12 It's All about Complex, Entangled Networks

We've come to the end of our short exploration of the brain, this most mysterious of organs. This book has tried to illustrate how mental processes are built from intricate interactions involving gray and white matter components. We've learned, hopefully, to appreciate some of the complexity of how the brain contributes to bringing forth the mind. In this last chapter, we return to some of the big questions and problems encountered previously, and some of the themes that will be important for advancing our understanding of mind and brain in the future.

What Does Brain Evolution Mean?

The geneticist Theodosius Dobzhansky famously stated that in biology, nothing makes sense unless it's in light of evolution. The same applies to neuroscience, a biological science. But evolution poses a conundrum. Vertebrates have been evolving for over 500 million years. A telencephalon, a midbrain, and a hindbrain are part of the general plan of their nervous system. Structures like the amygdala and the striatum are found in animals as diverse as a salmon, a crow, and a baboon. Thus, many parts of the brain are "conserved." But then, what is novel? Something *must* be new, after all.

In chapter 9, we reviewed how *homology* refers to relationships between traits shared as a result of common ancestry. The leaves of plants provide a good example.[1] The leaves of a pitcher plant, Venus flytrap, poinsettia, and a cactus look nothing alike and, in fact, have distinct functions. In the pitcher plant, the leaves are modified into pitchers to catch insects; in the Venus flytrap, they turn into jaws to catch insects; a poinsettia's bright red leaves resemble flower petals and attract insects and pollinators; leaves on a cactus plant have become modified into spines, which reduce water loss and can protect the plants from herbivores. Nothing alike—yet the four are homologous because they derive from a common ancestor.

A structure adopts new functions during evolution, while its ancestry can be traced to something more fundamental.[2] Take the hippocampus of rodents, monkeys, and humans. There is copious evidence indicating that the area is homologous in the three species—that is, it's a conserved structure. But does it perform the same function(s) in these species, or does it carry out *qualitatively* different function(s) in humans, for example? To many neuroscientists, this sounds implausible. However, the possibility need not be any more radical than saying that the forelimb does something qualitatively different in birds compared to turtles, say. If common ancestry precluded new functions, no species could ever take flight!

The ongoing discussion is particularly pertinent when we think of emotion and motivation, because researchers invoke "old" structures when studying these mental phenomena. Regions like the amygdala at the base of the forebrain and the periaqueductal gray (PAG) in the midbrain are invoked in the case of emotion, and the accumbens (part of the striatum) also at the base of the forebrain and the ventral tegmental area in the midbrain in the case of motivation. Because these regions are deeply conserved across vertebrates, they function in a similar way, or so the reasoning goes. If we entertain these areas in rodents, monkeys, and humans, closer as they are evolutionarily, the expectation would be that they work similarly. But rodents and primates diverged more than 70 million years ago. Are we to suppose that no *qualitative* differences have emerged? This seems rather implausible. (In chapter 9, we briefly reviewed some structural differences in the amygdala of rats, monkeys, and humans.)

The argument made in this book is that we should conceptualize evolution in terms of the *reorganization* of larger-scale connectional systems. Instead of more cortex sitting atop the subcortex in primates relative to rodents—which presumably allows the "rational" cortex to control "primitive" parts of the brain—more varied ways of interactions are possible, supporting more mental latitude.

The brain doesn't fossilize. Unfortunately, with time, it disintegrates, leaving no trace. So we simply don't have a way to know exactly what the brain of a common ancestor looked like. Without fossil remains, scientists tend to think of the brain of a common ancestor of rodents, primates, and humans as something like the current brain of a mouse, as this animal is the "rudimentary" one. But a mouse encountered today has had 75 million years to evolve from the ancestor in question, ample time to specialize to the particular niches it inhabits now.

Evolution is as much about what's preserved as what's new. Ever since science was transformed by the independent work of Charles Darwin and Alfred Russel Wallace in the late 1850s, biologists have sought to determine "uniquely human" characteristics. This has led to a near-obsession to identify one-of-a-kind nervous system features, from putative exclusively human brain regions to cell types. The cortex, in particular, has attracted much attention. The pallium of mammals is structured in a layered fashion, a quality not observed in other vertebrates. Well, not exactly, as some reptiles (such as turtles) have a dorsal pallium that is cortex-like, with three bands of cells. Mammals, however, have parts of the cortex that are much more finely layered, with six well-defined zones. In fact, a six-layered cortex is often referred to as "neocortex," with the "neo" part highlighting its sui generis property (in the book, the more neutral terminology "isocortex" was used in chapter 9 for this type of cortex).

I believe that the concept of reorganization of circuits is a much more promising idea. That is to say, what is unique about humans is the same that is unique about mice, or any other species: Their circuits are wired in ways that support survival of the species. This is not to deny that some more punctate differences play a role. But whatever the differences are, at least considering primates with larger body sizes, they are not staring us in the face—they are subtle. For example, all primates exhibit an isocortex that is massively expanded. Primates also have prefrontal cortices with multiple parts, including the lateral component, which neuroscientists often link to "higher cognitive" capabilities. More generally, direct evidence for human-specific cortical areas is scant.

Let's go back to Dobzhansky's call to consider biology in light of evolution—always. Biologists would vehemently agree. But evolution is so egregiously complex that the suggestion doesn't help as much as one would think. What we observe in practice is that neuroscientists who don't specialize in studying brain evolution are time and again cavalier, if not outright naive, about how they apply and think of evolution. By doing so, our explanations run the risk of becoming *just-so stories*.[3]

Fitting Behavior Inside a 40 × 40 × 40 Centimeter Box

The central question in neuroscience is to understand the physical basis of behavior. But what kinds of behavior can be studied in a lab? Mice and rats can be placed in chambers and mazes to perform tasks. One can then study

the effects of lesions on behavior. If cell recordings are conducted, the constraints are even more severe. Until just a few years ago, this required a fair amount of cabling to link the brain to signal amplifiers and other electronics. Experiments in primates are performed in a "monkey chair" that keeps the animal's body and head in place. Humans, of course, are studied inside magnetic resonance imaging (MRI) tubes that are anything but organic. With the technology available, getting closer to natural behaviors has simply not been possible.

A type of behavior that fits inside a $40 \times 40 \times 40$ centimeter box is classical conditioning. Indeed, it has been extensively studied by psychologists since the early twentieth century, and for those interested in the biological mechanisms of fear, the paradigm has been a godsend. It offers a window into this process while allowing careful control over study variables, a fundamental consideration in experimental science. The neuroscience of fear has been one of the most active areas of inquiry, thanks to the paradigm.

But the fixation with this task has led to a form of tunnel vision.[4] As Dennis Paré and Gregory Quirk, very prominent researchers in this area, state:

> When a rat is presented with only one threatening stimulus in a testing box that allows for a single reflexive behavioral response, one is bound to find exactly what the experimental situation allows: neuronal responses that appear tightly linked to the CS [conditioned stimulus] and seem to obligatorily elicit the conditioned behavior. (Paré and Quirk 2017, 6)

The very success of the approach has led to shortsightedness.

Placed inside a small, enclosed chamber the animal is limited to a sole response: Upon detecting the CS, it ceases all overt behavior and freezes in place. It can't consider other options, such as dashing to a corner to escape; it cannot try to attack the source of threat either, as there isn't another animal around—the shock comes out of nowhere! Now, when researchers study the rat's brain under such conditions, a close relationship between brain and behavior is established. But as Paré and Quirk warn, the tight link might be apparent insofar as it would not hold under more general conditions.

Neuroscience is experiencing a methodological renaissance. Advances in chemistry and genetics now allow precision in targeting regions and circuits in ways that would have sounded like science fiction a decade ago. But if we continue using the paradigms that have been the mainstay of the field, we will be cornering ourselves into a scientific cul-de-sac.[5] It's time to think outside the box.

A Thought Experiment

Try to contemplate a future device that allows registering in minute detail the behaviors of a cheetah and a gazelle during a chase, including all muscle and skeletal movements. At the same time, we are capable of recording billions of neurons across the two nervous systems while the entire chase unfolds from before the cheetah initiates the pursuit until its dramatic conclusion. What would we discover? How much of our textbooks would have to be altered?

A radical rethinking might be needed, and a lot would have to be rewritten. Alternatively, many of the experimental paradigms employed to date are quite effective in isolating critical mechanisms that reflect the brain's functioning in general settings. True, novel findings would be made with new devices and techniques, but they would extend current neuroscience by building naturally on current knowledge. The first scenario is not idle speculation, however.

So-called naturalistic experimental paradigms are starting to paint a different picture of amygdala function, for example. In one study, a rat was placed at one end of an elongated enclosure and a piece of food placed midway between the rat and a potential predator, a Lego-plus-motor device called a "Robogator" (figure 12.1) (Amir et al. 2019). To successfully obtain the food pellet, the rat had to retrieve it before being caught by the Robogator (don't worry, capture didn't occur in practice). The findings were inconsistent with the standard "threat-coding model" that says that amygdala responses reflect fear-like or other related defensive states. During foraging, when the rats were approaching the pellet, neurons reduced their firing rate and were nearly silent, not active, close to the predator. Clearly, responses did not reflect a threat per se.

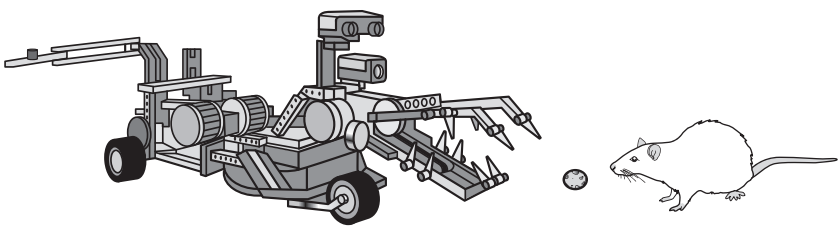

Figure 12.1
More naturalistic experimental paradigms are starting to be employed. Here, a rat can acquire a piece of food but must catch it before being chased by the "Robogator." *Source*: Model for the "Robogator" at left kindly provided by Jeansok J. Kim.

Another study recorded from neurons in the amygdala over multiple days as mice were exposed to different conditions (Gründemann et al. 2019). The mice were exposed to a small open field and were free to explore it. These creatures don't like to feel exposed, so they stayed at the corners of the box a good amount of time. But they also ventured out into the open and walked around the center of the box with some frequency. The researchers discovered two groups of cells: one engaged when the mouse was being more defensive in the corners (these "corner" cells fired vigorously at these locations), another when the mouse was in an exploratory mode visiting the center of the space ("center" cells fired strongly when the animal was around the middle of the field). The researchers recorded from the exact same cells during more standard paradigms, too, including fear conditioning and extinction. They then tested the idea that the firing of amygdala neurons tracks "global anxiety." For instance, they should increase their responses when the animal entered the center of the field in the open-field condition, as well as when they heard the CS tone used in the conditioning part of the experiment. Surprisingly, cells did not respond in this way. Instead, neuronal firing reflected moment-to-moment changes in the exploratory state of the animal, such as during the time window when the animal transitioned from exploratory (for example, navigating in the open field) to nonexploratory behaviors (for example, when starting to freeze).

The above two examples provide tantalizing inklings that there is a lot to discover—and revise—about the brain. It's too early to tell, but given the technological advances neuroscience is witnessing, examples are popping up all over the place. For example, a study by Karl Deisseroth and colleagues recorded activity of approximately 24,000 neurons throughout 34 brain regions (cortical and subcortical).[6] Whereas measuring electrical activity with implanted electrodes typically measures a few cells at a time, or maybe about 100 by using state-of-the-art electrode grids, the study capitalized on new techniques that record calcium fluorescence instead. When cells change their activity, including when they spike, they rely on calcium-dependent mechanisms. In genetically engineered mice, neurons literally glow based on their calcium concentration. By building specialized microscopes, it is possible to detect neuronal signaling across small patches of gray matter. In their study, when mice smelled a "go" stimulus, a licking response produced water as a reward. The animals were highly motivated to perform this simple task as the experimenters kept them in a water-restricted state. Water-predicting sensory stimuli (the "go" odor)

elicited activity that rapidly spread throughout the brain of the thirsty animals. The wave of activity began in olfactory regions and was disseminated within approximately 300 milliseconds to neurons in every one of the 34 regions they recorded from! Such propagation of information triggered by the "go" stimulus was not detected in animals allowed to freely consume water. Thus, the initial water-predicting stimulus initiates a cascade of firing throughout the brain only when the animal is in the right state—thirsty.

In another breakthrough study, researchers used calcium imaging to record from more than 10,000 neurons in the visual cortex of the mouse, while facial movements were filmed in minute detail (Stringer et al. 2019). They found that information in cortical neurons reflected over a dozen features of motor information (related to facial movements, including whiskers and other facial features), in line with emerging evidence from other investigations. These results are remarkable because, according to traditional thinking, motor and visual signals are only merged later in so-called higher-order cortical areas and definitely not in the primary visual cortex. But the surprises didn't stop there. The researchers also recorded signals across other parts of the forebrain, including cortical and subcortical areas. Surprisingly, information about the animal's behavior (at least as conveyed by motor actions visible on the mouse's face) was observed nearly everywhere they recorded. In considering the benefits of such ubiquitous mixing of sensory and motor information, the investigators ventured that behaving effectively depends on the combination of sensory data, ongoing motor actions, and internal-state variables. It seems that this is happening pretty much everywhere, including in parts of the brain believed for a long time to not mix them, like the primary sensory cortex.

The examples above hint that much is to change in neuroscience in the coming decades. Still, these results come from fairly constrained experimental settings. The amygdala study used a $40 \times 40 \times 40$ centimeter plastic box; the thirst study probed mice with their heads fixed in placed; and the facial movement study employed an "air-floating ball" that allowed mice to "run." Imagine what we'll discover in the future!

How Complicated Is It Again?

Throughout the book, we've been describing a systems view that encourages reasoning about the brain in terms of large-scale distributed and entangled circuits. But when we adopt this stance, pretty early on it becomes

clear that "things are complicated." Indeed, we might be asked if we need to entangle things that much. Shouldn't we attempt simpler approaches and basic explanations first? After all, an important principle in science is *parsimony*, often discussed in reference to what's called Occam's razor after the Franciscan friar William of Ockham's dictum that *pluralitas non est ponenda sine necessitate*: "Plurality should not be posited without necessity." In other words, keep it simple, or as Einstein is often quoted, "Everything should be made as simple as possible, but no simpler."

This idea makes sense, of course. Consider a theory T that tries to explain a given set of phenomena. Now suppose that an exception to T is described—say, a new experimental observation that is inconsistent with it. While not good for proponents of T, the finding need not be the theory's death knell if it's possible to extend T so that it can handle the exception, thereby avoiding the theory from being falsified. As T breaks down further and further, it could be gradually extended to explain the additional observations with a series of, possibly, ad hoc extensions. That's clearly undesirable. At some point, the theory in question is so bloated that simpler explanations would be heavily favored in comparison.

But whereas parsimony is abundantly reasonable as a general approach, what counts as "parsimonious" isn't exactly clear. That's where the rubber hits the road. Take an example from physics. In 1978, an American astronomer, Vera Rubin, noticed that stars in the outskirts of galaxies were rotating too fast, contradicting what would be predicted by our theory of gravity.[7] It was as if the mass observed in the universe was not enough to keep the galaxies in check, triggering a vigorous search for unaccounted sources of mass. Perhaps the mass of black holes had not been tallied properly? Other heavy objects such as neutron stars? When everything known was added up, the discrepancy was—and still is—huge. Actually, about five times more mass would be needed than what we have been able to put on the scale.

To explain the puzzle, physicists postulated the concept of dark matter, a substance previously unknown and possibly made of as-yet undiscovered subatomic particles. This would account for a whopping 85 percent of the mass of the universe. We see that to solve the problem, physicists left the standard theory of gravitation (let's call it G) unmodified, but they had to postulate an entirely new type of matter (call it m). This not a small modification! In 1983, the Israeli physicist Mordehai Milgrom proposed an alternative solution, a relatively small change to G. In his "modified gravity" theory, this force works as usual except when considering such

massive systems as entire galactic systems. Without getting into the details, the change essentially involved adding a new constant, called a_0, to the standard theory of gravity.[8]

So, our old friend G doesn't work. We can either consider $\{G+m\}$ or $\{G+a_0\}$ as potential solutions. Physicists have not been kind to the latter solution and considered it rather ad hoc. In contrast, they have embraced the former and devoted monumental efforts to finding new kinds of matter that can tip the scales in the right direction. At present the mystery is unsolved, and larger and more sophisticated instruments continue to be developed in the hope of cracking the problem. The point of this brief incursion into physics was not to delve into the details of the dispute but to illustrate that parsimony is easier said than done. What is considered frugal in theoretical terms depends very much on the intellectual mindset of a community of scientists. And, human as they are, they disagree.

Another example that speaks to parsimony relates to autonomic brain functions that keep the body alive—for example, regulating food and liquid intake, respiration, heart activity, and the like. The anatomical interconnectivity of this system has posed major challenges to deciphering how particular functions are implemented. One possibility, along the lines proposed in this book, is that multiregion interactions collectively determine how autonomic processes work. But consider an alternative position. In an influential review, Clifford Saper stated that "although network properties of a system are a convenient explanation for complex responses, they tell us little about how they actually work, and the concept tends to stifle exploration for more parsimonious explanations" (Saper 2002, 460). According to him, the "highly interconnected nature of the central autonomic control system has for many years served as an impediment to assigning responsibility for specific autonomic patterns" to particular groups of neurons.

It's not a stretch to say that thinking in terms of complex systems is not entirely natural to most biologists. Many, in fact, view it with a nontrivial amount of suspicion, as if this approach "overcomplicates" things. Their training emphasizes other skills, after all. Unfortunately, if left unchecked, the drive toward simple explanations can lead researchers to adopt distorted views of biological phenomena, as when proposing that a "schizophrenia gene" explains this devastating condition or that a "social cognition brain area" allows humans, and possibly other primates, to have behavioral capabilities not seen in other animals. Fortunately, neuroscience is gradually changing to reflect a more interactionist view of the brain. From the earlier

goal of studying how regions work, current research takes to heart the challenge of deciphering how circuits work.

A key aspect of any scientific enterprise is conceptual. Scientists decide the important questions that should be studied by accepting or rejecting papers in the top journals, funding particular research projects, selecting topics for conferences, and so on. Many of these judgments are subjective, in the sense that they are not inherent to the data collected by scientists. How one studies natural phenomena is based on accepted approaches and methods of practicing researchers. Accordingly, the position to embrace or shun complex systems is a collective viewpoint. To some, network-based explanations are too unwieldy and lacking in parsimony. In diametrical contrast, explanations heavily focused on localized circuits can be deemed as oversimplistic reductionism, or purely naive.

In the end, science is driven by data, and the evidence available puts pressure on the approaches adopted. Whereas mapping the human genome took $3 billion and a decade at first, gene mapping can be done routinely now for under a thousand dollars in less than two days, enabling unprecedented views of genetics. In the case of the brain, it's now viable to record thousands of neurons simultaneously, opening a window into how large groups of neurons generate behaviors in ways that weren't possible before. (Remember that most of what we know about neurophysiology relied on recordings of individual neurons or very small sets of cells at a time.) For example, in a trailblazing study, researchers recorded single neurons across the entire brain of a small zebrafish.[9] Their goal was to record not most but *all* of the creature's neurons! Perhaps one day in the not-so-distant future, the same will be possible for larger animals.

Causation in Complex Systems Is a Whole Different Thing

Nowhere else is the challenge of embracing complex systems greater than when confronting the problem of *causation*. "What causes what" is *the* central problem in science, at the very core of the scientific enterprise.

One of the missions of neuroscience is to uncover the nature of signals in different parts of the brain and ultimately what causes them. A type of reasoning that is prevalent is what I've called the *billiard ball* model of causation (Pessoa 2017a, 2018b). In this Newtonian scheme, force applied to a ball leads to its movement on the table until it hits the target ball. The

reason the target ball moves is obvious; the first ball hits it, and through the force applied to it, it moves. Translated into neural jargon, we can rephrase it as follows: A signal external to a brain region excites neurons, which excite or inhibit neurons in a second brain region through anatomical pathways connecting them. But this way of thinking, which has been very productive in the history of science, is too impoverished when complex systems—the brain for one—are considered.

We can highlight two properties of the brain that immediately pose problems for standard, Newtonian causation.[10] First, anatomical connections are frequently bidirectional, so physiological influences go both ways, from A to B and back. If one element causally influences another while the second simultaneously causally influences the first, the basic concept breaks down. Situations like this have prompted philosophers to invoke the idea of "mutual causality." For example, consider two boards arranged in a Λ shape so that their tops are leaning against each other; so, each board is holding the other one up. Second, *convergence* of anatomical projections implies that multiple regions concurrently influence a single receiving node, making the attribution of unitary causal influences precarious.

If the two properties above already present problems, what are we to make of the extensive cortical-subcortical anatomical connectional systems and, indeed, the massive combinatorial anatomical connectivity discussed in chapter 9? If, as advanced, the brain basis of behavior involves distributed, large-scale cortical-subcortical networks, new ways of thinking about causation are called for. The upshot is that Newtonian causality provides an extremely poor candidate for explanation in *non-isolable* systems like the brain.

What are other ways of thinking about systems? To move away from individual entities (like billiard balls), we can consider the temporal evolution of "multiparticle systems," such as the motion of celestial bodies in a gravitational field. Physicists and mathematicians have studied this problem for centuries, which was central in Newtonian physics. For example, what types of trajectories do two bodies, such as the earth and the sun, exhibit? This so-called two-body problem was solved by Johann Bernoulli in 1734. But what if we are interested in three bodies—say, we add the moon to the mix? The answer will be surprising to readers who think the challenge sounds easy given that two bodies were understood long ago. On the contrary, this problem has vexed mathematicians for centuries and

in fact cannot be solved! At least not in the sense of two bodies, because it doesn't admit to a general mathematical solution.

So, what can be done? Instead of analytically solving the problem, one can employ the laws of motion based on gravity and use computer *simulations* to determine future paths.[11] If we know the position of three planets at a given time, we can try to determine their positions in the near future by applying equations that explicitly calculate all intermediate positions. In the case of the brain, where we don't have comparable equations, we can't do the same. But we can extract a useful lesson and think of the joint state of multiple parts of the brain at a given time. How does this state, which can be summarized by the activity level of brain regions, change with time?

Before describing some of these ideas further, I'll propose an additional reason that contemplating dynamics is useful. For that, we need to go back in time a little.

The World Is Made of Processes, Not Things

Starting in the mid-fifth century BCE, Greek thinkers including Leucippus, Democritus, and Epicurus thought of nature as made of immutable atoms in empty space. Although the scientific revolution of the sixteenth and seventeenth centuries shunned older classic ideas, it enshrined atomism. That is to say, first and foremost the world consists of substantial particles or *things*. Science, therefore, must seek to explain organs, cells, molecules, and so on, to give a few biological examples. This statement appears so innocuous and obvious as to appear to be a truism. What could possibly be an alternative?

Perhaps not surprisingly, we can go back to the Greeks to entertain a second view, one encapsulated in the dictum *panta rhei* ("everything flows"), or the famous saying that "no person ever steps in the same river twice." As Heraclitus suggested:[12]

> Reality is not a constellation of things at all, but one of processes. The fundamental "stuff" of the world is not material substance, but volatile flux, namely "fire", and all things are versions thereof (puros tropai). Process is fundamental: the river is not an object, but a continuing flow; the sun is not a thing, but an enduring fire. Everything is a matter of process, of activity, of change (panta rhei).

In the first half of the twentieth century, this view was embraced by the so-called organicists, biologists attempting to build a science of life that was dynamic and systems oriented. As asserted by the geneticist Conrad

Waddington, biology does not study things; it studies processes occurring at various timescales (Dupré and Nicholson 2018, 9). In this view, thinglike entities don't necessarily need to be excluded; they can be considered "processes" stable and sustained enough to have substance—think of a person, an organ, or a cell. But embracing such a process-oriented mindset—a *process philosophy*—naturally leads to new ways of formulating and answering scientific questions. And whereas this view obviously helps with phenomena such as hurricanes, streams, and vortices, it encourages describing biological phenomena in terms of context-dependent, dynamic processes. For one, the compulsion of neuroscientists to define areas and subareas recedes, giving way to the goal of deciphering how processes involving multiple parts of the brain unfold temporally to support behaviors.

Transient Brain Dynamics

If the discussion sounds intriguing, it almost certainly feels vague. Let me illustrate some ideas in use by neuroscientists to understand *multiregion dynamics*. The objective is to describe the *joint state* of a set of brain regions and how it changes.

Imagine a system of n brain regions labeled 1, 2, . . . , each with an activation (or firing rate) strength that varies as a function of time denoted $x_1(t)$, $x_2(t)$, and so on. We can group these activities into a vector x. Recall that a vector is simply an ordered set of values, such as x, y, and z in three dimensions. At time t_1, the vector $x(t_1)$ specifies the *state* of the regions (that is, their activations) at time t_1. By plotting how this vector moves as a function of time, it is possible to visualize the temporal evolution of the system. We can call the succession of states at t_1, t_2, etc., visited by the system a *trajectory*. Now, suppose an animal performs two tasks, A and B, and that we collect responses across three brain regions, at multiple time points. We can then generate a trajectory for each task (figure 12.2), each providing a potentially unique *signature* for the task in question.[13] We have created a four-dimensional representation of each task, and it considers three locations in space (the regions where the signals were recorded from) and one dimension of time. Of course, we can record from more than three places; that only depends on what our measuring technique allows us to do. If we record from n spatial locations, then we'll be dealing with an $(n+1)$-dimensional situation (the +1 comes from adding the dimension of time). Whereas we can't plot

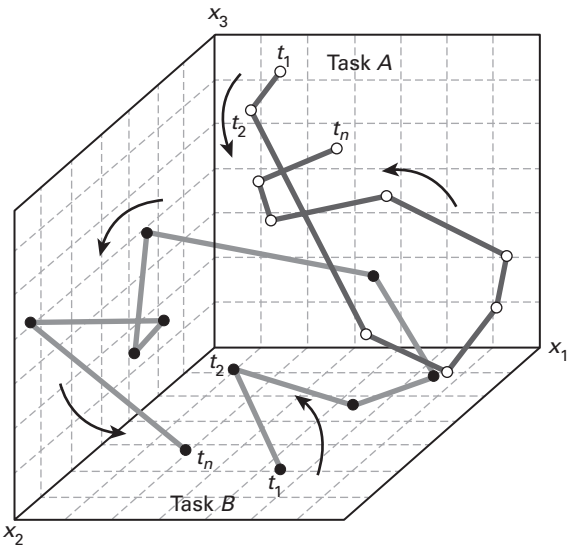

Figure 12.2
Spatiotemporal trajectories. The activity of three brain areas is shown across time. The three values at time t form a vector. In three dimensions, we can plot vectors (see figure 4.3b for examples) at each time point and join the endpoints. In this representation, time is only implicit and runs along the lines starting at t_1 and ending at t_n. The activities across the three regions as a function of time evolve distinctly for task A and task B. This particular evolution, the trajectory, provides a signature for the task in question.

the trajectory on a piece of paper, fortunately the mathematics is the same, so it poses no problems for data analysis.

When we think in terms of spatiotemporal trajectories, the object of interest—the trajectory—is spatially distributed and, of course, dynamic. It also encourages a process-oriented mindset instead of trying to figure out how a brain region responds to a brief stimulus. The process view also changes the typical focus on "billiard ball" causation—the white ball hits the black ball, or region A excites region B—as we are less obsessed about what single factor is responsible for a region's response. Experimentally, a central goal, then, becomes estimating trajectories robustly from available data.

Some readers may feel that, yes, trajectories are fine, but aren't we merely describing the system but not *explaining* it? Why is the trajectory of task A

different from that of task B, for example? Without a doubt, a trajectory is not the be-all and end-all of the story. Deciphering how it comes about is ultimately the goal, which will require more elaborate explanations, and here computational models of brain function will be key. In other words, what kind of system, and what kind of interactions among system elements generate similar trajectories, given similar inputs and conditions?

Final Thoughts

Neuroscience strives to elucidate the neural underpinnings of behaviors. Modern neuroscience has done so in a preponderantly reductionistic fashion for over a century and a half.[14] I would venture that progress has been stymied by such approach and that the time is ripe for the field to phase-transition into a period when a truly dynamic and networked view of the brain takes hold. Future research will need to strive to make progress along several fronts: dynamics, decentralized computation, emergence, and competition. At the same time, a science of the mind-brain must be developed by erecting it from a solid foundation of understanding behavior while employing computational and mathematical tools in an integral manner.

I believe the field of neuroscience needs to take stock and invest on the development of conceptual and theoretical sides. Bigger and shinier tools and techniques alone won't yield the necessary progress; we run the risk of being able to measure every cell (or subcellular component even) in the brain in a theoretical vacuum. To drive the point home, suppose that experimental physicists could measure every atom of a given galaxy. How would that advance understanding if not for a theory of gravitation that took more than 400 years of development? The current obsession in the field with causation is equally problematic. Without conceptual clarity (how should we even think of causation in highly entangled systems?), "causal" explanations in fact might miss the point.

Ultimately, to explain the cognitive-emotional brain, we need to dissolve boundaries within the brain[15]—perception, cognition, action, emotion, motivation—as well as outside the brain, as we bring down the walls between biology, psychology, ecology, mathematics, computer science, philosophy, and so on. Only then we will be on the right track.

The Entangled Brain

How Perception, Cognition, and Emotion Are Woven Together

By: Luiz Pessoa

Citation:

The Entangled Brain: How Perception, Cognition, and Emotion Are Woven Together
By: Luiz Pessoa
DOI: 10.7551/mitpress/14636.001.0001
ISBN (electronic): 9780262372107
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

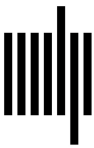

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Pessoa, Luiz, author.
Title: The entangled brain : how perception, cognition, and emotion are woven together / Luiz Pessoa.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Includes bibliographical references and index.
Identifiers: LCCN 2021061878 (print) | LCCN 2021061879 (ebook) | ISBN 9780262544603 (paperback) | ISBN 9780262372107 (pdf) | ISBN 9780262372114 (epub)
Subjects: LCSH: Perception. | Emotions and cognition. | Brain. | Neuropsychology.
Classification: LCC BF311 .P3767 2022 (print) | LCC BF311 (ebook) | DDC 153—dc23/eng/20220411
LC record available at https://lccn.loc.gov/2021061878
LC ebook record available at https://lccn.loc.gov/2021061879