© NewScientist

Feature

27 September 2017

The brain's 7D sandcastles could be the key to consciousness

We've glimpsed mind-bending geometric structures that fleetingly encode our thoughts, memories and feelings – and could solve the greatest mystery of all

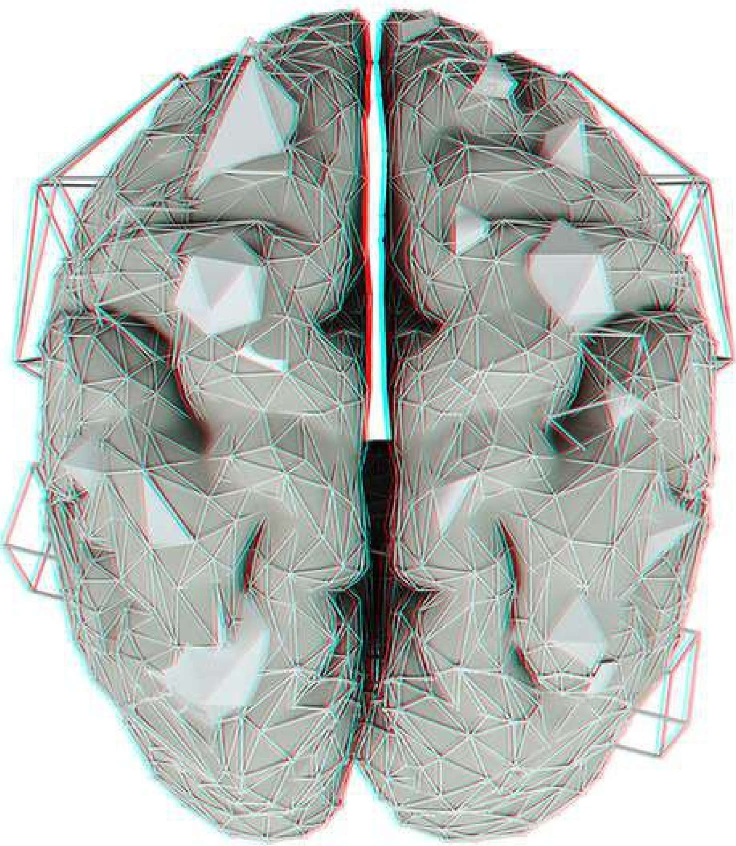

Steven Wilson Studio

By Anil Ananthaswamy

EDWIN ABBOTT, in his 1884 book *Flatland*, created a fictional 2D landscape full of lines, triangles, squares and circles that have no notion of up or down. One day, a 3D Sphere visits Flatland and whisks away a Square to a higher dimensional world. Square learns that Flatlanders are mere 2D projections of 3D beings. He then has the audacity to suggest that Sphere may be a shadow too – of a shape in four dimensions. "The very idea of it is utterly inconceivable," says the appalled Sphere.

Henry Markram thinks we might be suffering from a similarly blinkered perspective when considering the workings of our own brains. "We look at the brain, we see its immense complexity, but if it's a shadow projection from a higher dimension, we'll never understand it," Markram says. Those aren't idle words: he and his colleagues of the Blue Brain Project at the Swiss Federal Institute of Technology in Lausanne (EPFL) have been using algebraic topology, a field of mathematics used to characterise higher-dimensional shapes, to explore the workings of the brain.

What they have found beggars belief. As our brains think, learn and remember, they create elaborate but ephemeral structures in at least seven mathematical dimensions, and possibly many more. What's more, these transient structures, which appear and disappear like sandcastles on a beach, could help us understand how the brain creates our thoughts and feelings. They might even unravel the greatest mystery of them all: consciousness. "Algebraic topology is the mathematics to take neuroscience out of Flatland," says Markram.

The Blue Brain Project was launched in 2005, with the aim of simulating the entire human brain inside a computer. That's an ambitious goal and far from fruition. In late 2015, however, the team announced it had recreated a sliver of the rat brain that is involved in sensing touch. The real brain tissue is only 0.5 millimetres wide and 2 millimetres long, but its digital analogue consists of 31,000 neurons of more than 200 different types, with some 8 million connections between them (see "How to build a brain").

"How does a gigantic mass of identical cells produce such beautiful complexity?"

This is the most detailed digital reconstruction of part of a brain ever created. Not everyone thinks it's possible to understand a biologically complex organ like the brain by simply recreating it inside a computer, but for Markram, the project director, such simulations let you see how neurons work together at a level of detail unobtainable with an actual slice of brain tissue, let alone the whole brain. But he admits there's a problem: making sense of the data the simulations provide. That's where algebraic topology comes in.

Topologists study shapes as they undergo continuous deformations – things like pushing, pulling and stretching, but not breaking and reattaching. It's not always obvious if two shapes are similar. Push your finger into a ring doughnut made of clay and create an indentation, for example, and you can slowly deform this doughnut into a coffee cup. The indentation becomes the inside of the cup and the doughnut's central hole becomes the handle. The key is that both shapes have just one hole – the topology's invariant property. "People call topology rubber geometry," says Kathryn Hess, an algebraic topologist who also works on the Blue Brain Project. "Things can be deformed as if they are made of rubber or silly putty." The algebraic part refers to the use of algebra to represent and manipulate the properties of such objects.

Markram's fascination with the subject began in 1994, when he was neuroscientist at the University of Heidelberg in Germany. There he met algebraic topologist Ran Levi, and the two began discussing how this branch of mathematics might be used to understand the brain. Levi introduced Markram to Hess, and the three have spent years speculating about the topological shapes that might form in a working network of neurons, and what these shapes might have to do with brain function. “The algebraic topologists are very pure mathematicians, they live in these high-dimensional spaces, and they don’t really care about the realities of life,” says Markram. “So we had very, very abstract discussions.” The Blue Brain simulation provided an opportunity to test those abstractions on real data.

They were looking in particular for the appearance of structures called cliques. A network of neurons can be depicted as a graph, the mathematical name for a diagram like the map of the London underground. The neurons are like the stations on the map and the lines represent the connections between them. A clique is a dense type of graph in which every neuron is connected to every other neuron. They correspond to geometrical shapes: three neurons in a clique form a 2D triangle; four will form a 3D shape, a pyramid with triangular faces known as a tetrahedron. But if the cliques have more than four neurons, the geometric structures they represent exist in mathematical dimensions higher than we can visualise – four dimensions for five neurons, and so on (see “The multidimensional shapes of thought”).

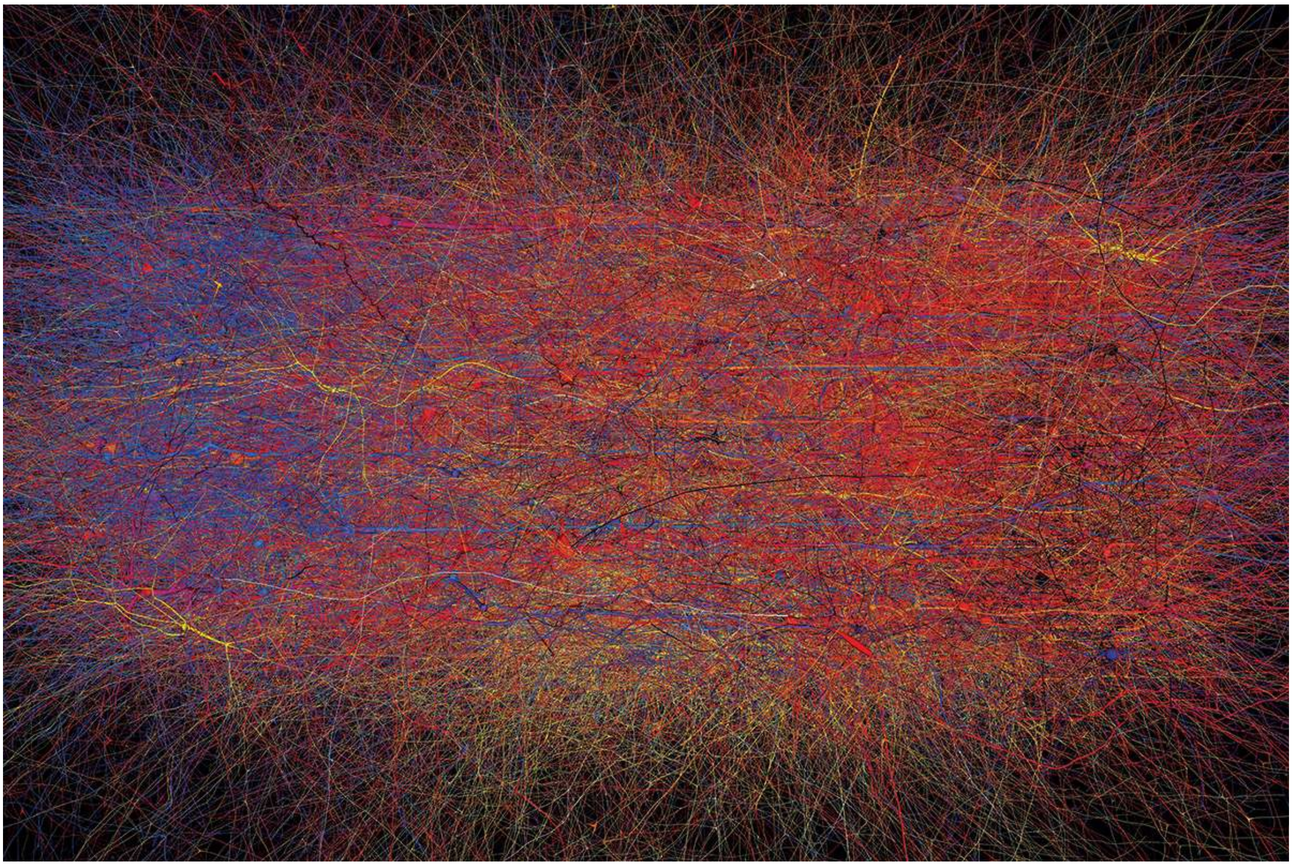

Neurons firing in the brain create tangled webs of connections
EPFL/Blue Brain Project

Other researchers had seen such cliques in real brains. For example, Chad Giusti at the University of Delaware in Newark and his colleagues found them when looking at the electrical activity of neurons in the hippocampus as a rat ran around its environment. But they were unable to discern the direction of information flow from one neuron to another within these cliques, which is crucial to understanding how they work.

This is a general problem when working with a real, functioning brain. "Directionality of information flow is very difficult to ascertain," says neuroscientist Olaf Sporns of Indiana University in Bloomington, who coined the term "connectome" for the brain's connectivity diagram. But it's not a problem when you're working with a digital brain.

Hess, Levi and their colleagues looked for "directed" cliques in the Blue Brain data, in which information enters via one neuron, passes through each of the other neurons and then exits via the last. So, for example, in a clique of three neurons, A, B and C, the information must flow from A to B to C, even though they are all connected to each other. You can tell whether this is the case by looking at the synapses connecting each pair of neurons, because information flows only one way across them.

The team was in for a surprise. The biologically inspired network had many times more directed cliques than a randomly constructed network would. "And, there were more of the higher dimensional ones," says Hess. They found directed cliques with up to eight all-to-all connected neurons, forming 7D cliques – a number Hess thinks will increase as the Blue Brain simulation grows in size. "I expect we'll find cliques with up to 15 neurons or 20 neurons," she says. But the complexity doesn't end there. The team saw that cliques come together into structures called cavities. For example, several 4D cliques can bound the surface of a 3D cavity. "This doesn't happen by chance," says Hess.

So far, so abstract. What do these structures have to do with brain function? Well, in a real brain, neurons that fire together wire together: the more two neurons work together, the stronger their connection becomes. And when the researchers let their simulated brain buzz with spontaneous activity, they found that pairs of neurons connected as part of a directed clique were more likely to fire together than pairs simply connected, but not part of a clique. What's more, the bigger the clique a pair of neurons belonged to, the more likely they were to fire together. "This was already an 'aha!' " says Hess. "Being connected is not enough. You have to be connected and be a part of a bigger structure. That was the first indication that we were on the track of something interesting."

The clincher would be to see how the digital brain would respond to the sort of stimuli that occur in a real brain. To find out, the Blue Brain team first recorded various neural signals that reach a real rat's somatosensory cortex – the part of the brain that processes touch – when its whiskers are tickled. Then they fed nine different sets of such signals to the digital simulation to see what would happen. They found that simple 1D and 2D cliques formed first, and then quickly grew into higher-dimensional cliques, sometimes reaching all the way to 7D. The stronger the tickling stimulus was and the more synchronised the input received by the neurons, the more dimensions the cliques formed. Once the peak was reached, the structures collapsed. "There's a culminating point, and poof, everything collapses," says Hess. Typically, the process would last a few tens of milliseconds.

The topological perspective shows how individual neurons work together to process information. “It’s only when you put on these glasses that suddenly you see this incredible sandcastle, a multidimensional structure,” says Markram. Neuroscientists have for decades been looking at the electrical activity in different neural networks and wondering what they all have in common. The cliques and cavities could be it. “When anything happens, the brain builds the most complex structure that it can. It climbs as high as it possibly can go, and then it collapses. All stimuli evoke the same stereotypical, multidimensional sandcastle building and collapsing,” says Markram.

Worm map

But could all this simply be an artefact of the digital model? To check this, the team applied algebraic topology to a real nervous system – that of the nematode worm *Caenorhabditis elegans*. The worm has only 302 neurons, and their connectivity has been completely mapped, allowing the team to look for directed cliques. What they found confirmed their simulation. “It’s far, far more complex than randomly connecting those few hundred neurons,” says Markram. “Even a worm has multidimensional structures, allowing those very few neurons to do incredibly sophisticated tasks. That’s why we think this is a universal principle of neuronal organisation.” If animals as diverse as rats and worms exhibit complex multidimensional cliques, then “it’s pretty likely that this is a highly general phenomenon across brains”, says Markram.

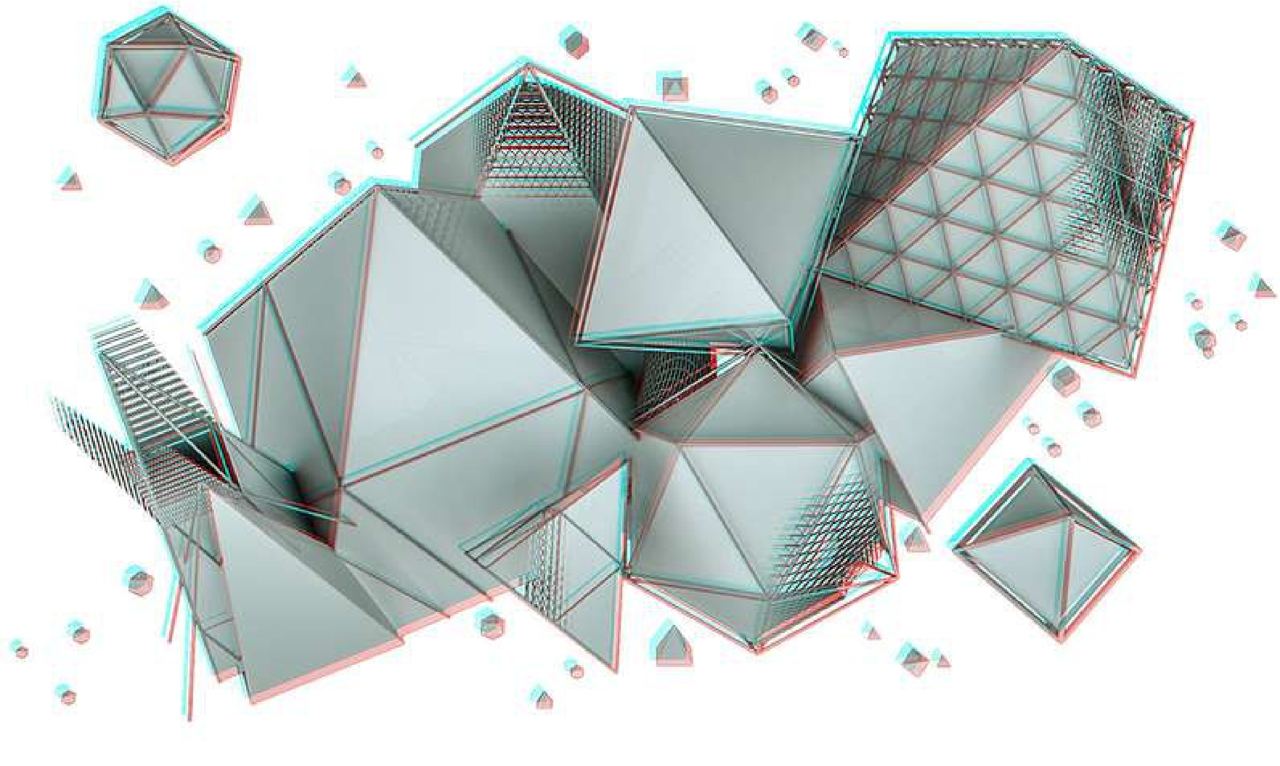

Steven Wilson Studio

If they are right, this study is a big deal, providing a way of analysing the transient connections that determine what an active brain is doing. So what do others think of it? Sporns says he is impressed that the research considers the direction of flow of information within the brain, which has been missing from connectome studies.

Karl Friston, a computational neuroscientist at University College London, agrees, but he also sees a problem with the approach. Trying to explain brain function by understanding its structure is circular reasoning, he says. “This overlooks the small fact that neural network structure emerges from function.” In other words, the cliques and other networks that form are determined by how the neurons have previously fired, and so become wired.

Nevertheless, Giusti thinks structures unearthed using algebraic topology will lead to a greater understanding of function – although it is early days. “The mathematics involved is technical enough that it’s not widely known,” he says, and the mathematical tools are still being developed. But they can potentially do amazing things, he says. For example, they could allow us to compare different people’s brains and different cognitive states. “I think we are at the beginning of a very exciting story,” says algebraic topologist Jacek Brodzki of the University of Southampton, UK.

Already, topological analysis is helping to solve some long-standing puzzles. For example, it is thought that the brain’s power comes from its “neural plasticity”, its ability to rewire itself as needed. This is a crucial ingredient for learning and forming memories. In theory, a brain is most plastic when there is a 50 per cent chance that one neuron will connect to another in its proximity. Yet in biological brains there is only about a 1 per cent chance that such connections occur, says Markram.

“Consciousness may itself be a shadow of a higher-dimensional structure”

On the face of it this makes no sense, but the topological structures provide a rationale: higher-dimensional cliques and cavities form only when the brain is sparsely connected. If these structures are a reflection of the brain’s ability to process information, then having a lower chance of making connections is better, not worse. “To form complex structures, you have to lose connections,” says Markram. “You have to try to find the lower bound of connections, which is completely radical thinking in neuroscience.”

Another puzzle that the topological lens addresses is how the brain, which looks so homogenous, nevertheless functions as though it were compartmentalised. “You see this tension: on the one hand, you have this gigantic mass of identical cells; and on the other hand, this beautifully complex array of ability of the various regions of the brain,” says Brodzki. Perhaps the cliques and cavities are the missing, emergent structures that influence function. “It’s a great result,” he says.

There are implications for artificial intelligence, too. Richard Granger, head of the brain engineering lab at Dartmouth College in New Hampshire, thinks the Blue

Brain Project is addressing a crucial gap in our knowledge about how the brain works. We know the anatomy and physiology at the level of single neurons and at the level of millions of neurons. But what if the intermediate scale is what matters when it comes to information processing? If that's the case, digitally simulating the brain and trying to find these mid-scale structures could help reveal the brain's powerful algorithms, which in turn could lead to powerful artificial intelligence.

"These are exciting and potentially groundbreaking studies," says Granger. "The scientific aim of understanding our brains and the engineering aim of duplicating them rely on our cracking the codes that make brains the best thinking machines we know of."

For Markram, the next step is to tie the ephemeral structures his team has discovered to learning and memory formation. For decades, neuroscientists have been looking at how synapses change when brains learn or store information, but they still have little idea what such changes mean. Maybe we have been doing Flatland mathematics all along. "If the changes that occur in the brain only make sense if you map them to a higher dimensional structure, then that's what you are going to have to do," he says. "Memory may be hiding in high-dimensional structures."

As the Blue Brain team continues its effort to create a larger and more accurate digital brain, Markram thinks that one day the topological approach could even help crack that hardest problem of all – consciousness. "When we see a phenomenon that looks mysterious and difficult and intractable, there is a scientific possibility that what we are seeing and experiencing is a shadow projection from higher-dimensional representations," he says. "We need mathematics to climb up into those higher dimensions. Then we'll understand how those shadows emerge. Consciousness may be a shadow."

How to Build a Brain

The goal is to recreate a human brain in a computer. There's still a long way to go, but the Blue Brain Project at the Swiss Federal Institute of Technology in Lausanne has made a start.

In 2015, the team published a digital simulation of a tiny slice of a rat's brain – the somatosensory cortex, which processes touch. Even this took years of painstaking work. More than 20,000 experiments on rat brains were used to meticulously model the shape of neurons, together with their properties such as electrical signalling and molecular mechanisms. Then, using anatomical details from five rat brains – factors such as the thickness of layers and the density of neurons in each – the neurons were assembled into a detailed digital model.

The next challenge was to figure out how these neurons would be connected. "No amount of experiments, even in the next 100 years, is going to give you all the data on all the connections that are inside a piece of brain the size of a pinhead," says Henry Markram, director of the Blue Brain project. Instead, the team had to rely on biological principles. For example, neurons must be within 3 millimetres of each other to connect.

But if all the neurons within shouting distance got interlinked, the network would be far more densely connected than it actually is in the brain. So the team applied algorithms to prune connections, to get the level of connectivity seen in real neural tissue.

Finally, they tested their simulation to see whether it responded to sensory inputs in the same way as the real thing. "The digital piece of tissue behaved very similarly to what we see in the brain," says Markram. "We see the same patterns of firing, with the same delay."

This article appeared in print under the headline "Throwing shapes"

Article amended on 2 October 2017

We have corrected Ran Levi's name

Anil Ananthaswamy is a consultant for *New Scientist*

Magazine issue 3145, published 30 September 2017