10. How is the whole nervous system organized?

In section 6.5, we introduced the notion of control mechanisms and noted that the nervous system consists of control mechanisms. In section 7.3, we raised the issue of how control mechanisms, in general, are organized. In this section we focus on how neural control mechanisms, in particular, are organized, considering two alternatives: organization in hierarchical pyramids or into heterarchical networks.

10.1 Hierarchical pyramid organization

In many social systems, such as corporations and the military, control is organized hierarchically in a pyramid, as in figure 24A. In this arrangement, multiple local controllers report to a smaller number of controllers at a higher level. This is iterated, culminating in a chief executive, a president or a general, with whom the buck stops (to use an expression coined by the U.S. President Harry Truman). In the control hierarchy, the lower-level control mechanisms are, adopting Dennett's (1991) term, bureaucratic. They function to provide information for the central control mechanism or to work out the implementation of its commands. On such a scheme, there are in principle no conflicts between lower-level and central controllers, since the lower-level controllers do not have their own agenda and function solely to serve the central controller.

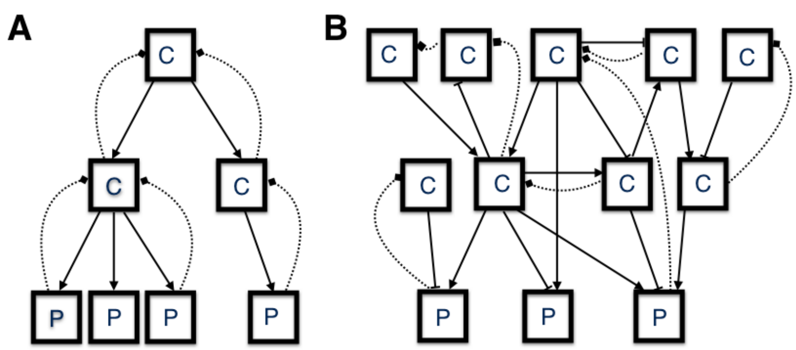

Figure 24. A. a typical hierarchical pyramid of control processes with information (dotted arrows) directed towards higher layers and commands directed to lower levels. Sensory inputs are represented by dotted arrows. Arrows and edge-ended lines indicate excitatory and inhibitory control. B. Heterarchical organization that violates several features of a hierarchical pyramid.

The nervous system is often conceptualized as organized as a hierarchical pyramid. Local nerve nets and pattern generators are brought under the control of individual ganglia (section 2.3). As a central brain evolved, peripheral ganglia were brought under the control of more central ones such as the nuclei in the vertebrate brain. Then, as the neocortex evolved, it assumed control: sensory inputs are fed up to it and it sends motor commands back down to subordinate levels of control. Except for the positioning of the basal ganglia and thalamus, Figure 15 presents such a picture.

One consequence of the pyramid structure is that entities at each level in the hierarchy face greater demands. In order to direct effectively a company or a military structured this way, the chief executive needs to acquire all the relevant information and use it in making decisions. Pyramid organization is also conservative as innovation is permitted only at the top. As a result, some social institutions relax the hierarchical pyramid, delegating to lower-level controllers the ability to act independently. Coordination between controllers then becomes a challenge. To some degree, this can be achieved by direct interaction of controllers at a given level (as between ganglia in the leech example discussed in is section 2.3). A further departure from the pyramid organization is to forego the strict layering of levels, allowing information and commands to skip layers or to be asserted within layers. At the extreme, there is no executive at the top of the pyramid and there are more control agents as one moves up levels (Figure 24B). McCulloch (1945) introduced the term *heterarchy* for situations in which one's preferences are not organized hierarchically (e.g., one prefers A to B, B to C, and C to A), and Pattee (1991) extended it to control systems in living organisms that significantly depart from the hierarchical pyramid organization.

10.2 Heterarchical network organization

The neocortex is generally understood to be organized hierarchically. But given its multiple processing streams (illustrated Figure 14B), it does not conform to the pyramid structure. When we turn to the nervous system more generally, departures from the hierarchical pyramid mode of organization become more notable. This is evident when we consider some of the organisms we discussed in other sections. In the jellyfish (section 2.2), multiple different control mechanisms, each responsive to conditions that require alterations in the default mode of swimming, act on the same nerve network. The various ganglia in the leech (section 2.3) and in *C. elegans* (section 4.2) act in relative independence. Despite becoming co-located in the brain over the course of evolution, the ANS and BNS (section 2.3) continue to operate independently to a significant extent. Regions of the hypothalamus, the arcuate nucleus, the lateral zone, and the SCN, each regulate different behaviors (feeding, attention vs. sleep, circadian rhythms respectively). The neocortex cannot dictate circadian rhythms (except by directing actions in the world such as exposing oneself to daylight).

How do control structures in the brain become organized in a heterarchical network? For a clue, consider one of the principles we discussed in section 6.4: in many real-world networks degree (number of connections per node) is distributed accordingly to a power law (section 6.4) in which a few nodes are extremely highly connected. Barabási and Bonabeau (2003) proposed a process by which networks come to exhibit this feature: in many contexts, when an edge is added from a given node to another, it is more likely that it will connect to an already highly connected node. We can couple this with a further factor—in control networks new edges are introduced to achieve better control, especially to overcome a shortcoming of the current control system, and this often involves connecting to another control mechanism. The process parallels that in revising computer software. Unlike in organisms, the initial code of a computer program may have been intelligently designed. But no matter how intelligently designed and well tested it is, it will likely fail in some contexts. When it does so, programmers do not redesign from scratch, but patch the current code by adding whatever new code will address the situation without obviously compromising other parts of the program (these are often referred to as *kludges*). Biological networks, including neural networks, evolve in the same manner—retaining new connections that happen to form between existing components when they will improve the organism's performance in the current environment. The check on these new connections is much like the check on software kludges—does the additional connection enhance the ability of the organism to maintain itself (or at least do not render it much less likely to survive)?

In Figure 24 we presented the heterarchical alternative to the hierarchical pyramid in terms of layers arranged hierarchically. A reason for maintaining a representation in terms of layers is that individual control mechanisms operate on specific other mechanisms and this seems well represented by putting a control mechanism at a higher level than the mechanism it controls. But the hierarchy is already breaking down with the inclusion of edges between nodes in the same layer and between nodes two layers apart. Adopting the perspective that evolution adds kludges to an existing control system, such departures from hierarchy are to be expected as there is no principled reason to maintain hierarchical organization. One could also add connections so that lower-level nodes act to control those at a higher level. Ultimately, control systems, especially those that have evolved, are better represented as networks.

10.3 How to achieve coherent and intelligent control in a heterarchical network?

One reason many people find the hierarchical pyramid structure to be intuitive is that it solves two important issues of control: coherence and intelligence. The issue of coherence concerns how neural control mechanisms manage to produce more or less coherent behaviors. Having a central executive making the call with other lower-level controllers merely implementing the details certainly seems like a good way to ensure coherence of behaviors. The issue of intelligence concerns how control mechanisms can generate adaptive behaviors in a wide range of novel contexts. As we discussed earlier, control mechanisms need to adjust their basic mechanisms to operate appropriately in a situation. Then, these different basic mechanisms can work together to produce adaptive behaviors. But how do control mechanisms know the appropriate control signals, especially for novel situations they have not encountered before? The central executive is supposedly the source of intelligence. It is often assumed to possess rich information about the world and to operate on this information in order to exercise context-appropriate control over more basic mechanisms. However, the brain does not appear to have a hierarchical pyramid structure. This requires us to address the question: how can heterarchical networks achieve coherent and intelligent control of organisms?

Theorists have identified strategies that enhance coherence between multiple control mechanisms (Clark, 2014). For example, the fact that multiple controllers all confront and get feedback from the same external world promotes coherence. Also, various neuromodulators and hormones dissipate broadly in neural and bodily systems, communicating information about conditions inside and outside the organism with different control and production mechanisms. Further, local communications between controllers can reduce conflicts between the decisions made by different controllers. We noted this with the leech (section 2.3): even though the decision to walk or swim is made in each ganglion, the control mechanisms in different ganglia act on each other so that they arrive at a coherent action. Finally, of special importance in promoting coherence are the basal ganglia as they provide a common structure in which evaluations made by multiple controllers distributed across the brain are brought together and outputs are sent back to these different controllers, modulating their individual behaviors. By forcing the integration of these initially diverse inputs, the basal ganglia have the infrastructure to enhance the coherence of the control decisions.

Outside of neuroscience, there are numerous examples of how intelligent behaviors can emerge from a network of controllers, each of which has only partial access to the information about the world. Marvin Minsky (1986), a pioneer in artificial intelligence, offered the metaphor of a "society of mind" that fits well with the view of the nervous system as a network of heterarchically organized control mechanisms.[1] Could a society give rise to a self that acts in a unified, intelligent way? Examples such as honey-bee swarms suggest that integrated intelligent behavior can arise from collective activity.[2] Modern democratic societies embrace the idea of determining courses of action through voting. In the following, we sketch how a heterarchical nervous system could produce intelligent behaviors by operating like a political democracy (see Huang, 2017, for a more detailed account).

Appealing to what is called the "wisdom of the crowd effect," social choice theory argues that under the right conditions, aggregating the decisions of multiple individuals (e.g., through voting) is more likely to result in a correct decision than relying on an individual's decisions. This can be demonstrated mathematically (List, 2013). Consider three people, each of whom has an independent reliability of 0.8 of getting the answer right. If the three people vote and go with the majority, the reliability increases to 0.896.[3] The likelihood of correctness increases further as more people are included. This result can be generalized to a wide range of conditions. An explanation for this is that the collective is integrating relatively reliable information stemming from different sources. This result applies as well to the nervous system. Multiple nuclei can integrate different sources of information, arriving collectively at a more reliable decision than if the organism relied on just one. In the vertebrate nervous system, the basal ganglia are organized to determine the winner in a manner comparable to voting: the inputs represent assessments of different alternatives and the competition to control the direct and indirect pathways culminates in a Go/NoGo decision. Thus, there is reason to think that a heterarchical system that makes decisions through processes such as voting could give rise to coherent and intelligent agency. If it can, there is less reason to assume that control mechanisms must be organized hierarchically.

---

[1] We develop the society of mind metaphor in terms of control—the human mind is composed of a massive number of control mechanisms. For a development in terms of representations, see Rupert (2011).
[2] For a relevant discussion on whether we could treat different types of distributed systems, such as bee swarms or human society, as a cognitive system, see Huebner (2014).
[3] The result is reached by adding up the probabilities of four different scenarios where the group gets the answer right ($0.8 \times 0.8 \times 0.8 + 0.8 \times 0.8 \times 0.2 + 0.8 \times 0.2 \times 0.8 + 0.2 \times 0.8 \times 0.8 = 0.896$).

10.4 Summary

Accepting that neural systems operate as control mechanisms, in this section we have considered two alternative patterns of organization: a hierarchical pyramid or a heterarchical network. Despite the challenge of maintaining coherence and achieving intelligence in a heterarchical scheme, the nervous system does appear to be organized heterarchically. We have briefly considered how coherent and intelligent actions might be generated in a heterarchically organized system.

11. What does neuroscience teach us about who we are?

We conclude this Element by considering a neurophilosophical question (section 1): what does the knowledge acquired in neuroscience, some of which we have reviewed in this Element, tell us about ourselves? When most people are asked to characterize themselves, they begin with traits such as race, sex, gender, age, height, hair color, etc. As biological organisms, they appeal in part to their history—they were born at a certain moment from specific parents and have followed a trajectory across the earth. Many will refer to important events in their personal history, their job or profession, and preferred activities. But beyond that, many people think that these are just the external expressions of something internal: their self. On the Temple of Apollo at Delphi in Greece was inscribed the injunction: "Know thyself." Given the role of the nervous system in controlling behavior, part of following the Oracle's injunction might be to consider who we are in light of what is known about our nervous system.

If one views the nervous system as a hierarchical control system with a central executive issuing commands, the executive implemented in the prefrontal cortex might be the self we need to know. As we have seen throughout this Element, however, the nervous system is organized heterarchically, with relatively independent control mechanisms located in many different parts of the nervous system. Activities such as eating, sleeping, and reproducing are controlled by nuclei in the hypothalamus. These are important features of who we are. Memories, especially memories for events in our lives, are developed in the hippocampus and ultimately laid down in the neocortex. Who we are is often revealed in our choice of actions, in which the basal ganglia play a major role. If the nervous system is really organized heterarchically, what sense can be made of *a* self to be found in our nervous system?

11.1 Reporting on our mental lives

In this section we explore what might seem a radical hypothesis: that there is not a self to learn about; rather each of us, drawing upon our ability to use language, constructs one. The key idea draws from Wilfred Sellars' (1956) Myth of Jones, which he offered, not as an historical account, but as a means to illustrate the status of our reports on our own mental states. According to his myth, before people developed the ability to describe their own mental states, they had developed natural science and made successful predictions about entities in the world. To explain how organisms behaved, they posited mechanisms within them. One scientist, Jones, turned this technique on human behavior: he hypothesized inner entities he called *thoughts* and developed a theoretical framework with which he could successfully predict how different people would behave on the basis of the thoughts he attributed to them. Other people learned to use this theoretical framework and even applied it to themselves, initially inferring what they thought based on their overt behaviors. Then Jones taught Dick to describe his thoughts without first consulting his overt behavior in much the way neuroscientists train non-human subjects to carry out tasks in their experiments—by giving him positive and negative feedback when his self-ascriptions fit those Jones made based on Dick's behavior. Dick successfully learned to do this; although he had no idea how he does so, he reports on his own decision making in terms of his thoughts. Sellars' point is that although there must be some basis on which Dick succeeds in doing this, there is no need to view him as reporting or "introspecting" internal states. Rather, he has learned to extract patterns, which, as we saw, is the forte of the neocortex.[4]

The theoretical framework Jones developed and Dick learned corresponds to what philosophers refer to as *folk psychology*. It characterizes humans in terms of attitudes towards what are often referred to as *propositions*. The idea is that in one's mental life one represents information in propositions which, for our purposes, we can treat as statements in a language such as English. An example proposition would be "My next-door neighbor has a cat." According to folk psychology, one can adopt different attitudes towards this proposition—one can believe that it is true, doubt that it is true, fear that it might be true, wish that it were true, etc. Further, one can reason in terms of these propositional attitudes: from the propositional attitudes of believing "there is yogurt in my refrigerator," "I would like to eat yogurt," and "If I go to my refrigerator, I can eat what is in it," one can infer "I should go to my refrigerator." On Sellars' account, one need not treat folk psychological statements as describing neural events. They are constructs in a story we tell about other people and ourselves. Nonetheless, in terms of them, we can provide useful accounts of how we and others behave. Moreover, we can update these accounts when they go astray and make better predictions in the future.[5] For example, if you think your friend believes it is going to rain, wants to stay dry, and believes bringing her umbrella will enable her to stay dry, you can infer that she will have her umbrella with her. If she shows up without it, you can inquire whether she didn't believe it was going to rain, didn't want to stay dry, or wasn't acting rationally.

11.2 Making norms for action explicit and living by them

As we discussed in section 6.5, control mechanisms can be viewed as implementing norms through their response to the measurements they make. In the case of the Watt governor, the norm that is implemented is built into its design—it was designed to maintain a steam engine at

---

[4] This analysis is developed further in Bechtel (2008, chapter 7). For recent philosophical and scientific development along this line, see Schwitzgebel (2019).

[5] Eliminativists such as Paul P. M. Churchland (1981) see this as a reason to repudiate folk psychology as a false theory.

a specific speed. The various control mechanisms in the brain likewise implement norms that are incorporated in them either through evolution or through learning. When we think of norms we often focus on moral norms. These are norms which we take to be capable of being expressed in language, rationally discussed, and chosen, at which point they influence one's actions. How do these moral norms relate to those implemented in neural control mechanisms in organisms?

The account of how we can report on our mental lives sketched in the previous section provides a framework for representing our mental life in language, but it doesn't directly address how the results of explicit adoption of norms can become efficacious. One might infer that if accounts of mental processes were constructions, they could not be efficacious. But this is wrong: we construct our account of our mental activities using our brains. As we construct such an account, we alter processes in our brains. This applies as well to our discourse about norms. As we adopt norms and remind ourselves of them, they may direct our behavior (Frankish, 2004). At present we know very little about the brain processes that figure in these activities and so are not in a position to spell out how they affect behavior. But given their reliance on language, it seems reasonable to assume that these processes can affect brain mechanisms that are engaged in selection of behavior.

We are aware that sometimes even when we commit ourselves explicitly to a given norm, we will violate it. We commit ourselves to leaving the party by 11pm, but end up staying to 1am. Philosophers characterize this as *weakness of the will*. Given what we have said about the heterarchical organization of the brain, this phenomenon is not surprising. Other neural control mechanisms compete and win out. But that doesn't mean explicit commitments to norms are always inefficacious. As emphasized by many moral theorists, one way to make our moral commitments efficacious is to turn them into habits by, for example, giving ourselves a reward when we fulfil our commitments repeatedly. Another is to draw attention to our commitments. If you commit yourself to a norm publicly, that might strengthen that norm when it is in competition with others. Or others may remind us of our commitment (for a recent development of the mind-shaping effect of the social practice of articulating our values, see McGeer, 2015).

11.3 Constructing a self

What should we make of the Delphic Oracle's injunction in light of our discussions of heterarchy and the constructive character of our accounts of our mental lives? On the account we have offered, one's concept of oneself would also be a construct.

Out of what do we construct ourselves? For most people, memories of episodes in our past are important elements. Tulving (1983) not only coined the term *episodic memory* for these memories but characterized them as mental time travel, thereby capturing how recalling an event seems like reliving it. Although at least with vivid memories people often have the sense of simply rehearsing the past, there is compelling evidence that recalling a memory does not involve retrieving a record of the past but rather reconstructing the past event from multiple sources of information. Such reconstruction is often affected by what happened on previous occasions during which one recalled the event (as shown by Loftus, 1975, a detail suggested by someone else during the recall can be subsequently remembered as part of the initial event). This ability to put material together into a past narrative also enables us to project ourselves into the future, characterizing ourselves in terms of what we hope to become. Beyond memories, we also characterize ourselves in terms of traits and abilities we take ourselves to have and norms we hope to uphold.

If our self is a construct, then perhaps the meaning of the Delphic oracle is that each of us needs to construct our self-concept, one that presents a narrative of our past, projects ourself into the future, and frames who we are. We can draw upon it in making major life decisions. Those decisions will also contribute to constructing our self-concept in the future. Unfortunately, our self-concept won't always be effective in guiding the decisions we make. Our nervous system is still heterarchical, and control mechanisms, whether linked to our self-concept or not, will generate many of our actions more or less independently. But, like the moral norms we discussed in the previous section, our self-concept can be invoked to direct and constrain the decisions we make, thereby providing focus to our lives.