Tipping Points in the Brain

If enough neurons fire in synchrony at one site, their message can spread far and wide without distortion

BY TARA THIAGARAJAN

KEY CONCEPTS

- The physical manifestation of a memory or percept is believed to be transient electrical activity among neurons. But identifying relevant electrical patterns from the chatter of the brain's billions of neurons has been elusive.
- When many neurons in a local area are saying the same thing, their message can spread to several regions of the cortex in a one-to-many cascade. These cascades are called coherence potentials.
- Coherence potentials may be able to reconcile the contradictions of distributed and localized function.

In his quest to unravel the relationship between brain function and intelligent behaviour, Karl Lashley, an American psychologist, conducted a set of experiments in the 1920s and 30s. The psychologist meticulously destroyed different parts of the cerebral cortex of rats, allowed them to recover from the surgery and then tested their ability to learn and remember certain tasks such as running a maze or distinguishing between two patterns. The degree of impairment depended only on how much of the cortex was destroyed—and not which part. It was not until he had removed at least half the cortex that learning and memory became severely impaired. Inspired by these findings, one of his students, Donald Hebb, went on to study the impairment resulting from brain damage and surgery in humans. Again, the results were astonishing. Clear-cut removal of parts of the cortex outside the speech area often had little, if any, detectable effect. A man who had a prefrontal lobe removed continued to score an extraordinary 160 or higher score on IQ tests. A woman who had lost the entire right half of the cortex continued to have an IQ of 115, better than two-thirds of the normal population. These were the striking cases, Hebb conceded, but it begged the question: how was it possible that while large brain injury often had severe impact on intelligence, sometimes it did not. Once a concept had been learned, Hebb noted, it was not easily lost by brain damage.

In light of these findings Lashley concluded that learning and memory were distributed across the cortex and not localized in any one place—that every part of the cortex was equal in its capability, or *equipotential*.

Yet, during these same decades, a rather different set of studies was taking shape, carving up the cortex into localized 'functional' areas. As early as 1907, German neurologist Korbinian Brodmann had divided up the

IMAGE AND ILLUSTRATIONS BY ALANKAR

> The discovery of coherence potentials, for the first time, provides an easy way to pick out groups of neurons engaged in a similar conversation.

cortex into 'functional areas' based on anatomical connectivity between the sensory apparatus and the cortex. Brodmann went as far as to propose that these areas of the cortex represented different functional 'organs'. By the 1940s the localization theory had gained further support. Largely consistent with the anatomical maps of Brodmann, researchers had constructed physiological maps. Application of pressure to different parts of the body elicited electrical potentials on distinct areas of the cortex. Every part of the body mapped to a distinct region of the cortex. Over the ensuing decades, the theory of localization found its way in the greater psyche of the scientific community. Neural pathways from the retina arrived in a particular region of the cortex which responded robustly with electrical potentials to visual cues. Neural pathways from the ear arrived at a particular region of the cortex which responded to sound. Conversely, electrical stimulations of those regions of the cortex resulted in corresponding movement or sensation. Such functional maps showing a visual cortex, auditory cortex, motor cortex and so on, are now printed in every basic neuroscience textbook.

How is it possible that sensorimotor function maps so clearly on to particular regions of the cortex while higher function such as learning and memory and intelligence that make use of these functions are distributed? How do the various modes of sensory input come together to create the integrated perceptual stream we experience? Almost a century after Lashley, these quandaries remain unresolved.

One of the most fundamental assumptions of neuroscience research is that the physical substrate of all thought and behaviour is the electrical activity transmitted among neurons or neural cells. In these terms, the question is whether *any* set of neurons can create the behavioural or perceptual response to a task at hand, or whether it must it be a *specific* set? Thus for any sort of understanding it is necessary to probe more deeply the elements of the brain, the neurons that make up the cortex, connecting with one another to form a complex circuitry of electrical communication. What do they each know? Who talks to whom? How do they share information? For many decades neuroscientists have painstakingly added pieces to this puzzle.

Individual neurons have now been probed in numerous ways. In a pioneering study in 1959 Torsten Wiesel and David Hubel found that neurons can be remarkably specific in what they care about. Measuring electrical activity from individual brain cells in the visual cortex of an anesthetised cat, they found that some fired vigorously when a bar was moved in front of the retina in one direction but not any of the others. In the ensuing decades a host of similar studies followed that explored neurons in various parts of the cortex. Some responded to specific colours but not others. Some responded to certain sound frequencies but not others. But not all neurons were so narrow in their world view. In 2005 Rodrigo Quiroga and colleagues found neurons that responded vociferously to pictures of Jennifer Aniston, but not to pictures of Brad Pitt or pictures of toilet brushes, suggesting that neurons are capable of larger concepts, recognizing complete images and not just select visual features. Still, the collective results seem to suggest that individual neurons hold highly specialized or selective knowledge, although of varying complexity. If knowledge is specialized among neurons and not generalized, how can that be reconciled with Lashley's findings on cortical damage? Surely damage of these locally special-

[THE BINDING PROBLEM]

HOW DOES IT ALL COME TOGETHER?

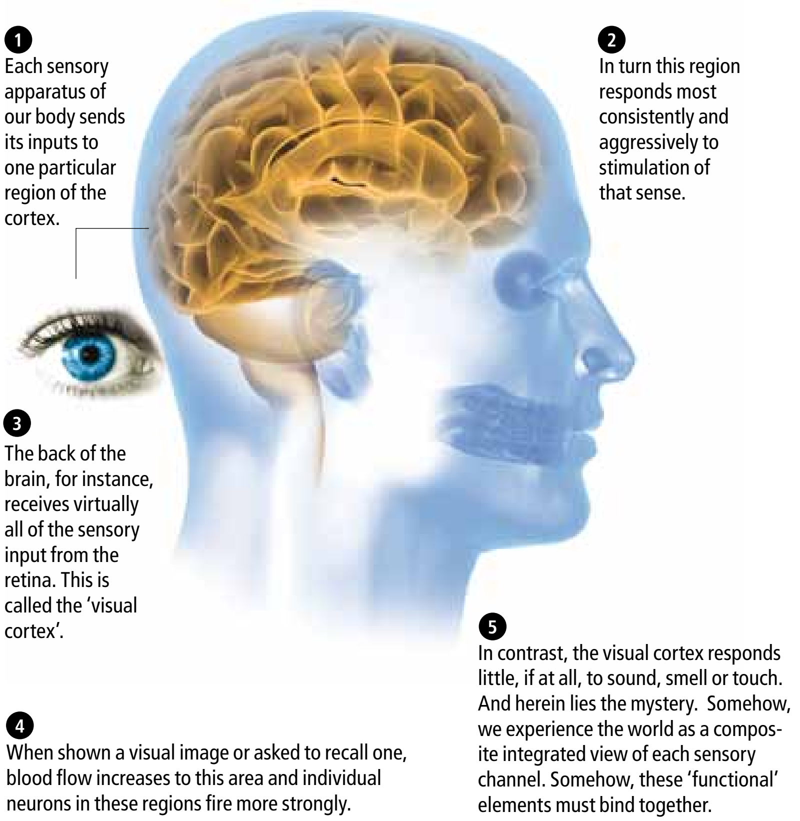

ized cells should completely destroy any abilities that draw from their knowledge. But extrapolating from Lashley's work, damage to those Jennifer Aniston cells in the visual cortex is not likely to get her off your mind.

This brings us back to the question of how neurons might share specialized knowledge to create an integrated perceptual whole. Hebb himself put forward one of the most influential hypotheses in this regard. Drawing from the findings of Lorente de No, a French physiologist who had identified closed loops of connected neurons that sometimes spanned multiple 'functional' areas of the cortex, he postulated that an individual experience, thought or memory manifests as electrical activity transiently reverberating among a subset of neurons in a closed loop, a 'cell assembly'. This is important because it suggests that what counts is not just what neurons hear and get excited about in the external world but how they share what they hear with one another, who shares with whom, and what memory this sharing creates. It frames the problem in the context of identifying structures of electrical activity among neurons that can persist in the network on the timescale of perception, typically a tenth of a second, without dying out. With this framework in mind, many sorts of simulated models have been proposed, some that are even significant departures from Hebb's own reverberating loops. However, empirical evidence of such 'cell assemblies' or persistent electrical activity among neurons has been elusive.

GLIMPSE OF AN IDEA

TECHNOLOGY HAS OFFERED US many ways of viewing the activity of the brain with increasingly better resolution, from measuring changes in the blood flow to different regions, to placing electrodes in the brain to listen in on the electrical chatter of surrounding neurons, and even to individual neurons in isolation. However, the intractable problem has been identifying which regions of the brain or which neurons are communicating with each other at any given time. There are billions of neurons in the brain and none are ever quiet for very long. Moreover, each neuron is connected to many others in a dense interconnected web, so it is difficult, if not impossible, to know who is talking to whom at any given time. Imagine if you could listen in on a whole bunch of people across the world by putting a microphone close to their mouths. Many would be talking at once. Sometimes they might be talking to their friends in their neighbourhood. Other times they may be talking on their phones to people far away. All together it would sound like an unintelligible cacophony. It is only if you could isolate individual conversations that you could begin to make sense of what they are saying and know which conversations mattered. Yet, isolating conversations among neurons has been tricky or near impossible.

This is the problem I set out to solve along with Dietmar Plenz at the National Institutes of Health. It could have turned out to be a blind alley but there were fortuitous clues that provided a starting point. For instance, numerous studies over the preceding few years had showed that during certain behaviours particular pairs of neurons are more likely to synchronize their electrical activity. Synchrony was important, although not perfectly consistent across trials of the same behaviour. Collective behaviour seemed to matter in other ways too. An electrode placed in any field of cortex measures the sum of the electrical activity of many neurons at once in a surrounding field. This summation of electrical

[NEURON ACTIVITY]

MAKING SENSE OF THE CHATTER

The most influential theory in neuroscience is that the physical manifestation of a composite perceptual experience or memory is electrical activity 'reverberating' for a brief period among a subset of neurons, a multiway conversation. Today, with technological advance we can listen in on the electrical chatter of hundreds of neurons at a time. Imagine being able to listen to hundreds of people all over the world talking but not knowing who is talking to whom. Unless you can listen to a single conversation in isolation, all the chatter would sound like one nonsensical cacophony. How could you make sense of it?

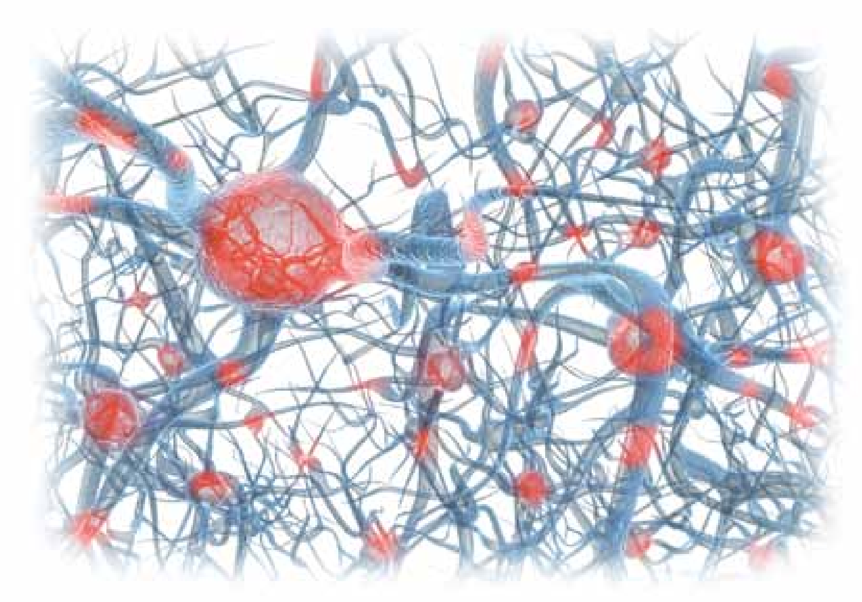

[THE AUTHOR]

Tara Thiagarajan is a Visiting Scientist at the National Centre for Biological Sciences in Bangalore, India, where she works on the network dynamics of the brain. She is also Chairperson of Madura Microfinance, an experiment beyond microfinance in reengineering trade and social networks of the rural poor. She completed her Ph.D. in Neuroscience at Stanford University and postdoctoral work at the National Institutes of Health in Bethesda, MD.

[COHERENCE POTENTIALS]

LISTENING IN ON TRANSIENT CONVERSATIONS

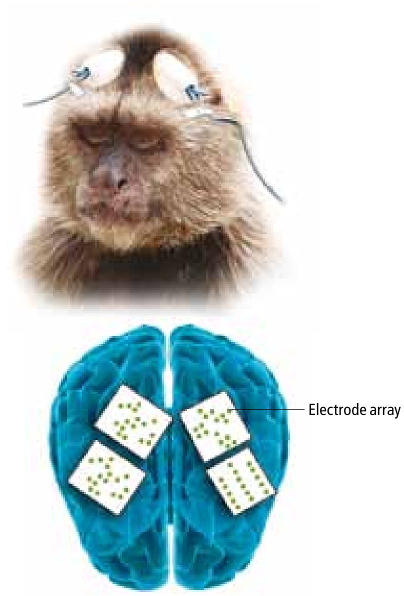

A

Electrodes placed on the surface of the brain measure the composite electrical activity of many neurons in their surround or local field. This measurement is called a local field potential or LFP. The more neurons in the field that are active at any given time, and the more the temporal pattern of activity of these neurons is similar or synchronized, the more it adds up. At each site, there are brief periods when there is greater synchronization of activity patterns, like a crowd pausing their various conversations to sing in unison.

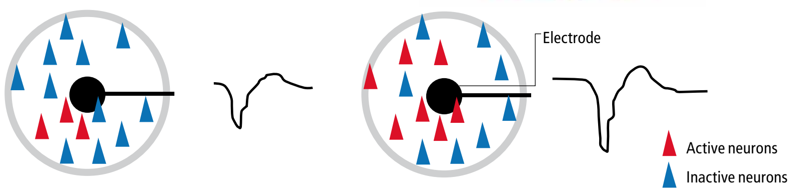

B

When enough neurons synchronize their patterns at one site, that pattern of activity can rapidly spread to multiple sites without any distortion. This means, if many neurons at one site are 'saying' the same thing at virtually the same time, their message can travel far and wide, rallying almost as many neurons at many other sites. We call these periods of coherence across multiple sites 'coherence potentials'. The patterns of these coherence potentials are highly diverse and can be quite complex. This means they could possibly encode a great deal of information.

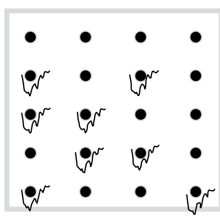

C

Over the course of our measurements, the electrical activity showed a stream of such associations, occasionally overlapping in time, each substantially different from the one before, and therefore, easily discernible. For the first time, this allows us to identify distinct transient associations in the cortex, i.e. groups of neurons engaged in a similar conversation. We can now begin to ask the important questions of what these transient conversations of the brain could mean and how different sites of the cortex come together to make us tick.

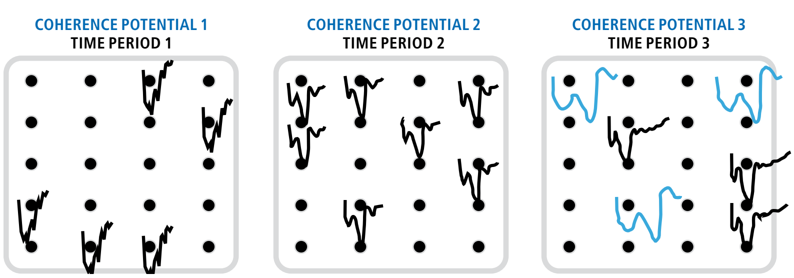

activity is called the 'local field potential' or LFP. It had been shown that the LFP at a particular location could carry more consistent information about a behaviour in its temporal structure than any individual neuron's activity. Temporal structure was important, but not well understood. But how were neuronal synchrony and collective temporal structure related to one another?

The first task was getting a firm understanding of the LFP. Because neurons signal by the movement of charge in and out across their membrane, this creates fluctuations that can easily cancel each other out in the aggregate. Complicating matters is that this measure is a composite not just of what the neurons 'say' but what they 'hear' as well. Because of these many ambiguities, the LFP signal was long considered uninformative. But what if for brief moments, many neurons were saying something similar at the same time? Like a crowd pausing their various conversations to sing in unison, the aggregate sound going from discordant chatter to coherence. Indeed, when we compared aggregate activity in the local fields around the electrode to the firing patterns of a small number of individual neurons in the same field, we found that the more the neurons fired in unison, the larger the amplitude or the electrical summation in the local field. And going further, we found that the more similar the temporal structure or pattern of the LFP waveform in different fields, the more similar the patterns of electrical activity of neurons in these fields. This meant that the amplitude of the signal was a proxy for synchronization of activity in a local region and temporal structure reflected the firing patterns of the neurons in that field. Typically, these periods of local synchronization lasted for a few hundred milliseconds, a time scale that matches with the temporal resolution of perception.

So what happens, we asked, when there is greater synchronization? Remarkably, we found that there appeared to be a 'tipping point', where when the signal was large enough at any one site, indicating more neurons 'talking' in unison, the pattern of activity they produced had a very high likelihood of spreading rapidly to other sites without any distortion. That means, if the crowd was speaking in unison, the message was much more likely to travel far and wide completely intact, instead of dissipating into smaller conversations and dying out. The way in which the messages travelled appeared to be in a one-to-many cascade often jumping across large stretches of cortex. We called these cascades coherence potentials, because they indicated brief periods of 'coherence' in the signal across multiple locations. Coherence potentials could initiate anywhere and spread over variable locations, though some locations certainly initiated and participated more frequently. Remarkably, physical distance didn't seem to matter, almost suggesting that there might be some mechanism of long range one-to-many broadcast at play. Of particular interest is that each coherence potential cascade was not the same as the one before. The diversity of the temporal patterns of coherence potentials spanned a wide range of frequencies indicating many degrees of freedom and therefore the ability to encode complex information. Simply put, these messages had few constraints in their structure and therefore could say quite a lot. The discovery of coherence potentials, for the first time, provides an easy way to pick out groups of neurons engaged in a similar conversation or in brief vocal agreement. And they are present in the cortical activity of rats and monkeys and, as we are now seeing, humans as well.

Over the course of our measurements, the electrical activity showed a stream of diverse coherence potentials, typically one after another in time. Importantly, each was substantially different from the one before in its temporal structure, and therefore easily discernible, providing a clear visualization of the cortical network shifting from one pattern of association to another. Indeed, the idea that our perceptual stream arises from the cortex shifting abruptly from one transient pattern of electrical activity to another in time has been proposed time and again by researchers across the globe.

But are these fast spreading messages the important conversations or simply neurons cheering at their version of a cricket match or soccer game? Do they have a role in driving behaviour? Studies by others indicating that the LFP in any particular location carries information about behaviour in its temporal structure certainly seem to suggest that they do. And at my lab at the National Centre for Biological Sciences in Bangalore, India, we are now finding more resounding and specific proof that they are indeed relevant to behaviour.

Most significantly, the discovery of coherence potentials provides ground for a new framework of thinking that could reconcile functional localization in the cortex with Lashley's early findings, that no specific function or memory could be ablated by removing any particular part of the cortex. The apparent tipping point in the spread of the message is analogous to the spread of ideas and innovations in society. Perhaps, just as ideas and innovations tend to originate or form in society at places of specialization—technology innovations in Silicon Valley, movies in Hollywood etc.—ideas or concepts are developed and originate in places of input specialization in the cortex. If they are good enough or get enough local support they then spread rapidly to other regions where they are talked about and adopted as part of common 'culture'. Once a concept has become part of common culture, it would be sufficiently distributed across the cortex such that damage to any particular part would not impact the ability to make use of that concept.

In my lab we are beginning to test this hypothesis using recordings from humans performing simple tasks. What seems to be unfolding is an organization of coherence potentials that correlates with behaviour in an unexpected way that parallels organizational structure in society. If you think about it, society arises from the interactions of brains and is therefore a meta network of neurons or network of neural networks. Designing experiments and analysis that draw from our understanding of the structure and function of society may therefore give us remarkable new insights into how the brain functions. ■

[A META NETWORK]

DOES THE BRAIN'S WORKING PARALLEL SOCIETY?

Society is a network of brains at work, a meta neural network. Perhaps, we can imagine, just as ideas tend to originate in society at places of specialization, technology innovations in Silicon Valley, movies in Hollywood etc., ideas are developed and originate in places of specialization in the cortex. If they are good enough or get enough local support then they spread rapidly to other regions where they become adopted. If enough of the cortex adopts an idea or innovation or agrees on something, it would become common culture, explaining why, while input is localized, function is not.

MORE TO EXPLORE

Coherence Potentials: Loss-less, all-or-none network events in the cortex, Thiagarajan et al, PLoS Biology, January 2010

The organizing principles of neuronal avalanches: Cell assemblies in the cortex? Dietmar Plenz and Tara Thiagarajan, Trends in Neuroscience March 2007

The Organization of Behavior, Donald O. Hebb, 1949 (A historical view of the debate)