see Section 10.2). Whether hierarchical or not, levels of control do not give rise to reduction, as each control mechanism has its own role to play in coordinating the activity of the mechanisms it controls.

7.4 Summary

We have identified three notions of level that figure in discussions of neuroscience. Only the mechanistic conception gives rise to a notion of reduction. Some theorists advocate advancing explanations at the lowest level possible. Most proponents of mechanist explanation, however, emphasize the importance of organization at each level, and so recognize contributions of both lower and higher levels.

8 Do Neural Processes Represent Anything?

Representation is perhaps the most contested term in philosophical discussions of neuroscience. A representation stands in for and can be used by a process as a surrogate for something it represents. Humans frequently make use of representations. The phrase “nervous system” represents nervous systems and is used in many sentences in this Element. Many of the figures in this Element represent parts of the nervous system while others represent processes thought to occur in it.

On many accounts, such as the one advanced by Marr (Section 7.1), the nervous system is a computational system. A key component of a computational account is that operations are performed on representations (philosophers often refer to this as the *computational theory of mind* – see Pitt, 2020). Accounts of neural activity commonly characterize that activity in terms of what it is supposed to represent. Thus, place cells (Section 5.2) are characterized as representing places – when they spike in sequence, they indicate a sequence of locations on an animal’s route. Neurons in different regions of the visual