Photo courtesy of Santa Clara University

Neurosciences 2018-2019

Georges M. Halpern, MD, DSC

With Yves P. Huin

The end of 2018 and the beginning of year 2019 have brought a flurry of exciting news and discoveries in the complex field of ***Neurosciences*** and related areas.

The first entry reports on ***buried genetic networks linked to mental illness(es)***. This opens a huge area of research, and –hopefully- an avenue for targeted, efficacious therapeutics.

The next one reports on the major role(s) our ***cerebellum*** (the *little* brain) play(s). Stay tuned for more!

And we cannot ignore –don't say **forget!** - ***Alzheimer's Disease***. We may well be closing in on the triggering mechanism(s) of this modern scourge. That is certainly not the *end of the story.*

Then it's ***cognition*** that is discussed at length, by several major actors, who actively study the ***mapping of ideas and spaces***.

Peter Carruthers claims that ***there is no such thing as conscious thought***. It is worth listening to him –even if you disagree. But ***computers*** and machine-learning determine ***states of consciousness***, which may require a reevaluation of our approach(es).

There are new clues on ***how the brain maps time*** and how we see, judge and evaluate ***art***. But to ***understand art, think biology***.

Neuroplasticity is rarely associated with the ***regrowth*** of a large portion of the brain; but it does occur. And it teaches us **a lot!**

A group of independent young, brilliant mathematicians work on the ***Foundations of Neural Networks*** –ad they are just getting started!

A warning: ***philosophy can make the previously unthinkable thinkable*** -even happening. Then, ***can intelligence buy happiness?***

Body and Mind talk to each other, and it's not ***that*** simple. But aren't we all humans? What makes us *human* –and differ from all other primates- is ***laughter*** -and it's **not** laughing matter!

To close this review, do not worry too much: ***the universe is a hologram***!

Enjoy!

Huge Brain Study Uncovers "Buried" Genetic Networks Linked to Mental Illness

Enormous genomic analysis yields tantalizing insights into mechanisms behind conditions such as schizophrenia and bipolar disorder - Courtesy of Getty Images

Brain conditions such as schizophrenia and autism spectrum disorder have long been known to have an inherited component but pinpointing how gene variants contribute to has been a major challenge. Now, some of the first findings from the most comprehensive genomic analysis of the human brain ever undertaken are shedding light on the roots of these disorders.

Among the discoveries are elements buried in the genome's 'dark matter' that seem to regulate gene expression. Researchers have also uncovered previously unidentified networks of genes and the buried elements, which might contribute to the chances of developing such disorders.

"*We're not claiming in the remotest way to have figured out the underlying mechanism of these diseases, or how you would go about designing drugs, but we are highlighting genes, pathways and also cell types that are associated with these diseases*," says Mark Gerstein, a molecular biophysicist at Yale University in New Haven, Connecticut, who was involved in many of the project's studies, selection of which were published this week in *Science*.

Unlike disorders caused by mutations in a single gene -such as cystic fibrosis or some types of muscular dystrophy- neuropsychiatric disorders including schizophrenia involve hundreds of genes that interact with environmental factors. Each gene contributes only a small amount to the overall disease risk.

Over the past decade, scientists have identified numerous genetic variant that are associated with such disorders. But in many cases, it is not clear how the sequence changes alter gene function- if at all. "*Typically, when we do a genetic study, we might find 50 associated genetic variants all clustered in the same region of the genome, and maybe only one of them is directly influencing the risk of disease*," says Michael O'Donovan, a psychiatric geneticist at Cardiff University, UK. Further complicating matters, some of these variants fall in regions of DNA that do not code for proteins. Until the past few years, scientists presumed these areas to be wastelands. But buried within them are the codes for elements that regulate gene expression, such as transcription factors and microRNAs, which can also have a powerful influence on a person's disease risk.

Beyond the Genes

The PsychENCODE Consortium, which was founded by the US National Institutes of Health in 2015, aims to join the dots between these genetic associations and actual changes in gene function, by taking samples of brain tissue from thousands of cadavers and studying them using multiple genomic-sequencing techniques. "*We know that [common neuropsychiatric] diseases are extremely heritable, but people still don't have a good idea of mechanism; the goal is to use functional genomics to try to figure out what's going on*," says Gerstein.

One of these studies combined multiple types of sequencing data from brain tissue taken from 1,866 dead people, as well as from single brain-cell types. Previous studies have revealed widespread variation in gene expression between brains, but by comparing sequencing data from specific cell types with that from whole brains, the team established that around90% of this variance is related to the relative proportion of different cell types in individuals' brains- something that seems to vary with age, and in conditions such as autism. "*We could even figure out key genetic variants that were linked to increases in these cell types*," says Gerstein.

The researchers also used these data to draw connections between specific genes and noncoding DNA variants that had previously been linked to neuropsychiatric disease. This narrows the search for those that actually influence how genes function and seem to be contributing directly to conditions such as schizophrenia. "*Some of these genes and cell types are well known, but there are also some new ones that we find, that people could potentially follow up on*," Gerstein adds.

The Growing Brain

Gerstein and his colleagues also explored how gene expression; chemical, or '*epigenetic*', modifications to genes that can alter their expression; and regulatory elements in various regions of the brain vary during brain development, using samples of tissue and single cells taken from 60 brains. They found that the greatest variation in gene expression occurs during fetal development and adolescence, which are known to be crucial periods for brain development. During these periods, genes previously associated with neuropsychiatric disease risk seem to form networks in certain brain regions. This could provide new insights into when and where to study these disease mechanisms and model them, says Nenad Sestan, a neuroscientist also at Yale, whose lab led this study.

In a separate paper, other members of the PsychENCODE Consortium focused on the potential role that gains or losses of large chunks of DNA sequences called copy number variations (CNVs) might have in neuropsychiatric disease. Previous studies have suggested that rare CNVs can strongly affect schizophrenia risk, although the mechanism by which they do so is unclear. "*In the past, we have always concentrated on CNVs affecting protein-coding genes, but there has been a blind spot which is CNVs in regions containing long-noncoding RNAs*," says Chunyu Liu, a specialist in psychiatry and behavioral sciences at SUNY Upstate Medical University in Syracuse, New York, who led this study. Although such molecules show no protein-coding potential, some of them are still capable of regulating gene expression and might contribute to schizophrenia risk in their own right.

Liu and his colleagues analysed brain tissue from 259 cadavers, focusing on long-noncoding RNAs (lncRNAs) in ten CNV-deletion regions that have previously been associated with an elevated risk of schizophrenia, looking to see if the expression of any of them correlated with that of protein-coding genes, which might imply a relationship. This led them to several lncRNAs that they suspect might help to regulate gene expression. One was called *DGCR5*; further experiments in neural progenitor cells revealed that it serves as a hub for several schizophrenia-related genes, potentially explaining why its absence is associated with an increased risk of the disease. In a related study, Liu and his colleagues analysed brain tissue from people with schizophrenia or bipolar disorder, and from healthy controls. They looked for microRNAs whose expression correlated with that of protein-coding genes. This led them to a network of microRNAs, transcription factors and genes that seem to work together to influence schizophrenia risk. By focusing on such networks, rather than simply on the influence of single genes, Liu hopes to improve understanding of the root causes of complex diseases such as schizophrenia. Even so, he emphasizes that this is just the beginning of a long journey to understand how variation in such regions affects gene expression, and how this, in turn, contributes to disease risk. O'Donovan agrees. "*These publications are important, but they do not* *provide the definitive answer to how genetic changes contribute to brain diseases,"* he says. "*These are reasonably substantial steps, but they are just steps- although we do hope that a lot more work of this sort will help us link the genetics to the biology of these disorders."*

The Cerebellum Is Your "Little Brain"—and It Does Some Pretty Big Things

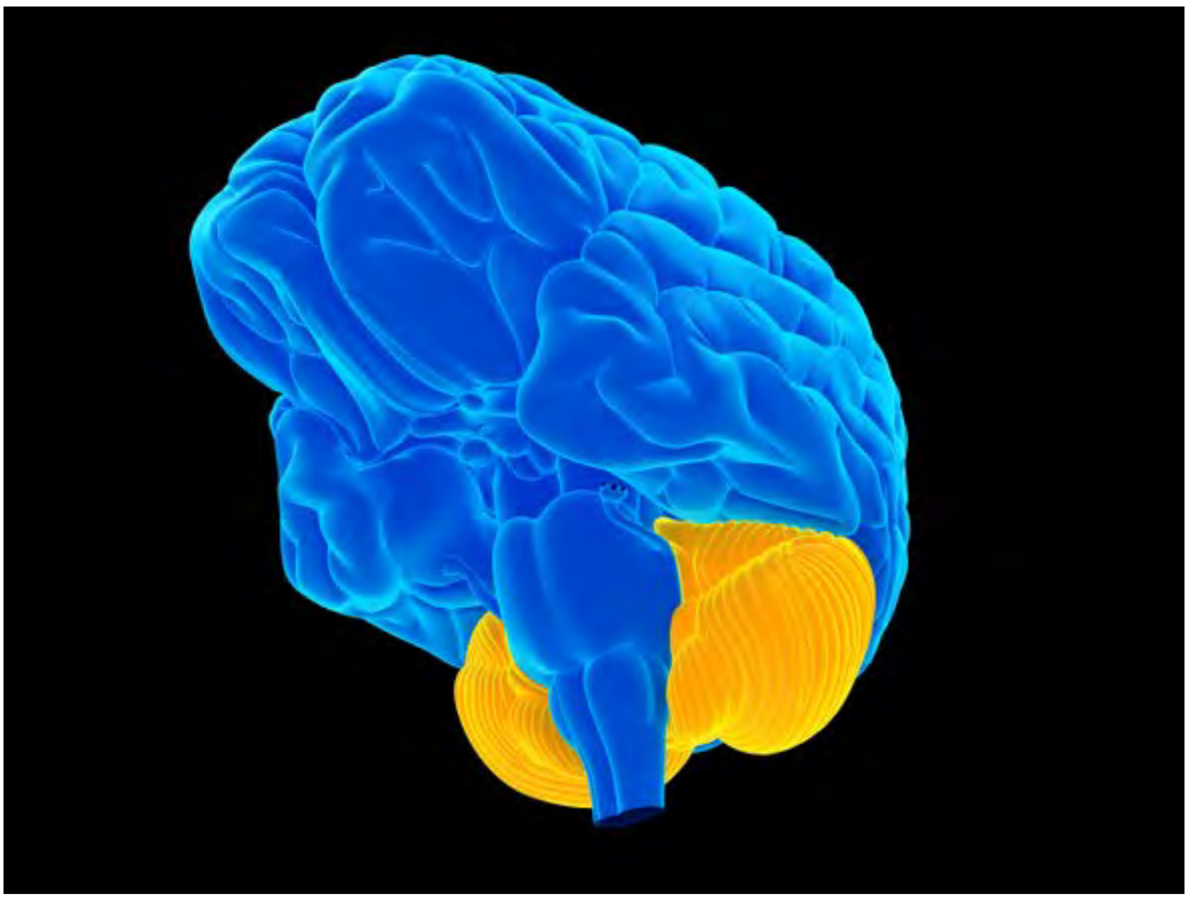

A newly identified circuit connecting the cerebellum to the brain's reward centers in mice - Courtesy of Getty Images

For the longest time the cerebellum, a dense, fist-size formation located at the base of the brain, never got much respect from neuroscientists. For about two centuries the scientific community believed the cerebellum (Latin for "*little brain*"), which contains approximately half of the brain's neurons, was dedicated solely to the control of movement. In recent decades, however, the tide has started to turn, as researchers have revealed details of the structure's role in cognition, emotional processing and social behavior. The longstanding interest in the cerebellum can be seen in the work of French physiologist Marie Jean Pierre Flourens (1794–1867). Flourens removed the cerebella of pigeons and found the birds became unbalanced, although they could still move. Based on these observations, he concluded the cerebellum was responsible for coordinating movements. *"[This] set the dogma that the cerebellum was involved in motor coordination*," says Kamran Khodakhah, a neuroscientist at Albert Einstein College of Medicine, adding: "*For many years, we ignored the signs that suggested it was involved in other things*."

One of the strongest pieces of evidence for the cerebellum's broader repertoire emerged around two decades ago, when Jeremy Schmahmann, a neurologist at Massachusetts General Hospital, described cerebellar cognitive affective syndrome after discovering behavioral changes such as impairments in abstract reasoning and regulating emotion in individuals whose cerebella had been damaged. Since then this line of study has expanded. There has been human neuroimaging work showing the cerebellum is involved in cognitive processing and emotional control -and investigations in animals have revealed, among other things, that the structure is important for the normal development of social and cognitive capacities. Researchers have also linked altered cerebellar function to addiction, autism and schizophrenia.

Although many of these findings suggested the cerebellum played an important part both in reward-related and social behavior, a clear neural mechanism to explain this link was lacking. New research, published this week in *Science*, demonstrates that a pathway directly tying the cerebellum to the ventral tegmental area (VTA) -one of the brain's key pleasure centers- can control these two processes. "*This work helps lay out the circuitry connecting the cerebellum to social and reward processing*," says Julie Fiez, a cognitive neuroscientist at the University of Pittsburgh who was not involved in this study. "*I think it's really exciting*."

Khodakhah, one of the study's authors, had focused his work on the cerebellum's role in motor coordination until he stumbled across the literature on the structure's nonmotor functions while reviewing grants. Intrigued by the cerebellum's links to conditions such as autism and addiction, he set out to investigate whether it may directly communicate with the VTA, an area of the brain previously linked to these disorders. Earlier investigations in his lab had hinted there might be unexpected connections between the cerebellum and other parts of the brain. Specifically, while examining the brain circuits underlying dystonia -a movement disorder that causes uncontrollable muscle contractions- in mice, Khodakhah's team discovered the cerebellum directly communicated with the basal ganglia (involved in movement, motivation and reward functions) to control complex movements. It was previously thought that to coordinate such actions, the two brain areas communicated via the cortex, the region responsible for higher-order tasks such as planning and decision-making. "That really fueled us to start looking at the direct cerebellar manipulation of other brain structures," Khodakhah says. To investigate the link between the cerebellum and the VTA, Khodakhah's team first injected the cerebellar cells of mice with herpes viruses, which act as mobile sentinels as they jump through synapses -the tiny gaps between brain cells- while carrying fluorescent tags. This experiment revealed several neurons in the VTA lit up with the glowing markers, indicating that cells in this brain region were, indeed, receiving direct connections from the cerebellum. Then, using optogenetics (a method that allows scientists to switch on or off specific cells in a neural pathway with flashes of light), the researchers demonstrated that stimulating the cerebellar neurons could activate cells in the VTA. Next, the team tested whether this circuit could influence both reward-related and social behaviors. They found that stimulating this pathway with optogenetics while mice explored one quadrant of a square enclosure caused them to develop a strong preference for the spot. By activating this pathway, the scientists were also able to condition the rodents -which are nocturnal- to favor exploring a bright compartment, despite their natural preference for dark places. "These findings suggest that this pathway could be involved in addictive behavior," Khodakhah says. He notes the latter experiment has been extensively used to study drug addiction in animals, and his group plans further studies. A future experiment might supply cocaine to rodents to see whether inhibiting the pathway between the cerebellum and the VTA can manipulate addictive behaviors.

When the researchers conducted similar mouse experiments using three interconnected chambers, they made an interesting discovery. The mice encountered a familiar animal that had been placed in one compartment (the "social chamber"). Adjoining it was an empty compartment (the "object chamber"). Mice typically spent more time in the social compartment. But after deactivating the cerebellum–VTA pathway using optogenetics, that preference disappeared, mirroring the behavior typically observed when scientists conduct the same test with animal models of autism.

Interestingly, the team found stimulating this circuit did not increase the rodents' interactions with an unfamiliar animal. According to the authors, this observation suggests the pathway does not necessarily increase pro-social behaviors but instead makes inanimate objects, for example, just as rewarding as interacting with others. "This [study] is one of the most clear and interesting demonstrations that the cerebellum is indeed involved in the control of high-level, nonmotor functions," says Egidio d'Angelo, a neurophysiologist at the University of Pavia in Italy who was not part of the work but penned a commentary accompanying the paper. "But this work is done in mice -now we have to see whether this happens in humans."

Schmahmann, who also did not take part in the study, notes these findings confirm the existence of a pathway first proposed by scientists several decades ago. "I was delighted to see [this research]," he adds. "They provide another really important building block in our ongoing attempt to [understand] the cerebellar contribution to cognition and emotion." Further probing the cerebellum–VTA circuit could one day help scientists treat various disorders, Khodakhah says. This circuit might be manipulated -using techniques such as transcranial magnetic stimulation or deep-brain stimulation- in individuals with addiction or autism. But more research is necessary before such interventions become reality -and for now, Khodakhah's team plans to test some of these methods in mice. "This is a really exciting time for cerebellar research," Khodakhah says. "I think over the next few years we'll see that the cerebellum plays a more and more prominent role in nonmotor functions, [such as] cognitive and emotional processing."

For Alzheimer's Sufferers, Brain Inflammation Ignites a Neuron-Killing "Forest Fire"

And it could also be the kindling sparking Parkinson's and other neurodegenerative maladies. For decades, researchers have focused their attacks against Alzheimer's on two proteins, amyloid beta and tau. Their buildup in the brain often serves as a defining indicator of the disease. Get rid of the amyloid and tau, and patients should do better, the thinking goes. - ***Courtesy Scientific American.***

But drug trial after drug trial has failed to improve patients' memory, agitation and anxiety. (I am thinking of a *"professor"* at our Department of Hong Kong Polytechnic University for whom poor **memory** was identical to Alzheimer's Disease and could be treated [even *cured*] with derivatives of *tacrine*!) One trial of a drug that removes amyloid even seemed to make some patients worse. The failures suggest researchers

were missing something. A series of observations and recently published research findings have hinted at a somewhat different path for progression of Alzheimer's, offering new ways to attack a disease that robs memories and devastates the lives of 5.7 million Americans and their families.
One clue hinting at the need to look further afield was a close inspection of the 1918 worldwide flu pandemic, which left survivors with a higher chance of later developing Alzheimer's or Parkinson's. A second inkling came from the discovery that the amyloid of Alzheimer's and the alpha-synuclein protein that characterizes Parkinson's are antimicrobials, which help the immune system fight off invaders. The third piece of evidence was the finding in recent years, as more genes involved in Alzheimer's have been identified, that traces nearly all of them to the immune system. Finally, neuroscientists have paid attention to cells that had been seen as ancillary— "*helper*" or "*nursemaid*" cells. They have come to recognize these brain cells, called microglia and astrocytes, play a central role in brain function—and one intimately related to the immune system.

All these hints are pointing toward the conclusion that both Alzheimer's and Parkinson's may be the results of neuroinflammation—in which the brain's immune system has gotten out of whack. "*The accumulating evidence that inflammation is a driver of this disease is enormous*," says Paul Morgan, a professor of immunology and a member of the Systems Immunity Research Institute at Cardiff University in Wales. "*It makes very good biological sense.*"

The exact process remains unclear. In some cases, the spark that starts the disease process might be some kind of insult—perhaps a passing virus, gut microbe or long-dormant infection. Or maybe in some people, simply getting older—adding some pounds or suffering too much stress could trigger inflammation that starts a cascade of harmful events.

This theory also would explain one of the biggest mysteries about Alzheimer's: why some people can have brains clogged with amyloid plaques and tau tangles and still think and behave perfectly normally. "*What made those people resilient was lack of neuroinflammation*," says Rudolph Tanzi, a professor of neurology at Harvard Medical School and one of the leaders behind this new view of Alzheimer's. Their immune systems kept functioning normally, so although the spark was lit, the forest fire never took off, he says. In Tanzi's fire analogy, the infection or insult sparks the amyloid match, triggering a brush fire. As amyloid and tau accumulate, they start interfering with the brain's activities and killing neurons, leading to a raging inflammatory state that impairs memory and other cognitive capacities. The implication, he says, is that it is not enough to just treat the amyloid plaques, as most previous drug trials have done. "*If you try to just treat plaques in those people, it's like trying to put out forest fire by blowing out a match*."

Lighting the Fire

One study published earlier this year found gum disease might be the match that triggers this neuroinflammatory conflagration—but Tanzi is not yet convinced. The study was too small to be conclusive, he says. Plus, he has tried to find a link himself and found nothing. Other research has suggested the herpes virus could start this downward spiral, and he is currently investigating whether air pollution might as well. He used to think amyloid took years to develop, but he co-authored a companion paper to the herpes one last year, showing amyloid plaques can literally appear overnight.

It is not clear whether the microbes—say for herpes or gum disease—enter the brain or whether inflammation elsewhere in the body triggers the pathology, says Jessica Teeling, a professor of experimental neuroimmunology at the University of Southampton in England. If microbes can have an impact without entering the brain or spinal cord—staying in what's called the peripheral nervous system—it may be possible to treat Alzheimer's without having to cross the blood–brain barrier, Teeling says.

Genetics clearly play a role in Alzheimer's, too. Rare cases of Alzheimer's occurring at a relatively young age result from inheriting a single dominant gene. Another variant of a gene that transports fats in brain cells, APOE4, increases risk for more typical, later-onset disease. Over the last five years or so large studies of tens of thousands of people have looked across the human genome for other genetic risk factors. About 30 genes have jumped out, according to Alison Goate, a professor of neurogenetics and director of the Loeb Center for Alzheimer's Disease at Icahn School of Medicine at Mount Sinai in New York City. Goate, who has been involved in some of those studies, says those genes are all involved in how the body responds to tissue debris—clearing out the gunk left behind after infections, cell death and similar insults. So, perhaps people with high genetic risk cannot cope as well with the debris that builds up in the brain after an infection or other insult, leading to a quicker spiral into Alzheimer's. "Whatever the trigger is, the tissue-level response to that trigger is genetically regulated and seems to be at the heart of genetic risk for Alzheimer's disease," she says. When microglia—immune cells in the brain—are activated in response to tissue damage, these genes and APOE get activated. "How microglia respond to this tissue damage—that is at the heart of the genetic regulation of risk for Alzheimer's," she says.

But APOE4 and other genes are part of the genome for life, so why do Alzheimer's and Parkinson's mainly strike older people? says Joel Dudley, a professor of genetics and genomics, also at Mount Sinai. He thinks the answer is likely to be inflammation, not from a single cause for everyone but from different immune triggers in different individuals.

Newer technologies that allow researchers to examine a person's aggregate immune activity should help provide some of those answers, he says. Cardiff's Morgan is developing a panel of inflammatory markers found in the blood to predict the onset of Alzheimer's before much damage is done in the brain, a possible diagnostic that could point to the need for anti-inflammatory therapy.

Like Threads

A similar inflammatory process is probably also at play in Parkinson's disease, says Ole Isacson, a professor of Neurology at Harvard Medical School. Isacson points to another early clue about the role of inflammation in Parkinson's: people who regularly took anti-inflammatory drugs like ibuprofen developed the disease one to two years later than average. Whereas other researchers focused exclusively on genetics, Isacson found the evidence suggested the environment had a substantial impact on who got Parkinson's.

In 2008–09, Isacson worked with a postdoctoral student on an experiment trying to figure out which comes first in the disease process: inflammation or the death of dopamine-producing neurons, which make the brain chemical involved in transmitting signals among nerve cells. The student first triggered inflammation in the brains of some rodents with molecules from gram-negative bacteria and then damaged the neurons that produce dopamine. In another group of rodents, he damaged the neurons first and then introduced inflammation. When inflammation came first, the cells died en masse, just as they do in Parkinson's disease. Blocking inflammation prevented their demise, they reported in *The Journal of Neuroscience*.

Other neurodegenerative diseases also have immune connections. In multiple sclerosis, which usually strikes young people, the body's immune system attacks the insulation around nerve cells, slowing the transmission of signals in the body and brain.

The spinal fluid of people with MS include antibodies and high levels of white blood cells, indicating the immune system is revved up—although it is not clear whether that immune system activation is the cause or result of MS, says Mitchell Wallin, who directs the Veterans Affairs Multiple Sclerosis Centers of Excellence. People with antibodies to the Epstein–Barr virus in their systems, especially if they caught the virus in late adolescence or early adulthood run a higher risk of developing MS—supporting the idea that an infection plays a role in MS.

Thanks to newer medications and improvements in fighting infections, people with MS are now living longer. This increased longevity puts them at risk for neurological diseases of aging, including Alzheimer's and Parkinson's, Wallin says. Lack of data has left it unclear whether people with MS are at the same, higher or lower risk for these diseases than the general population. "How common it is, we're just starting to explore right now," Wallin says.

Coming Soon?

It will be years before the concept of a neuroinflammatory can be fully tested, but there are already some relevant drugs in development. One start-up, California-based INmune Bio, recently received a $1-million grant from the Alzheimer's Association to advance XPro1595, a drug that targets neuroinflammation. The company is beginning its first clinical trial this spring, treating 18 patients with mild to moderate-stage Alzheimer's who also show signs of inflammation. The company plans to test blood, breath by-products and cerebral spinal fluid as well as conduct brain scans to look for changes in inflammatory markers. That first trial will just explore if XPro1595 can safely bring down inflammation and change behaviors such as depression and sleep disorders. Company CEO and co-founder Raymond Tesi says he expects to see those indicators improve, even in a short, three-month trial.

The best way to avoid Alzheimer's is to prevent it from ever starting, which might require keeping brain inflammation to a minimum, particularly in later life. Preventative measures are already well known: eat healthy foods, sleep well, exercise regularly, minimize stress and avoid smoking and heavy drinking.

You can't do anything about your genetics but living a healthy lifestyle will help control your inheritance, says Tanzi, who, along with Deepak Chopra, wrote a book on the topic, "*The Healing Self: A Revolutionary New Plan to Supercharge Your Immunity and Stay Well for Life*. "*It's important to get that set point as high as possible*."

The Brain Maps Out Ideas and Memories Like Spaces

Emerging evidence suggests that the brain encodes abstract knowledge in the same way that it represents positions in space, which hints at a more universal theory of cognition.

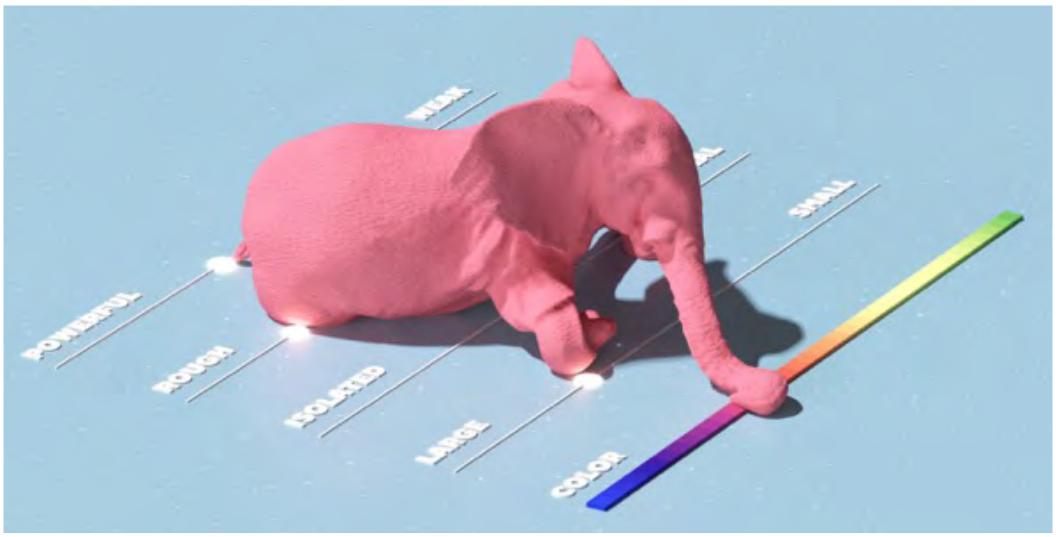

Photo Courtesy of Quanta Magazine

In the parable of the blind men and the elephant, each paid attention to a different aspect of the creature. The brain may do something similar by mapping out the qualities of perceptions, experiences and abstract concepts along various dimensions, with the help of the same system that it uses to map out physical spaces.

We humans have always experienced an odd -and oddly deep- connection between the mental worlds and physical worlds we inhabit, especially when it comes to memory. We're good at remembering landmarks and settings, and if we give our memories a location for context, hanging on to them becomes easier. To remember long speeches, ancient Greek and Roman orators imagined wandering through "memory palaces" full of reminders. Modern memory contest champions still use that technique to "place" long lists of numbers, names and other pieces of information.

As the philosopher Immanuel Kant put it, the concept of space serves as the organizing principle by which we perceive and interpret the world, even in abstract ways. "*Our language is riddled with spatial metaphors for reasoning, and for memory in general,*" said Kim Stachenfeld, a neuroscientist at the British artificial intelligence company DeepMind. In the past few decades, research has shown that for at least two of our faculties, memory and navigation, those metaphors may have a physical basis in the brain. A small seahorse-shaped structure, the hippocampus, is essential to both those functions, and evidence has started to suggest that the same coding scheme -a grid-based form of representation- may underlie them. Recent insights have prompted some researchers to propose that this same coding scheme can help us navigate other kinds of information, including sights, sounds and abstract concepts. The most ambitious suggestions even venture that these grid codes could be the key to understanding how the brain processes all details of general knowledge, perception and memory.

The Amnesiac and the Hexagons

On September 1, 1953, Henry Molaison, a 27-year-old man the world would come to know as "*Patient H.M.,*" went under the knife in a risky, experimental bid to cure a debilitating case of epilepsy. A neurosurgeon removed the hippocampus and surrounding tissues from deep within H.M.'s brain, alleviating some of his seizures but inadvertently leaving him a permanent amnesiac. Until his death more than half a century later, H.M. couldn't encode new memories: not what he'd had for breakfast, nor the most recent news headline, nor the identity of the stranger he'd been introduced to just a few minutes earlier. H.M.'s story, though tragic, revolutionized scientists' understanding of the role the hippocampus plays in how the brain organizes memory. Years later, another hippocampus-centered revolution transpired and earned its pioneers a Nobel Prize: the discoveries, decades apart, of two types of cells, which made it clear that the hippocampal region's fundamental functions included not just memory but also navigation and the representation of two-dimensional spaces.

Neuroscientist John O'Keefe's discovery of place cells, a major component of the brain's navigational system, jump-started research on cognitive maps in the hippocampus.
- Courtesy of John O'Keefe

The first of these came in 1971, when researchers uncovered "place cells," which essentially fire to indicate one's current location. John O'Keefe, a neuroscientist at University College London, and his colleagues monitored the brain activity of freely roaming rats and observed that some of their neurons fired only when they were in specific parts of their cages. Some became active as a rat sniffed around, say, its enclosure's northeast corner, but otherwise remained quiet; others fired in the cage's center. That is, the cells encoded a sense of place ("*you are here*") - and together, they created a map of the entire space. (When the rat was put in a different cage or room, these place cells "*remapped,*" encoding different local positions.) These findings inspired the proposal that the hippocampus might be creating and storing "*cognitive maps*" (an idea first put forth by psychologist Edward Tolman in the 1940s to explain how rats could suss out new shortcuts to rewards in mazes) beyond spatial ones. At the very least, the hippocampus seemed like a promising place to start looking for hints of such maps. That work eventually led a then-married pair of scientists at the Norwegian University of Science and Technology, May-Britt Moser and Edvard Moser, to direct their attention to the entorhinal cortex, located just next door to the hippocampus. The region provides major inputs to the hippocampus -and is also one of the first areas of the brain to deteriorate in Alzheimer's disease, which affects both navigation and memory. There, the researchers found what they called grid cells, which experts now think may be the most compelling candidate for cognitive mapmaker. Unlike the place cells, grid cells do not represent particular locations. Instead, they form a coordinate system that's independent of location. (As a result, they're popularly known as the brain's GPS.) Each grid cell fires at regularly spaced positions, which form a hexagonal pattern. Imagine the floor of your bedroom is tiled with regular hexagons, all the same size, and each hexagon is divided into six equilateral triangles. As you walk across the room, one of your grid cells fires every time you reach a vertex of any of those triangles.

Different sets of grid cells form different grids: grids with larger or smaller hexagons, grids oriented in other directions, grids offset from one another. Together, the grid cells map every spatial position in an environment, and any particular location is represented by a unique combination of grid cells' firing patterns. The single point where various grids overlap tells the brain where the body must be. This kind of grid network, or code, constructs a more intrinsic sense of space than the place cells do. While place cells provide a good means of navigating where there are landmarks and other meaningful locations to provide spatial information, grid cells provide a good means of navigating in the absence of such external cues. In fact, researchers think that grid cells are responsible for what's known as path integration, the process by which a person can keep track of where she is in space -how far she has traveled from some starting point, and in which direction -while, say, blindfolded.

"The idea is that the grid code could therefore be some sort of metric or coordinate system," said Jacob Bellmund, a cognitive neuroscientist affiliated with the Max Planck Institute in Leipzig and the Kavli Institute for Systems Neuroscience in Norway. "*You can basically measure distances with this kind of code*." Moreover, because of how it works, that coding scheme can uniquely and efficiently represent a lot of information. And not just that: Since the grid network is based on relative relations, it could, at least in theory, represent not only a lot of information but a lot of different types of information, too. "*What the grid cell captures is the dynamic instantiation of the most stable solution of physics*," said György Buzsáki, a neuroscientist at New York University's School of Medicine: "*the hexagon*." Perhaps nature arrived at just such a solution to enable the brain to represent, using grid cells,

any structured relationship, from maps of word meanings to maps of future plans.

The Brain's Navigational System

Specialized neurons in the hippocampus and neighboring brain areas enable navigation through space. Place cells encode specific local positions, while grid cells provide a longer-range coordinate system for determining position in the absence of external cues. Researchers think the brain may use a similar system to navigate through mental "spaces" of abstract knowledge and memories.

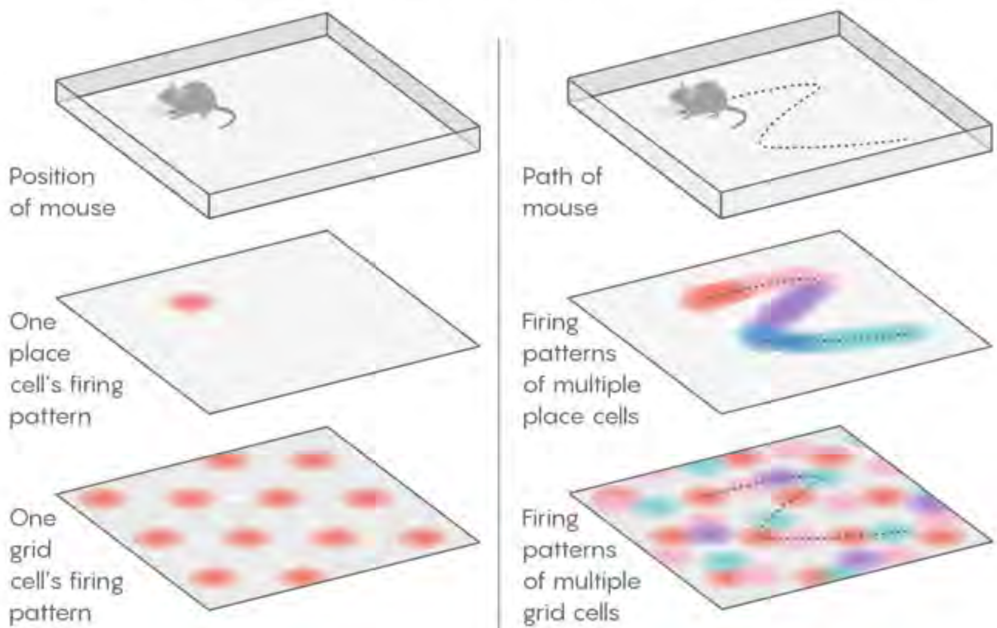

Mapping One Location

If a mouse is in one corner of a room, then there is one place cell that fires uniquely at that location. A grid cell that fires at that location also fires at other positions around it in a hexagonal array.

Mapping a Path

As the mouse moves, the activity of many place cells records the locations that it visited. Grid cell activity tracks how the mouse moved through overlapping hexagonal coordinate systems that tile the plane.

Lucy Reading- Copyright Ikkanda/Quanta Magazine

An Expanding Role for Grid Cells

"We've been thinking about how the hippocampus and entorhinal cortex machinery could have a more general purpose," Stachenfeld said. *"It's a really powerful idea, that you can have a [grid cell] representation of structure in general, and apply it more rapidly to new situations."* That, in turn, would allow one *"to behave more efficiently, to learn a lot faster."*

Since researchers usually could not take direct measurements of individual neurons in their test subjects, they had to get clever with their methodology. In 2010, for instance, neuroscientists figured out a certain kind of signal to look for in functional magnetic resonance imaging (fMRI) scans of the brain as an indirect signature of grid cell activity. This "*hexadirectional*" signal emerges in subjects navigating a virtual environment. As it turns out, it also characterizes other tasks, some spatial, some not so much.

Neuroscientists May-Britt Moser and Edvard Moser uncovered grid cells in the entorhinal cortex, which completed a picture of the brain's navigational system that O'Keefe had started decades earlier. – Copyright Geir Mogen/ Kavli Institute for Systems Neuroscience

One of the earliest examples came with behavior that fell somewhere between the two: the navigation of visual space. When monkeys, with their heads fixed in place, tracked images with just their eyes, researchers found evidence of grid cell activity in the entorhinal cortex. More recent work in humans has uncovered the same hexadirectional signature, and some experiments have even pinpointed other, more direct properties of the grid code already observed in physical navigation tasks. Similar principles may also guide how the brain encodes time. The hippocampus has already been found to contain place cells that also behave as "*time cell*" neurons in certain situations, activating to indicate successive moments in time (rather than successive positions in space). Rats would run through a maze, in which one section involved trotting in place on a wheel or treadmill for some predetermined number of seconds before continuing onwards.

Cognitive neuroscientist Jacob Bellmund studies the intersection between the hippocampus's two major functions, spatial navigation and episodic memory. He and his colleagues propose that the key to such a unifying framework lies in the brain's grid cells. – Copyright MPI CBS

During the interval when the rats ran in place, their actual location held constant, cells fired in their hippocampus to track their temporal progression: some neurons were active for the first few seconds, others for the next few and so on. The finding "*brings time as a different dimension into the equation*," Bellmund said.

More recently, work published in *Nature* last summer turned up evidence for a coding system that uniquely represents time in the context of memories or experiences. A team of researchers, led by the Mosers, uncovered a coding scheme for time that spanned multiple scales, from seconds to hours. Although no explicit link has yet been drawn between temporal organization and grid cells, scientists have seen hints of a connection: Grid cells signal elapsed time in rats running on treadmills, for instance. Last year, a team of scientists at Princeton University brought yet another potential dimension into the mix: ***sound***. They monitored brain activity in rats that were pushing a lever to change the frequency of an emitted tone to match one they had previously heard. Their observations hinted that the rats might be mentally navigating through an "*acoustic space*" in their minds to find the desired tone.

Perhaps most tantalizing of all, an experiment conducted in 2016 introduced a far more abstract context for grid cell behavior. Researchers led by Timothy Behrens, a computational neuroscientist at the University of Oxford, had people watch the silhouette of a bird on a screen as the length of its neck, the length of its legs or both were stretched and compressed. The hexadirectional signal arose in their fMRI data, in several areas of the brain; it varied just as if the test subjects were navigating a two-dimensional "*bird* space," where one axis denoted neck length, the other leg length. The finding suggested that the brain processes trajectories through physical spaces and conceptual spaces in much the same way. Now, researchers including Behrens, Bellmund and neuroscientist Christian Doeller propose that all knowledge can be plotted this way in terms of features of interest -that different objects, different experiences and different memories can be organized and traversed with the grid code. "*It seems to be quite arbitrary, what dimensions it can map*," Bellmund said. "*What's interesting is that it seems to be so general across domains, but the mechanism seems to be preserved*."

This work, added Thomas Wolbers, a cognitive neuroscientist at the German Center for Neurodegenerative Diseases, calls into question the idea that grid cells simply constitute "*a pure location signal*" -hardwired and specialized. "*So far, we'd only seen it in space because we'd only looked at navigation tasks and paradigms*," he said. "*It may be much more ubiquitous.*"

The Power of Analogy

One area that's seen some intriguing preliminary results is in social behavior. We think of society in spatial terms all the time: There are social ladders to climb, networks to build and expand, people we consider "*close*" or "*distant.*" Now, some research groups are probing social relationships for evidence of the grid code.

One recent study built up a two-dimensional space not unlike the bird experiment: People played a computer game, interacting with characters in ways that could change their levels of power or affiliation. The researchers found that the hippocampus seemed to track the positions of the characters in that space, relative to the test subject. Although the experiment did not determine whether the hippocampus is navigating that social information in a gridlike way, Matthew Schafer, a graduate student at the Icahn School of Medicine at Mount Sinai currently working on the project, expects to find the telltale hexadirectional signal. (He and others are now studying how that navigation might be disrupted or otherwise affected in people with conditions like autism spectrum disorder.)

Kim Stachenfeld, a research scientist on DeepMind's neuroscience team, hopes that understanding how the hippocampus and surrounding brain regions formulate a general representation of structure could inform better machine learning techniques.
- Copyright Mike Dodd

These ideas could make it worthwhile to pursue clues hidden in other kinds of spatial metaphors, too: Neurons beyond place cells and grid cells, after all, might also have something to contribute. There are head direction cells that fire when an animal points its head in a particular direction, and speed cells that indicate the rate at which one moves through space, and even boundary cells that represent the location of walls or other environmental borders. Studying these neurons in more abstract contexts might yield new insights. For instance, boundary cell activity has been reported for not just the borders of a physical space but also the borders between separate events in a temporal sequence. Could these neurons also play a role in forming borders between concepts, in creating distinct domains of knowledge in the brain? Or could head direction cells help one orient oneself within a given topic? The potential for such analogies is enormous.

The same goes for gaining a better understanding of diseases and other states. Wolbers studies aging, and in one recently published paper, he and his colleagues examined how the spatial navigation grid code changes in elderly people. They found that the signal became less stable, with the grid fluctuating between orientations - and that people with less stable grids were also much less adept at keeping track of their relative location when blindfolded and led along a circuitous course. Wolbers suggests that if the grid code is used to process many kinds of information and memories, it's possible that a pathology that destabilizes the spatial grid system might have a more general effect on the stability of memory and other areas of cognition. Still, "*at this stage*," he warned, "*the available data are scarce. We have to be cautious.*"

Kate Jeffery, a behavioral neuroscientist at University College London, agreed. Sure, the brain might use a common system to encode spatial and nonspatial knowledge, if the latter can be represented as varying continuously on a two-dimensional scale. But it's also possible that some cognitive tasks are so complicated and unnatural that the brain is forced to rely on a spatial analog as a crutch to get through them. Perhaps the experiments on sound frequency and stretched birds tapped into this feature, Jeffery said.

A Unifying Framework

To further cement the grid code's broader applications, then, researchers first hope to figure out how these cells may be working in more than two dimensions, given that higher-level knowledge tends to involve far more than pairs of qualities, like neck length and leg length, or power and association. This is something that's currently being examined in flying bats, which navigate through three dimensions rather than just two. Some researchers are making even bolder claims. Jeff Hawkins, the founder of the machine intelligence company Numenta, leads a team that's working on applying the grid code not just to explain the memory-related functions of the hippocampal region but to understand the entire neocortex - and with it, to explain all of cognition, and how we model every aspect of the world around us. According to his "*thousand brains theory of intelligence*", he said, "*the cortex is not just processing sensory input alone, but rather processing and applying it to a location.*" When he first thought of the idea, and how grid cells might be facilitating it, he added, "*I jumped out of my chair, I was so excited.*" Imagine closing your eyes and wrapping your hands around an unidentified object: in this case, a coffee cup. Hawkins posits that the brain takes in information about the position of each patch of skin touching the cup's surface, relative to the cup itself -much as the grid code allows you to know your body's position in space, relative to the room you're in. Each patch of skin generates an independent model of what it may be touching; all those models then get cross-referenced to reach the conclusion that the object is indeed a coffee cup. Hawkins thinks the same logic can apply to anything with a structured framework. "*Everything we do -planning, mathematics, physics, language -would be based on the same principle,*" he said. "*I think we're on a cusp here, where all of a sudden we're going to have a new paradigm for understanding how the brain works*."

While the hypothesis has piqued interest among other researchers, they remain skeptical that grid cells will be found beyond the vicinity of the hippocampus and say that Hawkins and his team have a long way to go to prove the power of their model. Still, it provides a good starting point for thinking about how to improve artificial intelligence. If the grid framework is indeed a general one, it could be mimicked to build machines that are far more flexible, creative, general and powerful. The field is just starting to grapple with these notions. For now, researchers are continuing to probe the activity of the hippocampus in a slew of different contexts, in hopes of finally uniting its memory and navigation functions once and for all. "*When conceptual and cognitive ideas really start to connect with the very low-level neural data*," Stachenfeld said, "*it's really very satisfying.*"

There Is No Such Thing as Conscious Thought

Philosopher Peter Carruthers insists that conscious thought, judgment and volition are illusions. They arise from processes of which we are forever unaware - Courtesy of Getty Images

Peter Carruthers, Distinguished University Professor of Philosophy at the University of Maryland, College Park, is an expert on the philosophy of mind who draws heavily on empirical psychology and cognitive neuroscience. He outlined many of his ideas on conscious thinking in his 2015 book "The Centered Mind: What the Science of Working Memory Shows Us about the Nature of Human Thought". More recently, in 2017, he *published a paper with the astonishing title of "The Illusion of Conscious Thought." In the following excerpted conversation, initially published in Gehirn & Geist, Carruthers explains the reasons for his provocative proposal.*

What makes you think conscious thought is an illusion?

I believe that the whole idea of conscious thought is an error. I came to this conclusion by following out the implications of the two of the main theories of consciousness. The first is what is called the Global Workspace Theory, which is associated with neuroscientists Stanislas Dehaene and Bernard Baars. Their theory states that to be considered conscious a mental state must be among the contents of working memory (the "*user interface*" of our minds) and thereby be available to other mental functions, such as decision-making and verbalization. Accordingly, conscious states are those that are "globally broadcast," so to speak. The alternative view, proposed by Michael Graziano, David Rosenthal and others, holds that conscious mental states are simply those that you know of, that you are directly aware of in a way that doesn't require you to interpret yourself. You do not have to read your own mind to know of them. Now, whichever view you adopt, it turns out that thoughts such as decisions and judgments should not be considered to be conscious. They are not accessible in working memory, nor are we directly aware of them. We merely have what I call "*the illusion of immediacy*"—the false impression that we know our thoughts directly.

One might easily agree that the sources of one's thoughts are hidden from view—we just don't know where our ideas come from. But once we have them and we know it, that's where consciousness begins. Don't we have conscious thoughts at least in this sense?

In ordinary life, we are quite content to say things like "*Oh, I just had a thought*" or *"I was thinking to myself."* By this we usually mean instances of inner speech or visual imagery, which are at the center of our stream of consciousness -the train of words and visual contents represented in our minds. I think that these trains are indeed conscious. In neurophilosophy, however, we refer to *"thought*" in a much more specific sense. In this view, thoughts include only nonsensory mental attitudes, such as judgments, decisions, intentions and goals. These are amodal, abstract events, meaning that they are not sensory experiences and are not tied to sensory experiences. Such thoughts never figure in working memory. They never become conscious. And we only ever know *of* them by interpreting what does become conscious, such as visual imagery and the words we hear ourselves say in our heads.

So, consciousness always has a sensory basis?

I claim that consciousness is always bound to a sensory modality, that there is inevitably some auditory, visual or tactile aspect to it. All kinds of mental imagery, such as inner speech or visual memory, can of course be conscious. We see things in our mind's eye; we hear our inner voice. What we are conscious of are the sensory-based contents present in working memory.

In your view, is consciousness different from awareness?

That's a difficult question. Some philosophers believe that consciousness can be richer than what we can actually report. For example, our visual field seems to be full of detail -everything is just there, already consciously seen. Yet experiments in visual perception, especially the phenomenon of inattentional blindness, show that in fact we consciously register only a very limited slice of the world. [*A person experiencing inattentional blindness may not notice that a gorilla walked across a basketball court while the individual was focusing on the movement of the ball.*] So, what we think we see, our subjective impression, is different from what we are actually aware of. Probably our conscious mind grasps only the gist of much of what is out there in the world, a sort of statistical summary. Of course, for most people consciousness and awareness coincide most of the time. Still, I think, we are not directly aware of our thoughts. Just as we are not directly aware of the thoughts of other people. We interpret our own mental states in much the same way as we interpret the minds of others, except that we can use as data in our own case our own visual imagery and inner speech.

You call the process of how people learn their own thoughts interpretive sensory access, or ISA. Where does the interpretation come into play?

Let's take our conversation as an example -you are surely aware of what I am saying to you at this very moment. But the interpretative work and inferences on which you base your understanding are not accessible to you. All the highly automatic, quick inferences that form the basis of your understanding of my words remain hidden. You seem to just *hear* the meaning of what I say. What rises to the surface of your mind are the results of these mental processes. That is what I mean: the inferences themselves, the actual workings of our mind, remain unconscious. All that we are aware of are their products. And my access to your mind, when I listen to you speak, is not different in any fundamental way from my access to my own mind when I am aware of my own inner speech. The same sorts of interpretive processes still must take place.

Why, then, do we have the impression of direct access to our mind?

The idea that minds are transparent to themselves (that everyone has direct awareness of their own thoughts) is built into the structure of our "*mind reading*" or "*theory of mind*" faculty, I suggest. The assumption is a useful heuristic when interpreting the statements of others. If someone says to me, "*I want to help you,*" I must interpret whether the person is sincere, whether he is speaking literally or ironically, and so on; that is hard enough. If I also had to interpret whether he is interpreting his own mental state correctly, then that would make my task impossible. It is far simpler to assume that he knows his own mind (as, generally, he does). The illusion of immediacy has the advantage of enabling us to understand others with much greater speed and probably with little or no loss of reliability. If I had to figure out to what extent others are reliable interpreters of themselves, then that would make things much more complicated and slow. It would take a great deal more energy and interpretive work to understand the intentions and mental states of others. And then it is the same heuristic transparency-of-mind assumption that makes my own thoughts seem transparently available to me.

What is the empirical basis of your hypothesis?

There is a great deal of experimental evidence from normal subjects, especially of their readiness to falsely, but unknowingly, fabricate facts or memories to fill in for lost ones. Moreover, if introspection were fundamentally different from reading the minds of others, one would expect there to be disorders in which only one capacity was damaged but not the other. But that's not what we find. Autism spectrum disorders, for example, are not only associated with limited access to the thoughts of others but also with a restricted understanding of oneself. In patients with schizophrenia, the insight both into one's own mind and that of others is distorted. There seems to be only a single mind-reading mechanism on which we depend both internally and in our social relations.

What side effect does the illusion of immediacy have?

The price we pay is that we believe subjectively that we are possessed of far greater certainty about our attitudes than we actually have. We believe that if we are in mental state X, it is the same as being in that state. As soon as I *believe* I am hungry, I am. Once I *believe* I am happy, I am. But that is not really the case. It is a trick of the mind that makes us equate the act of thinking one has a thought with the thought itself.

What might be the alternative? What should we do about it, if only we could?

Well, in theory, we would have to distinguish between an experiential state itself on the one hand and our judgment or belief underlying this experience on the other hand. There are rare instances when we succeed in doing so: for example, when I feel nervous or irritated but suddenly realize that I am actually hungry and need to eat.

You mean that a more appropriate way of seeing it would be: "I think I'm angry, but maybe I'm not"?

That would be one way of saying it. It is astonishingly difficult to maintain this kind of distanced view of oneself. Even after many years of consciousness studies, I'm still not all that good at it.

Brain researchers put a lot of effort into figuring out the neural correlates of consciousness, the NCC. Will this endeavor ever be successful?

I think we already know a lot about how and where working memory is represented in the brain. Our philosophical concepts of what consciousness actually are much more informed by empirical work than they were even a few decades ago. Whether we can ever close the gap between subjective experiences and neurophysiological processes that produce them is still a matter of dispute.

Would you agree that we are much more unconscious than we think we are?

I would rather say that consciousness is not what we generally think it is. It is not direct awareness of our inner world of thoughts and judgments but a highly inferential process that only gives us the impression of immediacy.

Where does that leave us with our concept of freedom and responsibility?

We can still have free will and be responsible for our actions. Conscious and unconscious are not separate spheres; they operate in tandem. We are not simply puppets manipulated by our unconscious thoughts, because obviously, conscious reflection does have effects on our behavior. It interacts with and is fueled by implicit processes. In the end, being free means acting in accordance with one's own reasons—whether these are conscious or not.

Briefly Explained: Consciousness

Consciousness is generally understood to mean that an individual not only has an idea, recollection or perception but also *knows* that he or she has it. For perception, this knowledge encompasses both the experience of the outer world ("it's raining") and one's internal state *("I'm angry").* Experts do not know how human consciousness arises. Nevertheless, they generally agree on how to define various aspects of it. Thus, they distinguish "phenomenal consciousness" (the distinctive feeling when we perceive, for example, that an object is red) and "access consciousness" (when we can report on a mental state and use it in decision-making).

Important characteristics of consciousness include subjectivity (the sense that the mental event belongs to me), continuity (it appears unbroken) and intentionality (it is directed at an object). According to a popular scheme of consciousness known as Global Workspace Theory, a mental state or event is conscious if a person can bring it to mind to carry out such functions as decision-making or remembering, although how such accessing occurs is not precisely understood. Investigators assume that consciousness is not the product of a single region of the brain but of larger neural networks. Some theoreticians go so far as to posit that it is not even the product of an individual brain. For example, philosopher Alva Noë of the University of California, Berkeley, holds that consciousness is not the work of a single organ but is more like a dance: a pattern of meaning that emerges *between* brains.

Computers Determine States of Consciousness

A machine learning algorithm uses EEG traces to find a patient's odds of waking - Courtesy of Getty Images

Consciousness is a peculiar, even supernatural idea. From three pounds of flesh emerges an awareness of the body that houses it and the world around it. We all recognize consciousness when we see it, but what is it, really? And where does it go when it's gone? Neuroscience doesn't have the tools to answer these questions -if they're really possible to answer at all- but in a hospital, doctors need to be able to diagnose consciousness. They need to know if a patient with a brain injury is aware of himself or surroundings. This diagnosis is still mostly made with a simple bedside exam. Is the patient following commands? Is he gesturing or verbalizing purposefully, etc.?

For patients at the edge of consciousness -not lucid but not comatose either- defining the state of consciousness is difficult. Purposeless movements and sounds can look a lot like purposeful ones. Awareness comes and goes. In many, a high stakes diagnosis will be made. The patient is either in a minimally conscious state, where there's some likelihood of recovery, or the patient is given a diagnosis of unresponsive wakefulness syndrome, where the actions are deemed random and purposeless and there's little hope of recovery. Troublingly, these diagnoses are mixed up in as many as 40% of cases.

With a great deal at stake, a recent study in the journal Brain tries to give doctors a little help. The article details a machine learning algorithm that distinguishes unresponsive wakefulness syndrome from a minimally conscious state using EEG brainwave recordings. The algorithm, if put into use, would take some of the guesswork out of this diagnosis, and likely perform better than most human doctors. But diagnosing state-of-mind with an algorithm raises ethical concerns. How comfortable are we with turning over this kind of life-or-death diagnosis to a machine, especially since our handle on consciousness, as an idea, is so minimal?

Looking into the brain for traces of consciousness is not a new idea. For decades, researchers have been studying how brain scanning techniques like PET and fMRI could be used to study the edge of consciousness. In a landmark 2014 study, PET scans showed that brains could respond to cues in some patients given a (mis)diagnosis of unresponsive wakefulness syndrome. What's more is that the patients with an active PET scan were more likely to make a meaningful recovery. This finding argues that PET scans should be used if there's any doubt about a patient's state of consciousness. PET scans, though, aren't available in every hospital. They're also expensive, prone to artifact, and difficult to interpret. A more accessible alternative is electroencephalography or EEG, where electrical sensors are placed on the patient's scalp, picking up activity through the skull. EEG registers brain activity as waves when enough neurons fire in unison. In a healthy person, these waves undulate at predictable frequencies. After a brain injury, the pattern is less predictable.

In the new study, a group at Pitié-Salpêtrière Hospital in Paris took EEG recordings from 268 patients diagnosed with either unresponsive wakefulness syndrome or a minimally conscious state. The EEGs were recorded before and during a listening task designed to pick up on the conscious processing of sounds. Dozens of aspects of the data were fed into a machine learning algorithm called a DOC-Forest. The DOC-Forest performed relatively well at this complex task. Roughly 3 out of 4 cases were diagnosed properly. (Note: instead of accuracy, the authors use a better performance metric called AUC. AUC takes into account the rate of false positive classification, which has profound consequences here.) The authors also took care to push the DOC-Forest into real world scenarios. They introduced random noise into the data, simulating what differences in data collection procedures might look like. They considered different arrangements of sensors on the skull. They also used the algorithm on a different set of patients from a hospital in Liège, Belgium. In each case, the DOC-Forest performed well, with roughly the same performance measure.

From a certain perspective, this machine learning algorithm is a significant advance. EEG data is complex and contains multiple dimensions -time, frequency, testing condition, sensor locations, etc. Think pages and pages of squiggly waves on a computer screen. Typically, researchers would focus on a handful of easy-to-interpret aspects of the data, say the appearance of a specific brainwave during the listening task. This focus on interpretation excludes potentially important aspects of the data, though. Machine learning doesn't have this human bias toward interpretability and communicability. It just focuses on classifying the data correctly, which is all that's needed here. If put into practice, the DOC-Forest could be a helpful tool for an inexperienced neurologist. The DOC-Forest would run through the squiggly lines of EEG data and provide odds that the patient has some level of consciousness that the inexperienced doctor missed in his or her bedside tests. There's a circularity here, though. The algorithm is "*trained*" on cases that human neurologists diagnosed with bedside tests. While the group at Pitié-Salpêtrière could track patients for some time to minimize misdiagnoses, the algorithm just associates EEG signals with those -albeit more expert- bedside diagnoses. What, though, of a form of consciousness that's not revealed in any of these tests, EEG or otherwise? Keep in mind we don't really know where and how consciousness emerges. We don't have much sense of the forms conscious experience may take outside of the ones we experience for ourselves. One could argue our minimal understanding of the problem means that we shouldn't get the machines involved quite yet. On the other hand, it's not clear that we'll ever have satisfying answers to these questions. So, why not let a carefully designed tool, like the DOC-Forest, help make decisions within our current understanding of consciousness. There's no easy answer, but it's something that should probably be discussed as these tools push closer to everyday use.

New Clues to How the Brain Maps Time

The same brain cells that track location in space appear to also count beats in time. The research suggests that our thoughts may take place on a mental space-time canvas
– Illustration Courtesy of Quanta Magazine

Our brains have an extraordinary ability to monitor time. A driver can judge just how much time is left to run a yellow light; a dancer can keep a beat down to the millisecond. But exactly how the brain tracks time is still a mystery. Researchers have defined the brain areas involved in movement, memory, color vision and other functions, but not the ones that monitor time. Indeed, our neural timekeeper has proved so elusive that most scientists assume this mechanism is distributed throughout the brain, with different regions using different monitors to keep track of time according to their needs. Over the last few years, a handful of researchers have compiled growing evidence that the same cells that monitor an individual's location in space also mark the passage of time. This suggests that two brain regions -the hippocampus and the entorhinal cortex, both famous for their role in memory and navigation- can also act as a sort of timer.

In research published in November 2018, Howard Eichenbaum, a neuroscientist at Boston University, and collaborators showed that cells in rats that form the brain's internal GPS system, known as grid cells, are more malleable than had been anticipated. Typically, these cells act like a dead-reckoning system, with certain neurons firing when an animal is in a specific place. (The researchers who discovered this shared the Nobel Prize in 2014.) Eichenbaum found that when an animal is kept in place -such as when it runs on a treadmill- the cells keep track of both distance and time. The work suggests that the brain's sense of space and time are intertwined. The findings help to broaden our understanding of how the brain's memory and navigation systems work. Perhaps both grid cells and other GPS-like cells aren't tuned only to space but are capable of encoding any relevant property: time, smell or even taste. "*It probably points to a broad thing the hippocampus does*," said Loren Frank, a neuroscientist at the University of California, San Francisco, who studies memory and the hippocampus. "*It figures out the relevant axis for encoding experiences and then uses the cells to map those experiences*."

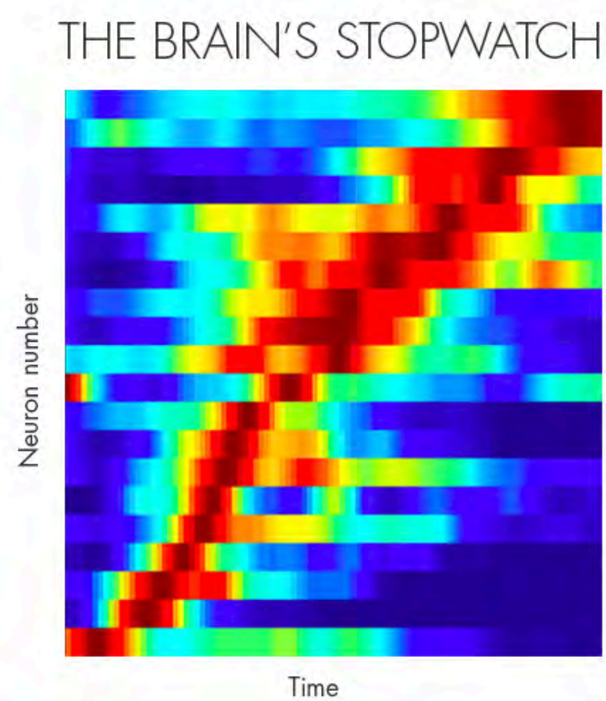

Different cells fired at different times during the 15-second stretch that rats ran on a wheel. Each row represents activity of an individual neuron (red indicates peak activity).
- Copyright Pastalkova, E et al. Science 2008

These maps in turn construct a framework for memory, providing an organizing system for our never-ending series of past experiences. "The hippocampus is this grand organizer of memories in space and time," Eichenbaum said. "*It provides a spatiotemporal framework onto which other events are applied*."

Time Tiles

To study how the hippocampus monitors time, scientists train rats to run on a wheel or tiny treadmill. This setup holds the animal's location and behavior constant, so that researchers can focus on the neural signals linked to time. (Rats are too fidgety to sit still, so running helps standardize their normally twitchy behavior.) Electrodes implanted deep in the brain record when different cells fire.

Howard Eichenbaum, a neuroscientist at Boston University, is exploring how parts of the brain that map space can also track time – Copyright Rohan Chitracar

In Eichenbaum's experiments, a rat runs on the treadmill for a set period -say, 15 seconds- and then gets a reward. As the animal repeats the cycle over and over, its brain learns to track that 15-second interval. Some neurons fire at one second, others at two seconds, and so forth, until the 15 seconds have elapsed. "*Each cell will fire at a different moment in time until they fill out the entire time interval*," Eichenbaum said. The code is so accurate that researchers can predict how long an animal has been on the treadmill just by observing which cells are active. Eichenbaum's team has also repeated the experiment, varying the treadmill's speed, to make sure the cells aren't simply marking distance. (Some of the cells do track distance, but some seem linked solely to time.) Although these neurons, dubbed "*time cells*," are clearly capable of marking time, it's still not clear how they do it. The cells behave rather like a stopwatch -the same pattern of neural activity repeats every time you start the clock. But they are more adaptable than a stopwatch. When researchers change the conditions of the experiment, for instance by extending the running duration from 15 to 30 seconds, cells in the hippocampus create a new firing pattern to span the new interval. It's like programming the stopwatch to follow a different time scale altogether.

Moreover, time cells rely on context; they only mark time when the animal is put into a situation in which time is what matters most. When other variables come into play, the same cells behave differently. Allow a rat to explore a new environment, for example, and these same cells will map themselves to space; a particular cell will fire every time the animal is in a specific location rather than doing so at a certain time.

The Brain's Space-Time Matrix

Eichenbaum's work dovetails with a 15-year trend in neuroscience research that suggests the hippocampus is more flexible than scientists expected. Researchers traditionally thought of it as a mapmaker -place-encoding cells were discovered 40 years ago- but growing evidence suggests it can encode other types of information as well. According to the newest picture, place cells can map not just space but other relevant variables. Time is one of them, but others are possible. For example, "*a wine taster might have a space of wine tastes and smells,*" Frank said. But many scientists still view the hippocampus as a largely spatial structure. According to their argument, neural circuitry evolved to keep track of location, and everything else is just recorded on top of it. "*The hippocampus provides a code that is fundamentally spatial in nature*," said Bruce McNaughton, a neuroscientist at the University of California, Irvine.

Eichenbaum's findings challenge this viewpoint, but they don't bury it. "*What's definitely clear is that place cells can represent information beyond place*," said David Foster, a neuroscientist at Johns Hopkins University. "*But what's less clear is whether they can code for the pure passage of time*." In the timed treadmill experiments, the rats appear to be doing something very much like counting. But are these cells marking the passage of time itself, or are they responding to something else that merely looks like time? "*We don't know the driving principle that tells cells to fire at a specific point, but I don't think it's time*," said Eva Pastalkova, a neuroscientist at the Howard Hughes Medical Institute's Janelia Research Campus in Ashburn, Va. "*It's not* *precise enough; they are not like ticking clocks*."

György Buzsáki, a neuroscientist at New York University's Neuroscience Institute whose lab did some of the first experiments exploring how the hippocampus tracks time, proposes that rather than monitoring time itself, these cells are doing something else -remembering a path through a maze or plotting the animal's next move. Both memories and future plans unfold in time, so time cells may simply reflect this mental activity. "*That's the number-one problem for me: Are there dedicated neurons in the brain doing nothing else but keeping track of time*?" Buzsáki said. "*Or do all neurons have functions that happen in sequential order, which for the experimenter can be translated into time*?" Buzsáki points out that it may not even make sense to think of hippocampal cells as independently coding for space or time. The human brain often considers time and distance interchangeably. "*If one asks how far New York is from LA, the answers you get vary: 3,000 miles, six hours by flight,"* he said. "*In older language, distances were typically given by time -the days it takes to go from one valley to another- since it was not distance but the number of sunsets that was easy to calculate*." For Buzsáki, the issue goes beyond neuroscience and reaches into physics. Physicists consider space-time as a cohesive, four-dimensional entity, a fabric upon which the objects and events of the universe are embedded. "*Neuroscience must converge back to the old problem of physics: Are there place and time cells? Or is there only a single time-space-continuum representation in the brain?"* Buzsáki said.

Memory Maps

Eichenbaum is less concerned with these abstract questions. His goal is to unpack the role time plays in forming memories. "*When you recall what you did this morning, you remember events in the order in which they occurred,"* he said. "*How does the hippocampus organize memories in time?"* People with damage to the hippocampus often can't make new memories -the famous patient H.M., who had a lobotomy to remove much of that part of the brain, introduced himself to his doctor over and over again each day. But these patients also have trouble remembering the sequence of words or objects presented in an experiment. "*How does the hippocampus support the ability to remember the temporal order of a sequence of events?"* Eichenbaum said. Eichenbaum envisions time cells as providing a timeline onto which sequential events are attached to represent an experience. If memories are a movie, he said, time cells are what puts the individual frames in order. His team is planning experiments that will intersperse time delays with different events, to see how time cells modify their code to remember the order in which the events occurred. "*I don't think the hippocampus is a clock,"* he said. "*But it's using a clock to map out when things happened in a memory to keep*"

The Strange Geometry of Thought

Published in the February 8th, 2019 issue of Nautil.us, Adithya Rajagopalan, a second-year graduate student in the department of neuroscience at Johns Hopkins University & Janelia Research Campus

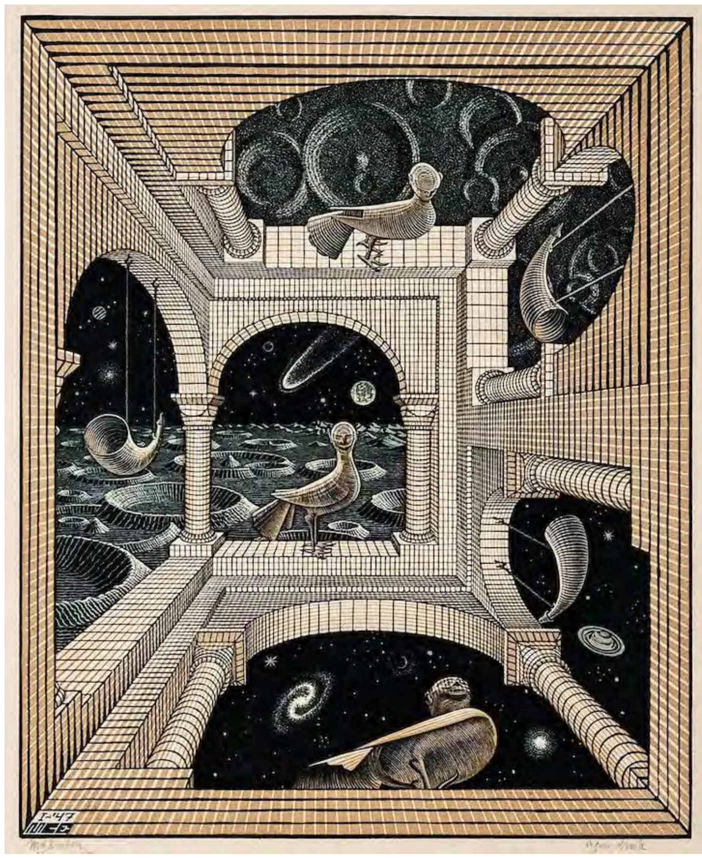

The brain may represent concepts in the same way that it represents space and your location, by using the same neural circuitry for the brain's "inner GPS"
– Copyright Sharon Mollerus / M.C. Escher / Flickr

In 2014, the Swedish philosopher and cognitive scientist Peter Gärdenfors went to Krakow, Poland, for a conference on the mind. He was to lecture at Jagiellonian University, courtesy of the Copernicus Center for Interdisciplinary Studies, on his theory of conceptual, or "*cognitive,*" spaces. Gärdenfors had been working on his idea of cognitive spaces, which explain how our brains represent concepts and objects, for decades. In his book, *Conceptual Spaces*, from 2000, he wrote, "*It has long been a* *common prejudice in cognitive science that the brain is either a Turing machine working with symbols or a connectionist system using neural networks*." In Krakow, Gärdenfors pushed against that prejudice. In his talk, "*The Geometry of Thinking,"* he suggested that humans can do things that today's powerful computers can't do—like learn language quickly and generalize from particulars with ease (to see, in other words, without much training, that lions and tigers are four-legged felines)—because we, unlike our computers, represent information in geometrical space.

In a 2018 *Science* paper, co-authored with Jacob Bellmund, Christian Doeller, and Edvard Moser—neuroscientists from the Max Planck Institute in Leipzig and the Kavli Institute in Trondheim—Gärdenfors, of the University of Lund, buttressed his idea with recent advances in brain science. He argued that the brain represents concepts in the same way that it represents space and your location, by using the same neural circuitry for the brain's *"inner GPS*". "*Cognitive spaces are a way of thinking about how our brain might organize our knowledge of the world*," Bellmund said. It's an approach that concerns not only geographical data, but also relationships between objects and experience. "*We were intrigued by evidence from many different groups that suggested that the principles of spatial coding in the hippocampus seem to be relevant beyond the realms of just spatial navigation*," Bellmund said. The hippocampus' place and grid cells, in other words, map not only physical space but conceptual space. It appears that our representation of objects and concepts is very tightly linked with our representation of space.

Work spanning decades has found that regions in the brain—the hippocampus and entorhinal cortex—act like a GPS. Their cells form a grid-like representation of the brain's surroundings and keep track of its location on it. Specifically, neurons in the entorhinal cortex activate at evenly distributed locations in space: If you drew lines between each location in the environment where these cells activate, you would end up sketching a triangular grid, or a hexagonal lattice. The activity of these aptly named "*grid*" cells contains information that another kind of cell uses to locate your body in a particular place. The explanation of how these '*place'* cells was stunning enough to award scientists John O'Keefe, May-Britt Moser, and Edvard Moser, the 2014 Nobel Prize in Physiology or Medicine as mentioned earlier. These cells activate only when you are in one particular location in space, or the grid, represented by your grid cells. Meanwhile, head-direction cells define which direction your head is pointing. Yet other cells indicate when you're at the border of your environment—a wall or cliff. Rodent models have elucidated the nature of the brain's spatial grids, but, with functional magnetic resonance imaging, they have also been validated in humans.

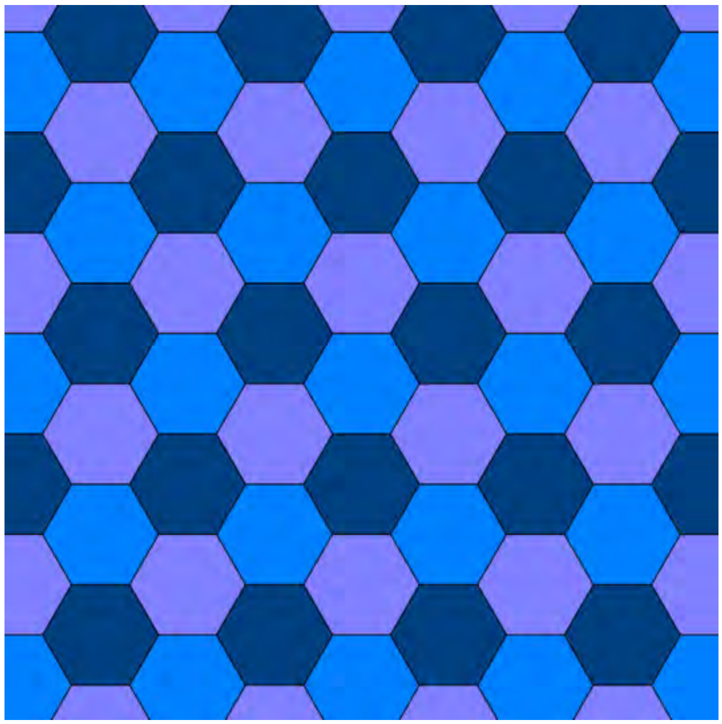

Hexagonal lattice – Copyright Wikicommons

Recent fMRI studies show that cognitive spaces reside in the hippocampal network—supporting the idea that these spaces lie at the heart of much subconscious processing. For example, subjects of a 2016 study—headed by neuroscientists at Oxford—were shown a video of a bird's neck and legs morph in size. Previously they had learned to associate a particular bird shape with a Christmas symbol, such as Santa or a Gingerbread man. The researchers discovered the subjects made the connections with a "*mental picture*" *t*hat could not be described spatially, on a two-dimensional map. Yet grid-cell responses in the fMRI data resembled what one would see if subjects were imagining themselves walking in a physical environment. This kind of mental processing might also apply to how we think about our family and friends. We might picture them "*based on their height, humor, or income, coding them as tall or short, humorous or humorless, or more or less wealthy*," Doeller said. And, depending on whichever of these dimensions matters in the moment, the brain would store one friend mentally closer to, or farther from, another friend.

But the usefulness of a cognitive space isn't just restricted to already familiar object comparisons. "*One of the ways these cognitive spaces can benefit our behavior is when we encounter something we have never seen before*," Bellmund said. "*Based on the features of the new object we can position it in our cognitive space. We can then use our old knowledge to infer how to behave in this novel situation*." Representing knowledge in this structured way allows us to make sense of how we should behave in new circumstances. Data also suggest that this region may represent information with different levels of abstraction. If you imagine moving through the hippocampus, from the top of the head toward the chin, you will find many different groups of place cells that completely map the entire environment but with different degrees of magnification. Put another way, moving through the hippocampus is like zooming in and out on your phone's map app. The area in space represented by a single place cell gets larger. Such size differences could be the basis for how humans can move between lower and higher levels of abstraction—from "*dog*" to "*pet*" to "*sentient being*," for example. In this cognitive space, more zoomed-out place cells would represent a relatively broad category consisting of many types, while zoomed-in place cells would be more narrow. Yet the mind is not just capable of conceptual abstraction but also flexibility—it can represent a wide range of concepts. To be able to do this, the regions of the brain involved need to be able to switch between concepts without any informational cross-contamination: It wouldn't be ideal if our concept for bird, for example, were affected by our concept for car. Rodent studies have shown that when animals move from one environment to another—from a blue-walled cage to a black-walled experiment room, for example—place-cell firing is unrelated between the environments. Researchers looked at where cells were active in one environment and compared it to where they were active in the other. If a cell fired in the corner of the blue cage as well as the black room, there might be some cross-contamination between environments. The researchers didn't see any such correlation in the place-cell activity. It appears that the hippocampus can represent two environments without confounding the two. This property of place cells could be useful for constructing cognitive spaces, where avoiding cross-contamination would be essential. "*By connecting all these previous discoveries,*" Bellmund said, "*we came to the assumption that the brain stores a mental map,* *regardless of whether we are thinking about a real space or the space between dimensions of our thoughts*."

Scientists still need to experimentally verify the link between the hippocampus and higher-order cognitive functions in humans. fMRI studies like the ones from the group in Oxford are, as yet, only suggestive. "*Although the coarse nature of the fMRI signal urges caution in making conclusions at the level of neuronal code*s," the researchers concluded, "*we have reported an unusually precise hexagonal modulation of the fMRI signal during nonspatial cognitio*n." It is also unknown whether place cells can actually represent objects at particular locations in a cognitive space. Revealing this in experiments with human subjects is hard, since they require very fine-resolution brain imaging. But recent advances in higher-resolution fMRI could possibly provide a solution. Bellmund pointed out that rodent research could also reveal the existence of cognitive spaces. A 2017 paper, for example, found that place cells in rats can form a map of sound frequencies. Different cells in the hippocampus respond to different frequencies of sound—forming a cognitive space of sound. What's more, studies in humans that have seen grid-like activity in the hippocampus have also seen this activity in other parts of the cortex. Therefore, it is highly likely that complicated, higher-order cognitive abilities arise from interactions between several parts of the brain.

Gärdenfors' theory highlights a fruitful path, not only for cognitive scientists, but for neurologists and machine-learning researchers. It is a kind of incomplete, generic sketch on a canvas that invites refinement and elaboration. Cognitive spaces are, as Gärdenfors and Bellmund put it, a "*domain-general format human thinking for,"* an "*overarching framework" that can help unravel the causes of neurodegenerative diseases, like Alzheimer's, and "to inform novel architectures in artificial intelligence."*

Whys of Seeing

Christ and the Disciples at Emmaus (1937) by Han Van Meegeren
- Photo courtesy Museum Boijmans Van Beuningen/Wikipedia

Many philosophical questions about the arts would benefit from some serious empirical research. One thinker who welcomed empirical findings was the art historian E H Gombrich (1909-2001), who was influenced by findings in experimental psychology showing that perception is a matter of inference rather than direct seeing. But all too often philosophers have relied on intuitions and hunches without seeking information about how people actually interact with works of art. If we want to understand the arts, it's time to take experimental psychology seriously. Today, experimental philosophers and philosophically inclined psychologists are designing experiments that can help to answer some of the big philosophical questions about the nature of art and how we experience it – questions that have puzzled people for centuries, such as: why do we prefer original works of art to forgeries? How do we decide what is good art? And does engaging with the arts make us better human beings?

Christ and the Disciples at Emmaus, believed to have been painted by Johannes Vermeer in the 17th century, hung in the Museum Boijmans Van Beuningen in Rotterdam for seven years; in 1937, it was admired by the Vermeer expert Abraham Bredius as '***the** masterpiece of Johannes Vermeer of Delft*'. But in 1945, Han van Meegeren confessed that he had forged this painting (and many others), and should thus be deemed as great an artist as Vermeer. But this did not happen. The same work formerly revered was now derided. There are two kinds of art forgeries: *invented forgeries* in the style of an established artist, and *copy forgeries*, which are reproductions of existing works. Most commonly, forgers such as van Meegeren produce invented forgeries. Copy forgeries are less common; these are more difficult to get away with since it is often known where the original resides. Moreover, because it is impossible to make a perfect copy by hand, one can always see (or hope to see) differences between the original and the copy, and use these differences to disparage the copy. The art critic Clive Bell in 1914 suggested that exact copies always lack life: the lines and forms in the original are caused by emotions in the artist's mind that are not present in the copier's mind. The philosopher Nelson Goodman in 1976 argued that, even if we can detect no physical difference between the original and the copy, just knowing that one painting is the original and the other is the copy tells us that there could be subtle differences that we cannot see now but that we might learn to see later. This knowledge shapes our aesthetic experience of what we believe to be a direct copy. The puzzle posed by forgery is this: why does our perception and evaluation of an artwork change simply by learning it is a forgery? After all, the work itself has not changed. Philosophers have taken two broad positions on this question.

According to the *formalist* position, when the original and the forgery are visually indistinguishable, they are not aesthetically different. For example, Monroe Beardsley in 1959 argued that we should form our aesthetic judgments only by attending to the perceptual properties of the picture before us, and not by considering when or how the work was made or who it was made by. So why did people change their evaluation of the Vermeer painting once van Meegeren confessed to being the artist? According to Alfred Lessing, writing in 1965, this response can be chalked up to social pressures: '*Considering a work of art aesthetically superior because it is genuine, or inferior because it is forged, has little or nothing to do with aesthetic judgment or criticism. It is rather a piece of snobbery.*' This view assumes that artworks have perceptual properties that are unaffected by our knowledge about the background of the work.

According to the *historicist* position, what we perceive in a work is influenced by what we know about it. Despite the original and the forgery being visually indistinguishable, they are aesthetically different precisely because of what the formalists deny is relevant – our beliefs about who made the work, when, and how. The German critic Walter Benjamin in the 1930s argued that our aesthetic response considers the object's history, '*its unique existence in a particular place*'. He believed that a forgery has a different history and thus lacks the '*aura*' of the original. The philosopher and critic Arthur Danto took a similar historicist position in 1964 when he asked us to consider why a Brillo box by Andy Warhol that is visually identical to a Brillo box in a supermarket is a work of art. To determine that the box in the museum is a work of art '*requires something the eye cannot descry – an atmosphere of artistic theory, a knowledge of the history of art: an artworld'*. Denis Dutton claimed in 2009 that we perceive a forgery to be aesthetically inferior to an original because we consider the kind of achievement the work represents – the artist's process – and a forgery represents a lesser kind of achievement than an original.

Psychologists have stepped into the fray to determine how much the label *'forgery'* affects our response to a work of art – and, if so, why. The first question is easier to answer than the second. Studies that just telling people that a work is a forgery (or even the less-charged term '*copy'*) causes them to rate that work lower on a host of aesthetic dimensions. Artworks labelled forgeries or copies are rated as less good, less beautiful, less awe-inspiring, less interesting, lower in monetary value, and even physically smaller than the same image shown to other respondents as an '*original*'. In addition, brain activation changes: while the visual areas of the brain didn't change in response to whether Rembrandt paintings were labelled '*authentic*' or '*copy*', the label '*authentic*' resulted in greater activation of the orbitofrontal cortex – an area associated with reward and monetary gain.

Clearly, people don't behave how the formalists thought that they should. What is causing their appreciation to be diminished? One possibility is that our sense of forgery's moral evil unconsciously influences our aesthetic response. Another is that our knowledge of forgery's worthlessness on the art market has the same kind of unconscious effect. But if we could strip forgery of its connection with deception and lack of monetary value, would it still be devalued? And, if so, can we demonstrate that the historicist position is correct?

With her research team, Ellen Winner, Professor of Psychology at Boston College, put this to the test by showing people two duplicate images of an unfamiliar art work side by side, telling them that the painting on the left was the first in a planned series of 10 identical works by a painter. Participants were then told one of three different stories about who made the work on the right: that it was by the artist, by the artist's assistant, or by a forger. For those told it was the artist's assistant, we specified that the assistant's copy had the artist's stamp on it, and that having a team of assistants was typical artistic practice (hence not fraudulent). The auction price of $53,000 was listed below all images (right and left) except for the forgery, which was listed at only $200. We asked people to rate the copy relative to the original on six dimensions:

> *Which one is more creative?*
> *Which one do you like more?*
> *Which one is more original?*
> *Which one is more beautiful?*
> *Which one is the better work of art?*
> *Which one is more likely to be influential?*

Responses fell into two categories: *broadly evaluative* (what formalists called aesthetic) – regarding beauty, goodness and liking; and *historical-evaluative* (what historicists called historical) – with reference to creativity, originality and influence. We reasoned that forgeries would always be the most devalued of the three kinds of copies because of their immorality and their lack of monetary worth. The artist's copy, however, is like a forgery without these two marks against it. Thus, our key comparison was between responses to the artist's versus the assistant's copy, relative to the original. We found that, on broadly evaluative dimensions, the artist's and the assistant's copies were rated identically – with no distinctions in beauty, liking or goodness. Thus, our participants behaved like formalists. Previous studies reporting lower beauty ratings for images labelled forgeries had presented works one at a time. But here, when the original and the forgery were presented simultaneously, people were forced to concede that there was no beauty difference. On historical-evaluative ratings, however, the story was different. People rated the assistant's copy as less creative, original and influential than the artist's copy – even though both works were copies, both signed by the artist, and both worth the same monetarily. People now behaved as historicists, consistent with Danto's position that visually identical Brillo boxes are not artistically identical. These findings tell us that, when moral and monetary considerations are ruled out, there is still something wrong with a forgery. It's not quite what Dutton thought, because while an original certainly represents a different kind of achievement from a forgery, there is really no difference in achievement between an artist's copy and an assistant's copy. Both are copies, after all. So, what is it that's wrong then? Ellen Winner submits that it's the aura that Benjamin talked about, which is dependent most critically on *who* made the work. Benjamin's idea of '*aura*' is consistent with what psychologists call essentialism – the view that certain special objects (e.g., my wedding ring, or my childhood teddy bear) gain their identity from their history, and have an underlying nature that cannot be directly observed, a view developed extensively by the psychologist Susan Gelman. This is why we reject perfect replicas of such objects: we want the original. We appear to treat works of art this way too – as if they contain the essence of the artist, or the artist's mind. We prefer the copy by the artist to the copy by the assistant because only the former contains that essence. This leads to the conclusion that just knowing that we are looking at a painting by Vermeer (even if it is a copy of a Vermeer by Vermeer) makes us feel like we are communing with Vermeer. Do we really want to find out that we were actually communing with van Meegeren? These findings predict that we will not respond well to what the future is bringing us: three-dimensional prints of paintings virtually indistinguishable from the originals, and works of art generated by computers. These works will not allow us to infer the mind of the human artist.

The American art critic Peter Schjeldahl put this well when he wrote in 2008:

> *The spectre of forgery chills the receptiveness – the will to believe – without which the experience of art cannot occur. Faith in authorship matters. We read*

> *the qualities of a work as the forthright decisions of a particular mind, wanting to let it commandeer our own minds, and we are disappointed when it doesn't.*

If we read into a work of art the artist's decisions, as Schjeldahl writes, then we are inferring a mind behind the work. Can we do this for abstract art? And, if so, can this help us distinguish art by great abstract expressionists from superficially similar art by children and animals? Tension between those who revere and those who deride abstract art can be seen even among the most highly regarded art historians. In *Art and Illusion* (2000), Gombrich focused on representational art as a great human achievement, and disparaged abstract art as a display of the artist's personality rather than skill. Contrast this to the attitude of the late American art historian Kirk Varnedoe, who was chief curator of painting and sculpture at the Museum of Modern Art from 1988 to 2001. In *Pictures of Nothing* (2006), Varnedoe responds explicitly to Gombrich's challenge, writing that abstract art is a signal human achievement created in a new language, and filled with symbolic meaning. The '*mind-boggling range of drips, stains, blobs, blocks, bricks, and blank canvases' seen in museums of modern art are not random spills*, he wrote. Rather, like all works of art, they are '*vessels of human intention'* and they '*generate meaning ahead of naming*'. They represent a series of deliberate choices by the artist, and involve invention and evoke meanings – for example, energy, space, depth, repetition, serenity, discord. Chimps, monkeys and elephants have all been given paints, brushes and paper on which to make marks. And their paintings, like those of preschoolers, bear a superficial resemblance to abstract expressionist paintings. Who hasn't heard someone deride abstract art as requiring no skill at all, with statements such as '*My kid could have done that!*'

We wanted to find out whether people see more than they think they do in abstract art – whether they can see the mind behind the work. Ellen Winner's team created pairs of images that looked eerily alike at first glance. Each pair consisted of a painting by a famous abstract expressionist whose works were found in at least one major art-history textbook (e.g., Mark Rothko, Hans Hofmann, Sam Francis, Cy Twombly, Franz Kline and others) and a painting either by a child or a nonhuman animal (chimp, gorilla, monkey or elephant). The question we asked was whether people would prefer, and judge as better, works by artists over works by children and animals. And, if so, on what basis?

We set up the study so that people first saw 10 pairs of paintings without any labels revealing who made them, followed by 20 more pairs with authorship information under each image: one labelled '*artist*' and the other labelled '*child', 'chimp', 'gorilla', 'monkey'* or '*elephant*'. Of these 20, half were randomly labelled correctly, and half incorrectly (thus a Hofmann painting might have been labelled '*child*'). If people can't distinguish between works by artists and by untrained children and animals, we should expect chance to play a part in their responses to the first 10 unlabeled pairs, which would mean choosing works by artists they liked more or saw as better only 50 per cent of the time. But this did not happen. For both the *like* and the *better* questions, participants selected the works by artists at an above-chance level. And when our respondents chose a work by an artist as *better* or *preferred*, they often explained their choice by referring to the mind behind the work, saying that it looked more thought-out, intentional and planned. Such mentalistic explanations were significantly more frequent following the choice of the artist's actual work (whatever the label).

Winner's team has repeated this experiment in different ways – for example, presenting the works one at a time rather than in pairs and asking which is by the artist *vs* the child or animal – and they found that, overall, people are correct about two-thirds of the time, a rate that is significantly above chance responding. Most importantly, when they asked people to rate each image on the dimension of perceived intentionality (without telling them that some were by children and animals), they found that those by artists received significantly higher intentionality ratings. This led them to conclude that people see more than they think they see in abstract art – they see the mind behind the work, and this is what leads them to distinguish between works by artists and by children or animals, and what leads them to classify artists' works as better.

Moreover, when people do make mistakes, they are operating based on perceived intentionality: those artists' works rated as low in intentionality get misidentified as having been done by children or animals; and those works by children and animals rated high in intentionality get misidentified as done by artists. What this shows is that we evaluate abstract art by inferring (rightly or wrongly) the mind behind the art. Just as we evaluate an original as better than a forgery because we infer the mind of the original master behind the original work. Another area in the philosophy of art that can be illuminated by the findings of experimental psychology is the question of whether studying the arts has the *'transfer'* effect of making us better people. It is indeed plausible to assume that reading particular works of literature about particular kinds of injustices might arouse in us feelings of empathy, not just for the characters, but also for real individuals who might be in the same kinds of situations as the fictional ones. This is a view that is often promulgated, but rarely with evidence.

Colson Whitehead, the author of *The Underground Railroad* (2016), a powerful novel about slavery, reported that a stranger told him: '*Your book made me a more empathetic person*.' The claim that the narrative arts make us more empathic seems highly plausible: we project ourselves into fictional characters and simulate what they are experiencing. Where better to step into the shoes of another than with works of literature, where we can meet a wide variety of people, all so different from ourselves? Referring to the novels of Charles Dickens and George Eliot, Martha Nussbaum wrote in *Cultivating Humanity* (1997): '*It is impossible to care about the characters and their wellbeing in the way the text invites, without having some very definite political and moral interests awakened in oneself*.' Not all would agree. The literary critic Harold Bloom wrote in 2000 that: '*The pleasures of reading indeed are selfish rather than social. You cannot directly improve anyone else's life by reading better or more deeply.*' The philosopher Gregory Currie speculates that when we expend empathy on fictional characters, our empathy for actual people is depleted. William James seems to have had the same idea when in 1892 he asked us to imagine '*the weeping of the Russian lady over the fictitious personages in the play, while her coachman is freezing to death on his seat outside*'. This is reminiscent of Kwame Anthony Appiah's description in 2008 of Catholic seminarians who have just heard a lecture on the Good Samaritan – and then rush right past someone in need when they realize they're late for the next lecture. After leaving the fictional world, have we paid our empathy dues?

Experimental psychologists have begun to seek evidence of fiction's power to induce empathy. In one study after reading a story about an injustice committed against an Arab-Muslim woman, participants were less likely to categorize angry ambiguous-race faces as Arab. But did this translate into kinder behavior? This was not examined. In another study, after reading an excerpt from *Harry Potter* about stigma, children reported more positive attitudes about immigrant children at their school. But note that this change of mind (or heart) occurred only for children who identified with Harry, and only after a discussion of the reading with their teacher – and this might have been the deciding factor. Meanwhile, after reading a story about a character behaving prosocially, participants were more willing to help the experimenter pick up accidentally dropped pens. But note that picking up dropped pens is very low-cost helping, and we have no idea how long such an inclination to help lasts. If we just go by the studies carried out so far, Winner would have to come down on the side of the skeptics. The evidence is there only for very near and very immediate transfer – prosocial stories written by experimenters make us think and act more prosocially (often quite similarly to how the characters thought and acted) immediately after reading and, in one case, also after discussing the story's themes. These near-transfer connections are not particularly generalizable to actual literature, since most literature does not carry a moral lesson. And remember that the Germany which led to Hitler was one of the most literate societies, reading Goethe and listening to Beethoven's '*Ode to Joy'* by night.

And yet, Ellen Winner is not ready to conclude that there might be no link between reading certain works of literature and becoming more empathic. I make here a plea for more and better research on the effects of literature on compassion. Great literature allows us to get inside the lives of people we would never meet in real life. Winner would wager that if we could do the study right, we could show that reading Dickens not only makes us feel what it is like to be a penniless, hungry and unjustly treated child, but also more likely to help children in such conditions too – if the opportunity to help presents itself. Similarly, what might be the effect of reading Victor Hugo's *Les Misérables* (1862) – especially the scene where Jean Valjean, just released from years of hard labour for stealing food to feed his sister's child, steals silverware from the bishop who has taken him in – and when Valjean is caught, the bishop gives him two candlesticks too, as a gift? Here we see the cruel effects of injustice and the potential life-giving results of kindness. Might this make a judge more likely to be lenient in sentencing certain kinds of crimes? Of course, there would be cases of people who don't react this way. No one would claim that there is a necessary link between reading Dickens and Hugo, and compassion for the poor and outrage at injustice. But reading Dickens and Hugo just might nudge some of us to be more compassionate, and that is all we would need to make the case that literature has the power to induce compassion.

Ellen Winner has discussed three philosophical questions about art that can be addressed empirically: why do we disparage once-revered artworks when they are outed as forgeries? How do we judge quality in abstract art? And does entering a narrative world make us more empathetic once we exit that world? Of course, not all philosophical claims about art can be studied empirically. No experiment could give conclusive answers to ontological questions such as *What is art?* or *What is beauty?* But the philosophical problems discussed here are inherently psychological ones. When a philosophical claim is about how the mind works, then that claim is not immune to psychological data, and it is incumbent upon philosophers to consider whether psychologists' answers are satisfactory.

To Understand Art, Think Biology

Just as cells are the building blocks of the human body, a painting's points, lines, colors, and tensions are the building blocks of its life.

Visitors look at an untitled painting by Jean-Michel Basquiat at the Christie's auction house in London in June 2013 - Copyright FRANK AUGSTEIN / AP

An artwork is a living organism. If you visually break down a work of art into its various components and systems, you will begin to understand how each of its elements functions and how those elements work together in harmony, just as you would if you were learning gross anatomy or dissecting a body. In this way, you can begin to see not just what an artwork looks like, but how it's structured, what its elements and systems do, how they interrelate, and how they contribute to the life of the artwork as a whole. You can begin to understand that just as cells are the building blocks of an organism's life, so too an artwork's elements are the building blocks of its life. Essential to the life of an organism, such as a human body, are the tensions among ligaments and muscles, the circulation of fluids, the strength and density of bone, the functions of organs, the elasticity and porousness of skin, the rhythms of breathing and heartbeat. Those interdependent elements of the body, if they are not purposeful, healthy, and working together, could become useless, if not dangerous, to the organism as a whole. So too an artwork's unique, interdependent elements (its points, lines, movements, shapes, forms, colors, structures, energies, tensions, light, and rhythms) must be present, healthy, functional, and purposefully fused -working together in harmony, subservient to the greater whole- in order for that work of art to have life.

Paul Klee suggested that a line is a point going for "*a walk … a point, shifting its position forward*." In representational works, we often experience a line as designating some particular thing or place: a strand of hair, the horizon, the contour of a form, such as the edge of a cheek or an apple. But a line can also represent abstract ideas about boundary, meeting place, energy, fusion, and fission. A line can be a spine or a vein. It can be lightning, or stress, or striving. When two forms meet and press against each other, they create a new line -a juncture, an offspring- out of the encounter and merging of their two separate boundaries: new life and energy generated out of interface. Each element in a work of art has the potential, through relationships and contact with other elements, to foster new forms and life -to convey a sense of growth and transformation; to convey that the artwork, though its forms might be literally unchanging, is not *fixed* but continually in motion, interaction, creation. These are the qualities and energies and actions we notice and therefore make happen in an artwork through the acts of looking and experience.

The more you begin to understand about what an artwork's elements are doing -and can do- the more you'll begin to realize what is possible in art. You'll begin to see how subtle pressures and spatial shifts operate and move you through an artwork, not just vertically and horizontally but also from front to back, in depth, and from inside to outside. And as you move through an artwork, letting your eyes hop and glide from form to form, you'll begin to pick up on the rhythms and melodies of the artwork. You'll begin to feel your own eyes dancing through the artwork's elements: moving faster here, slower there, pausing, and resting; traversing long, lyrical arcs and feeling the rat-a-tat effects -as in Piet Mondrian's *Broadway Boogie Woogie*, Byzantine mosaics, and the checkerboard patterns of abstract medieval manuscript pages- of staccato pulses. If you accept that these forces are at work, you'll begin to sense not just palpable forms in a picture, but also the palpable space in which those forms exist: a sense of distance and atmosphere and tension and air among forms; a sense that those spaces are not only open and navigable, but charged with light and energy, and that those spaces, like the forms themselves, seem to breathe and provide air, that they are believable, purposeful -alive.

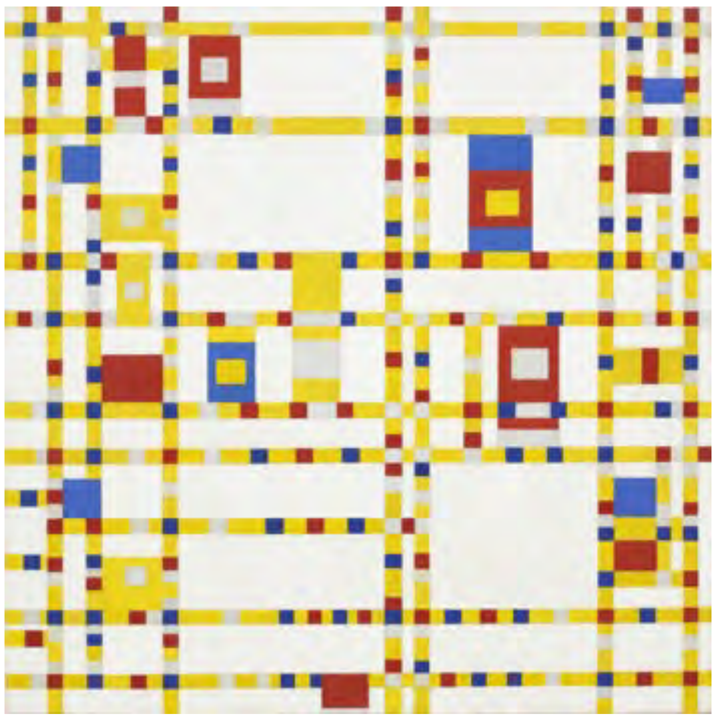

Courtesy of New York MOMA

Consider the work of Wassily Kandinsky (1866–1944), the Russian painter and professor who wrote a number of groundbreaking books. In Kandinsky's early landscape-based abstractions from 1911, line, color, and shape free themselves from acting as descriptive nouns into becoming active verbs, forces, and energies: Contours and colors mix and dissolve, suggesting stained glass and children's finger painting; a horse and rider are expressed as a lyrical bolt of lightning. By the time he paints the purely abstract masterpiece *Black Lines* (1913), all sense of the landscape has vanished: Color, shape, volume, space, and line are completely independent, free agents, and his spatial arabesques simultaneously remain on the surface of his flat picture and also transform that flat surface into a stretchy membrane, as if the painting were made out of rubber or taffy, or could be poked and twisted, or even turned inside out, like the surface of a balloon.

Courtesy of Guggenheim Museum. New York

Kandinsky gave his lines, forms, and shapes pictographic immediacy and energy. Sometimes, as in the later masterpieces *Blue World* (1934), *Striped* (1934), *Thirty* (1937), *Various Parts* (1940), *Sky Blue* (1940), and *Various Actions* (1941), his forms feel as if they are painted or incised hieroglyphics, while they also suggest unfamiliar creatures and microscopic organisms, as well as veins and jolts of electric current. They look like signs and symbols yet feel alive and in motion.

Or consider Hans Hofmann (1880–1966), a renowned German American abstract painter and an influential teacher of many of America's most celebrated mid-century artists, who came up with the English term *push and pull* to describe the dynamic interplay between flatness and depth in a painting or drawing. Even in an abstract painting made of colored rectangles, no two shapes of different sizes (even if they are the same hue) can really seem to exist in the exact same spatial location relative to each other. Larger forms will tend to advance, perhaps, and smaller forms will tend, perhaps, to recede. Or the opposite can happen. Other qualities besides size -such as brightness, intensity, opacity, translucency, density, warmth, and coolness, and especially location relative to other elements -all contribute to how we read where, exactly, a color or form is in relation to another color or form, and to how we experience the rhythmic pulse and speed and dynamic push and pull of those forms as they seem to shift within the spatial universe of the picture.

Courtesy of Guggenheim Museum, New York

We can experience the visceral dynamics of push and pull in Hofmann's abstract painting *The Gate* (1959-1960). Here, Hofmann's color shapes are heavy with opacity and impasto -thickened paint that feels as if it has been troweled onto the canvas. The speed of the colored shapes feels at first to be somewhat sluggish, as if forms are still congealing. Yet Hofmann gives those forms various speeds and levels of elasticity. Each form seems to be stepping forward and/or backward in space; the forms all seem to be competing, perhaps, for frontality. If we try to ascertain which of Hofmann's forms is actually the closest to us in space and which is the farthest away, we find we are muscling in and out through the picture space -moving back and forth where the painted forms move. Hofmann's *Gate* -pushing and pulling, exerting its will on us- feels alive. His notion of push and pull, which refers to more than spatial dynamics, gets at the heart of what living is: The forms are interacting and struggling within the universe of the picture, just as we interact and struggle within our own universe. We relate to and sense in an artwork its own tensions and movements and vitality, its will to live -its ability, for instance, to defy the downward pull of gravity or to strive vertically against the opposite tension of horizontal. We sense that an artwork's forms endeavor to move, to assert themselves, to change and grow -to move within, affect, and break free from the binding flat plane or the solid marble.

It is neither important nor necessary, usually, to name and identify the symbolic meanings of those dynamics, but it is important to feel them and to understand how they operate on and relate to one another in an artwork, how they build into and contribute to and fuse as an organic whole. It's important, I believe, to begin to ascertain what the elements of a work are doing and why the artist created them and put them there. It is in these ways that we move beyond the mere physical, or formal, aspects of an artwork's elements and get to something closer to the artwork's philosophy -its raison d'être.

The Brain That Remade Itself

Doctors removed one-sixth of this child's brain—and what was left did something incredible

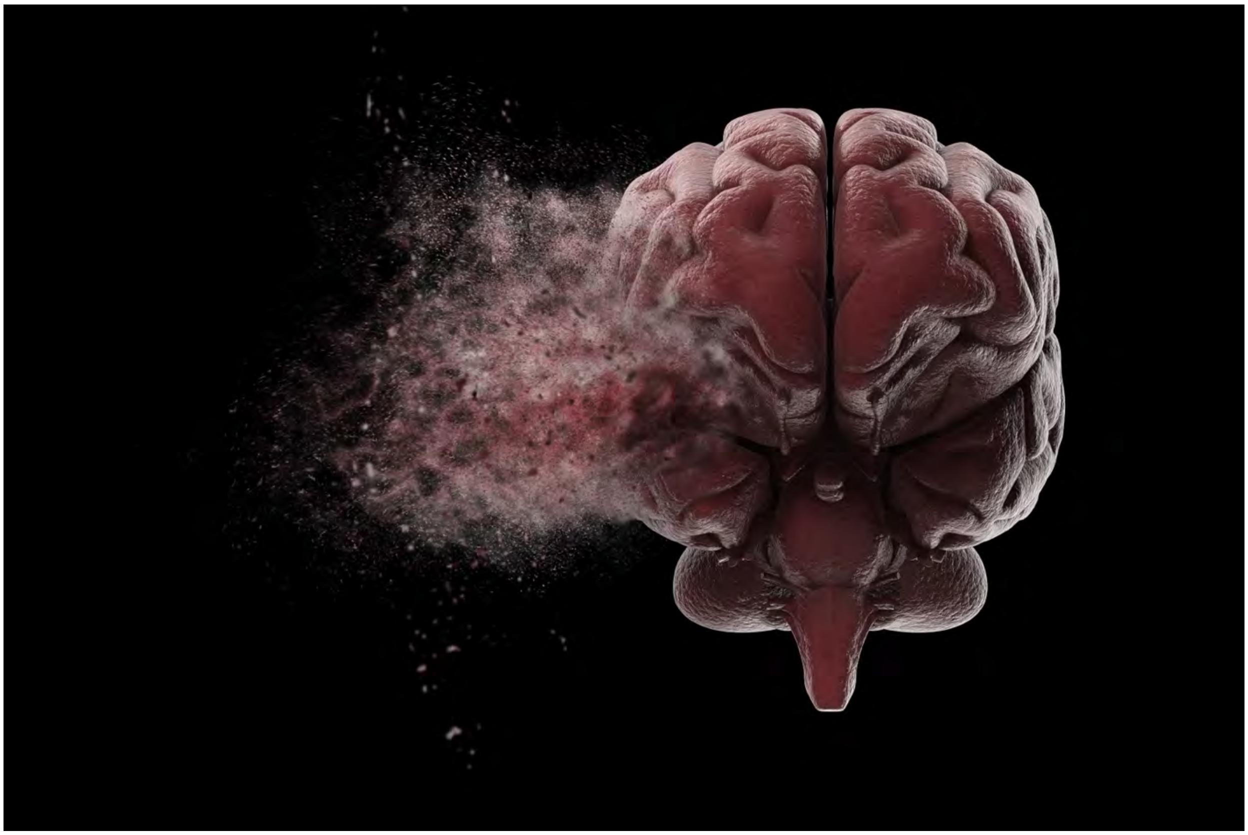

Credit: iLexx/Getty Images

I put my hand on a bishop and slide it several squares before moving it back. "Sho*uld I move a different piece instead?"* I wonder to myself.

"You have to move that piece if you've touched it," my opponent says, flashing a wry grin.

Fine. I move the bishop. It's becoming increasingly obvious to me now—I'm going to lose a game of chess to a 12-year-old.

My opponent is Tanner Collins, a seventh-grade student growing up in a Pittsburgh suburb. Besides playing chess, Collins likes building with Legos. One such set, a replica of Hogwarts Castle from the *Harry Potter* books, is displayed on a hutch in the dining room of his parents' house. He points out to me a critical flaw in the design: The back of the castle isn't closed off. "*If you turn it around*," he says, "*the whole side is open. That's dumb.*"

Tanner Collins, Credit: Courtesy of Nicole Collins

Though Collins is not dissimilar from many kids his age, there is something that makes him unlike most 12-year-olds in the United States, if not the world: He's missing one-sixth of his brain.

Collins was three months' shy of seven years old when surgeons sliced open his skull and removed a third of his brain's right hemisphere. For two years prior, a benign tumor had been growing in the back of his brain, eventually reaching the size of a golf ball. The tumor caused a series of disruptive seizures that gave him migraines and kept him from school. Medications did little to treat the problem and made

Collins drowsy. By the day of his surgery, Collins was experiencing daily seizures that were growing in severity. He would collapse and be incontinent and sometimes vomit, he says.

When neurologists told Collins' parents, Nicole and Carl, that they could excise the seizure-inducing areas of their son's brain, the couple agreed. "*His neurologist wasn't able to control his seizures no matter what medication she put him on,"* Nicole says. *"At that point, we were desperate… His quality of life was such that the benefits outweighed the risks*."

Surgeons cut out the entire right occipital lobe and half of the temporal lobe of Collins' brain. Those lobes are important for processing the information that passes through our eyes' optic nerves, allowing us to see. These regions are also critical for recognizing faces and objects and attaching corresponding names. There was no way of being sure whether Collins would ever see again, recognize his parents, or even develop normally after the surgery.

And then the miraculous happened: Despite the loss of more than 15 percent of his brain, Collins turned out to be fine.

"We're looking at the entire remapping of the function of one hemisphere onto the other."

The one exception is the loss of peripheral vision in his left eye. Though this means Collins will never legally be able to drive, he compensates for his blind spot by moving his head around, scanning a room to create a complete picture. "*It's not like it's blurred or it's just black there. It's, like, all blended,"* Collins tells me when I visit him at home in January*. "So, it's like a Bob Ross painting*."

Today, Collins is a critical puzzle piece in an ongoing study of how the human brain can change. That's because his brain has done something remarkable: The left side has assumed all the responsibilities and tasks of his now largely missing right side.

"*We're looking at the entire remapping of the function of one hemisphere onto the other,"* says Marlene Behrmann, a cognitive neuroscientist at Carnegie Mellon University who has been examining Collins' brain for more than five years.

What happened to Collins is a remarkable example of neuroplasticity: the ability of the brain to reorganize, create new connections, and even heal itself after injury.

Neuroplasticity allows the brain to strengthen or even recreate connections between brain cells—the pathways that help us learn a foreign language, for instance, or how to ride a bike.

The fact that the brain has a malleable capacity to change itself isn't new. What's less understood is how exactly the brain does it. That's where Behrmann's study of Collins comes in. Her research question is twofold: To what extent can the remaining structures of Collins' brain take over the functions of the part of his brain that was removed? And can science describe how the brain carries out these changes, all the way down to the cellular level? Previous neuroplasticity research has shed light on how the brain forms new neuronal connections with respect to memory, language, or learning abilities. (It's the basis for popular brain-training games meant to improve short-term memory.) But Behrmann's research is the first longitudinal study to look closely at what happens in the brain after the regions involved in visual processing are lost through surgery or damaged due to a traumatic brain injury.

"We know almost nothing about what happens in the visual system after this kind of surgery," she says. *"I think of this as kind of the tip of the iceberg."*

So far, Behrmann's findings are turning medical dogma on its head. They suggest that conducting brain surgeries in kids suffering seizures shouldn't be viewed as the last available option, as it was for Collins. The surgery he underwent, while successful roughly 70 percent of the time, is still uncommon, which means that many people with similar brain tumors may be suffering unnecessarily. And depending on what Behrmann discovers, we may learn more than we ever have before about the brain's capacity to bounce back.

The first-time Collins collapsed because of a seizure, he was four and being minded by a babysitter. Over time, his symptoms grew more varied and more severe. "*It's like my brain froze*," he says. *"I was really confused, and then I'd get really nauseous, throw up, and then I'd be kind of acting normal again*."

A daily ritual ensued: Collins would go to school, have a seizure, collapse, and go home. Still, despite the misery, the seizures were a blessing in disguise. They led to the discovery of the tumor slowly enveloping a piece of his brain.

"These are some of the most common tumors we see in children," says Christina Patterson, MD, a pediatric epilepsy neurologist and part of the medical team that prepared Collins for surgery at the UPMC Children's Hospital of Pittsburgh. "*Taking out the tumor is ultimately the cure*."

The deeper problem with pediatric tumors like the one Collins developed—beyond the nausea, headaches, and confusion that he experienced—is that the seizures they produce can damage the electrical networks of the brain.

"We know that the pediatric brain has plasticity, [and] that we're constantly creating new algorithms in the brain to live life," Patterson says*. "But when you have seizures on top of that, those disrupted electrical networks that are the seizures prevent any kind of meaningful remapping."*

Inside our brains are about 100 billion neurons. These neurons build thousands of connections with one another and communicate with their cellular brethren by converting electrical signals into chemical neurotransmitters, which are responsible for carrying information between the brain cells. As we master new skills, the brain's neurons form new connections and strengthen old ones that aided in learning that information. Instead of discrete regions carrying out specific tasks, the brain depends on groups of neural networks talking to each other across multiple regions. (Behrmann says a single neuron can communicate with 50,000 other cells.) If the network is damaged, the brain cells can't communicate effectively.

Picture a map of the United States that shows a phone company's LTE network crisscrossing the country, and you have a rough approximation of how the human brain operates. Surgery for Collins, in this case, was akin to repairing a downed cell tower.

Before Collins' surgery to remove the tumor, doctors opened up his head and placed electrodes on the surface of his brain and inside his visual cortex. For seven days, Collins lay in a hospital bed as the electrodes mapped his brain's electrical activity, creating what was essentially a schematic diagram showing doctors where the seizures were originating and which brain areas needed to be cut out.

Collins recognized his parents after the surgery, but he couldn't match their faces to their names. The problem resolved itself in a couple of days, but the episode left Nicole and Carl concerned: How was their son's brain going to function with a missing part?

Consider, for a moment, a page from a *Where's Waldo?* book. When your eye focuses on the crowded image, you're actually only receiving two types of feedback: the light that falls on the retina and the color of that light. "*That's all your eye can pick up,*" Behrmann says. *"Yet somehow, almost instantaneously, you get an interpretation of the scene."*

Patterson put the Collins family in touch with Behrmann, who studies how brain plasticity relates to vision at her lab at Carnegie Mellon. Collins was the ideal candidate for Behrmann's research. Children's brains are young and still developing and therefore have the most potential for neuroplastic change. Because Collins' tumor formed in the part of the brain crucial for visual processing, Behrmann could track his progress over time to determine whether there were any lingering deficits in his ability to interpret images. Because Collins was a child, his brain was also in a critical period of development where it builds the capacity to recognize faces, something that happens gradually and becomes more finely tuned throughout our teenage years.

As University of Toronto psychiatrist Norman Doidge notes in his 2007 book, *The Brain That Changes Itself*, the notion that there is a critical period of brain development is one of the most important discoveries in neuroplasticity—and one for which we have kittens to thank. In the 1960s, as Doidge recounts, scientists David Hubel and Torsten Wiesel mapped the visual cortex of kittens—much in the same way Collins' surgical team mapped his own brain—to learn how vision is processed. Then, in an admittedly grisly procedure, the scientists sewed shut the eyelid of one of the kittens in the study. Upon opening the eyelid, they found that the visual areas of the kitten's brain responsible for processing images from that eye didn't develop, leaving the kitten blind in that eye, even though nothing was biologically wrong with the eye. The researchers discovered that if kittens' brains were to develop normally, they had to be able to see the world around them between their third and eighth weeks of life.

But another discovery from the study proved even more important—and earned Hubel and Wiesel the Nobel Prize. "*The part of the kitten's brain that had been deprived of input from the shut eye did not remain idle,*" Doidge writes. "*It had begun to process visual input from the open eye, as though the brain didn't want to waste any* 'cortical real estate' *and had found a way to rewire itself.*"

In Collins' case, the question was whether the fully intact left hemisphere of his brain would pick up the functionality of the missing third of his brain, especially the task of facial recognition, which is typically carried out by the right hemisphere.

Collins' left brain not only looked and performed the way his left brain should; it also looked similar in scans to other kids' intact right brains.

Starting just before Collins was seven and continuing for three years, Behrmann administered a series of tests roughly every six months. In one challenge, he was shown photos of faces in intervals of roughly 30 seconds. If he remembered a face, he clicked a button. A similar test was administered using photos of houses, and if Collins saw the same photo back to back, he clicked a button. Each test occurred while he was inside a functional MRI machine, which allowed Behrmann to measure the flow of blood and oxygen to different regions of the brain. The more active an area of the brain, the more blood it draws.

Throughout these experiments, Behrmann compared Collins' brain function to a control group of kids his own age without brain abnormalities. The results, published last August in *Cell Reports*, were striking: His neurological function was "*absolutely normal*," with no subtle delays or deviations in development.

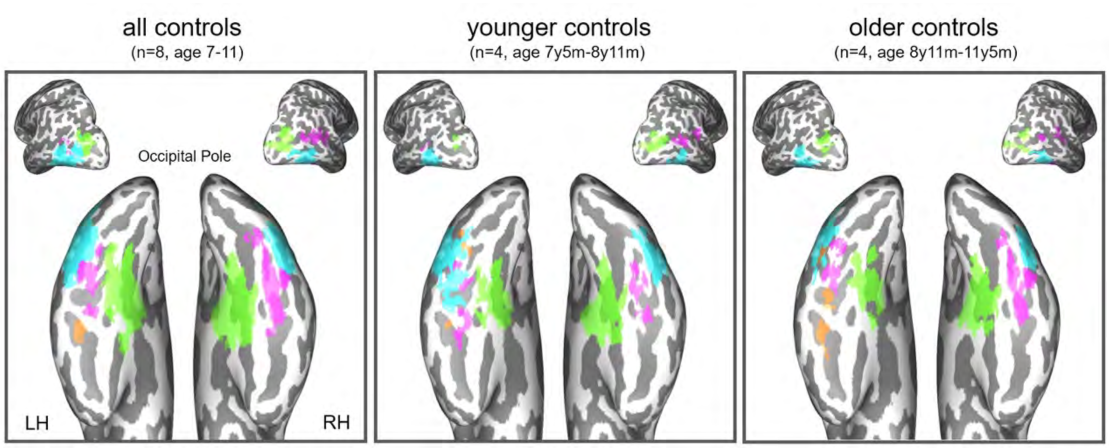

This figure shows the brain images of control groups of children around Tanner Collins' age. The images show what normal brain development looks like at a given age
- Credit: Liu et al., 2018, Cell Reports

Over coffee in the kitchen of her Pittsburgh home, Behrmann showed me successive scans of Collins' brain that told the tale. "*When he was eight, you can see the first glimmerings of face recognition in the brain*," she says. "*By the time he got to 10, you can see that his left hemisphere looks really like the right hemisphere of the controls.*"

In scans, Collins' left brain not only looked and performed the way his left brain should; it also looked similar in scans as other kids' intact right brains. That's because the functions of the visual cortex he lost by having one-third of his right brain removed—the ability to see objects and know what they are, and the ability to recognize faces—were subsumed by his left brain. Also, fascinating to Behrmann was how the left brain could accommodate two different skills: word recognition, which is the domain of the left brain, as well as facial recognition. Indeed, part of the surprise was that the left brain could keep doing what it normally does in addition to the newly added right-brain activity.

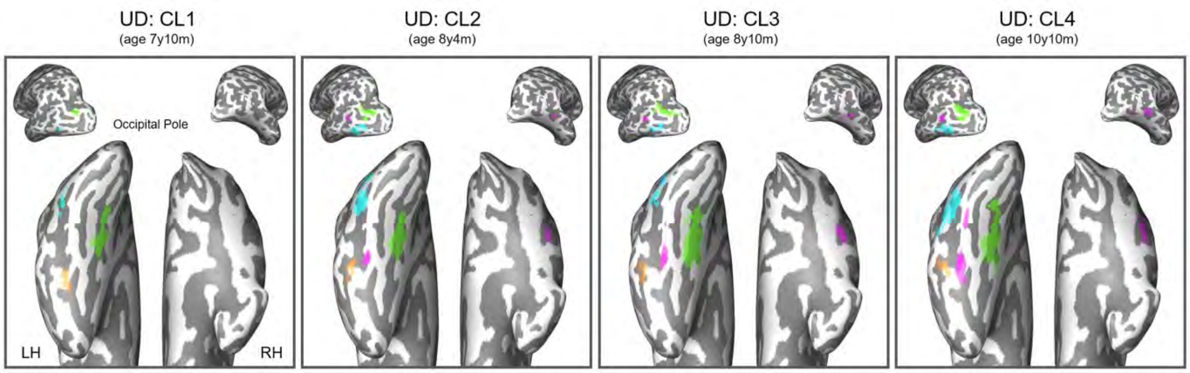

This figure is of Tanner Collins' brain. The images show that the left hemisphere is successfully assuming the right hemisphere responsibilities that we would typically see in children his same age. The only difference here is that those responsibilities have all shifted over to Collins' left brain - Credit: Liu et al., 2018, Cell Reports

In other words, Behrmann's work revealed that Collins' brain rewired itself, like the brain of the kitten that Hubel and Wiesel studied.

Just how the brain accomplishes this feat remains a central question. By analyzing brain scans using a neuroimaging technique known as diffusion tensor imaging, which shows how water travels along the brain's white-matter tracts, Behrmann has found initial glimmerings that the white matter of the brain—the electrical wiring that underlines communication between multiple neurological regions—actually changes. Areas of the brain that weren't connected before create new links, an example of neuroplasticity in action that may preserve brain functionality. But scientists still don't know what triggers the cells of the white matter to behave in this way.

"When Tanner is 20, I think we'll know a lot more about the overall wiring," Behrmann

says. *"The one thing that we will not know in humans, and I don't know how we will ever know it, are the changes that occur at the level of the cells themselves."*

Every three to six months, Collins returns to Behrmann's lab to undergo tests and be examined for any visual deficits. Behrmann hopes that following him over time will lead to more definitive answers, not only about how his visual system finally reorganizes itself but also the process by which it does so. "*We've got a long way to go, but the work, I think, is really exciting*," she says.

In a follow-up study Behrmann conducted with Collins and nine other children—all of whom are missing areas of either their left or right hemisphere—eight of them, including Collins, showed absolutely normal vision function. The two who did not are children whose brain damage from seizures was more severe prior to their surgeries.

This sort of insight is needed to gauge when to perform a brain surgery like the one Collins had. At what age should parents agree to remove a tumor that's causing epileptic seizures? Sometimes, resective surgery that removes brain tissue can make it difficult for a person to use and understand words; it can also, as it did in Collins' case, result in visual impairment.

"Once we have a better picture of exactly what happens after we remove large segments of the brain, we may be able to counsel families more effectively," says Taylor Abel, MD, a pediatric neurosurgeon who specializes in epilepsy surgery and arrived at the Children's Hospital of Pittsburgh last summer to begin collaborating with Behrmann. *"The goal should be to do whatever you can to stop the seizures and get off medications as early in your life as possible. The sooner you do that, the sooner you can return to a normal developmental trajectory."*

It may even be the case, Abel and Behrmann point out, that some of the reorganization that took place in Collins' brain started prior to his scheduled surgery. It's not something Behrmann can prove, since all the research conducted on Collins has taken place post-surgery.

"When you have an abnormality in your brain that's causing seizures, that abnormality can actually cause the brain to reorganize or start reorganizing before the surgery actually takes place," Abel says. *"But the other thing that sometimes happens is that the seizures affect the functions in the brain, and the brain doesn't reorganize."*

Behrmann says one of the fundamental goals of her research is to study a large enough population of children to determine if there are patterns of optimal recovery based on the age they had their surgery. Reorganization to the degree Collins has experienced is impossible for adults undergoing similar surgery, Behrmann says, as they lack the neuroplasticity seen in children.

For Nicole and Carl, the surgery was unequivocally the right decision. *"What was happening before the surgery was pretty awful,"* Nicole says. *"After surgery, the changes were only for the better. Yeah, he has his visual deficits. But everything else was for the better."*

In late 2017, a follow-up MRI at the Children's Hospital of Pittsburgh showed that Collins' tumor grew back. This time, though, it was the size of a pea. Two months later, in February 2018, surgeons opened his brain a second time. Collins says the prospect of a second surgery didn't bother him; he just wanted the pea-size tumor out of his head so he wouldn't have to worry about it. (The surgery went well, and he's still tumor-free.)

As we close in on minute 24 of our chess match, I move my king in the corner of the board, still certain of my impending doom. Collins scans his remaining white pieces and then takes a look at where his king sits.

"*Mate*," he says, looking up at me.

Checkmate for me, I realize, surprised by a victory I did not expect. Collins begins breaking down the moves he made, retracing some of his steps. It seems he forgot about a pawn of mine that was still on the board.

"I like losing," he says. *"Obviously, I like winning, too. But when you lose, you gain the knowledge."*

Even after losing a portion of his brain, Collins is still learning. His brain is still growing, still adapting—and, even if it's not readily apparent, still changing.

Foundations of Neural Networks

Neural networks can be as unpredictable as they are powerful. Now mathematicians are beginning to reveal how a neural network's form will influence its function.

Credit: Koma Zhang for Quanta Magazine

When we design a skyscraper, we expect it will perform to specification: that the tower will support so much weight and be able to withstand an earthquake of a certain strength.

But with one of the most important technologies of the modern world, we're effectively building blind. We play with different designs, tinker with different setups, but until we take it out for a test run, we don't really know what it can do or where it will fail. This technology is the neural network, which underpins today's most advanced artificial intelligence systems. Increasingly, neural networks are moving into the core areas of society: They determine what we learn of the world through our social media feeds, they help doctors diagnose illnesses, and they even influence whether a person convicted of a crime will spend time in jail.

Yet "*the best approximation to what we know is that we know almost nothing about how neural networks actually work and what a really insightful theory would be*," said Boris Hanin, a mathematician at Texas A&M University and a visiting scientist at

Facebook AI Research who studies neural networks. He likens the situation to the development of another revolutionary technology: the steam engine. At first, steam engines weren't good for much more than pumping water. Then they powered trains, which is maybe the level of sophistication neural networks have reached. Then scientists and mathematicians developed a theory of thermodynamics, which let them understand exactly what was going on inside engines of any kind. Eventually, that knowledge took us to the moon. *"First you had great engineering, and you had some great trains, then you needed some theoretical understanding to go to rocket ships,"* Hanin said.

Within the sprawling community of neural network development, there is a small group of mathematically minded researchers who are trying to build a theory of neural networks — one that would explain how they work and guarantee that if you construct a neural network in a prescribed manner, it will be able to perform certain tasks. This work is still in its very early stages, but in the last year researchers have produced several papers which elaborate the relationship between form and function in neural networks. The work takes neural networks all the way down to their foundations. It shows that long before you can certify that neural networks can drive cars, you need to prove that they can multiply.

The Best Brain Recipe

Neural networks aim to mimic the human brain — and one way to think about the brain is that it works by accreting smaller abstractions into larger ones. Complexity of thought, in this view, is then measured by the range of smaller abstractions you can draw on, and the number of times you can combine lower-level abstractions into higher-level abstractions — like the way we learn to distinguish dogs from birds. "*For a human, if you're learning how to recognize a dog you'd learn to recognize four legs, fluffy*," said Maithra Raghu, a doctoral student in computer science at Cornell University and a member of Google Brain. "*Ideally we'd like our neural networks to do the same kinds of things*."

Maithra Raghu, a member of Google Brain, has been identifying principles that explain how neural networks operate - Credit Arun Chaganty

Abstraction comes naturally to the human brain. Neural networks have to work for it. As with the brain, neural networks are made of building blocks called *"neurons*" that are connected in various ways. (The neurons in a neural network are inspired by neurons in the brain but do not imitate them directly.) Each neuron might represent an attribute, or a combination of attributes, that the network considers at each level of abstraction.

When joining these neurons together, engineers have many choices to make. They must decide how many layers of neurons the network should have (or how "deep" it should be). Consider, for example, a neural network with the task of recognizing objects in images. The image enters the system at the first layer. At the next layer, the network might have neurons that simply detect edges in the image. The next layer combines lines to identify curves in the image. Then the next layer combines curves into shapes and textures, and the final layer processes shapes and textures to reach a conclusion about what it's looking at: woolly mammoth!

"The idea is that each layer combines several aspects of the previous layer. A circle is curves in many different places, a curve is lines in many different places," said David Rolnik, a mathematician at the University of Pennsylvania. Engineers also must decide the "*width*" of each layer, which corresponds to the number of different features the network is considering at each level of abstraction. In the case of image

recognition, the width of the layers would be the number of types of lines, curves or shapes it considers at each level. Beyond the depth and width of a network, there are also choices about how to connect neurons within layers and between layers, and how much weight to give each connection.

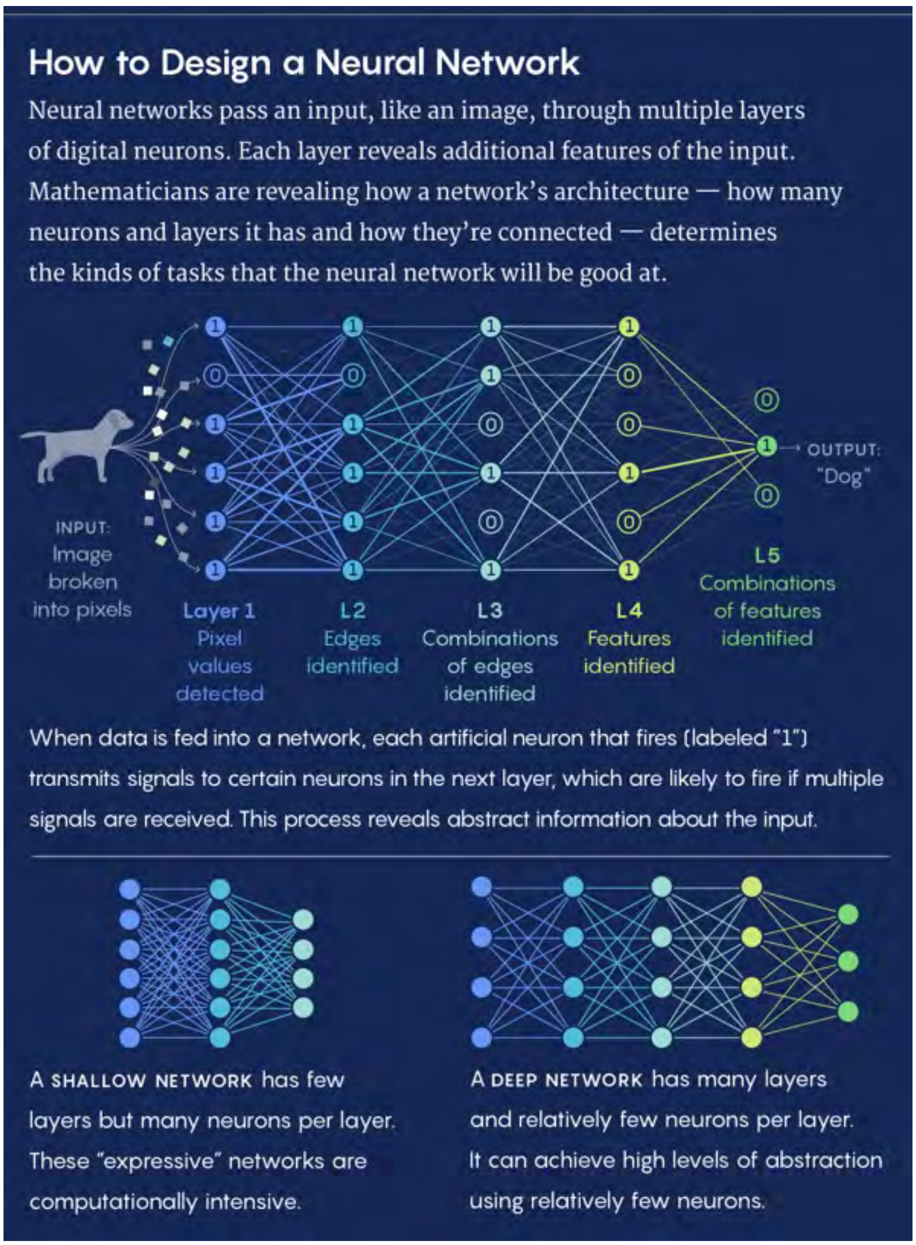

Copyright Lucy Reading-Ikkanda/Quanta Magazine

So, if you have a specific task in mind, how do you know which neural network architecture will accomplish it best? There are some broad rules of thumb. For image-related tasks, engineers typically use "*convolutional*" neural networks, which feature the same pattern of connections between layers repeated over and over. For natural language processing — like speech recognition, or language generation — engineers have found that "*recurrent*" neural networks seem to work best. In these, neurons can be connected to non-adjacent layers.

Beyond those general guidelines, however, engineers largely must rely on experimental evidence: They run 1,000 different neural networks and simply observe which one gets the job done. "*These choices are often made by trial and error in practice,*" Hanin said. "*That's sort of a tough [way to do it] because there are infinitely many choices and one really doesn't know what's the b*est."

A better approach would involve a little less trial and error and a little more upfront understanding of what a given neural network architecture gets you. A few papers published recently have moved the field in that direction. "*This work tries to develop, as it were, a cookbook for designing the right neural network. If you know what it is that you want to achieve out of the network, then here is the recipe for that network,*" Rolnick said.

To Rope a Red Sheep

One of the earliest important theoretical guarantees about neural network architecture came three decades ago. In 1989, computer scientists proved that if a neural network has only a single computational layer, but you allow that one layer to have an unlimited number of neurons, with unlimited connections between them, the network will be capable of performing any task you might ask of it. It was a sweeping statement that turned out to be fairly intuitive and not so useful. It's like saying that if you can identify an unlimited number of lines in an image, you can distinguish between all objects using just one layer. That may be true in principle, but good luck implementing it in practice.

Researchers today describe such wide, flat networks as "expressive," meaning that they're capable in theory of capturing a richer set of connections between possible inputs (such as an image) and outputs (such as descriptions of the image). Yet these networks are extremely difficult to train, meaning it's almost impossible to teach them how to actually produce those outputs. They're also more computationally intensive than any computer can handle.

Boris Hanin, a mathematician at Texas A&M University, has studied the tradeoff between depth and width in neural networks - Copyright Intel AI One Tree Studio

More recently, researchers have been trying to understand how far they can push neural networks in the other direction — by making them narrower (with fewer neurons per layer) and deeper (with more layers overall). So maybe you only need to pick out 100 different lines, but with connections for turning those 100 lines into 50 curves, which you can combine into 10 different shapes, which give you all the building blocks you need to recognize most objects.

In a paper completed last year, Rolnick and Max Tegmark of the Massachusetts Institute of Technology proved that by increasing depth and decreasing width, you can perform the same functions with exponentially fewer neurons. They showed that if the situation you're modeling has 100 input variables, you can get the same reliability using either 2^{100} neurons in one layer or just 2^{10} neurons spread over two layers. They found that there is power in taking small pieces and combining them at greater levels of abstraction instead of attempting to capture all levels of abstraction at once. "*The notion of depth in a neural network is linked to the idea that you can express something complicated by doing many simple things in sequence*," Rolnick said.

"*It's like an assembly line*."

Rolnick and Tegmark proved the utility of depth by asking neural networks to perform a simple task: multiplying polynomial functions. (These are just equations that feature variables raised to natural-number exponents, for example $y = x^3 + 1$.) They trained the networks by showing them examples of equations and their products. Then they asked the networks to compute the products of equations they hadn't seen before. Deeper neural networks learned the task with far fewer neurons than shallower ones. And while multiplication isn't a task that's going to set the world on fire, Rolnick says the paper made an important point: "*If a shallow network can't even do multiplication then we shouldn't trust it with anything else*."

David Rolnick, a mathematician at the University of Pennsylvania, proved that increasing a network's depth allowed a network to accomplish tasks with exponentially fewer neurons - Credit Stephanie Ku

Other researchers have been probing the minimum amount of width needed. At the end of September, Jesse Johnson, formerly a mathematician at Oklahoma State University and now a researcher with the pharmaceutical company Sanofi, proved that at a certain point, no amount of depth can compensate for a lack of width. To get a sense of his result, imagine sheep in a field, except these are punk-rock sheep: Their wool has been dyed one of several colors. The task for your neural network is to draw a border around all sheep of the same color. In spirit, this task is similar to image classification: the network has a collection of images (which it represents as points in higher-dimensional space), and it needs to group together similar ones.

Johnson proved that a neural network will fail at this task when the width of the layers is less than or equal to the number of inputs. So, for our sheep, each can be described with two inputs: an *x* and a *y* coordinate to specify its position in the field. The neural network then labels each sheep with a color and draws a border around sheep of the same color. In this case, you will need three or more neurons per layer to solve the problem. More specifically, Johnson showed that if the width-to-variable ratio is off, the neural network won't be able to draw closed loops — the kind of loops the network would need to draw if, say, all the red sheep were clustered together in the middle of the pasture. "*If none of the layers are thicker than the number of input dimensions, there are certain shapes the function will never be able to create, no matter how many layers you add*," Johnson said. Papers like Johnson's are beginning to build the rudiments of a theory of neural networks. At the moment, researchers can make only very basic claims about the relationship between architecture and function — and those claims are in small proportion to the number of tasks neural networks are taking on. So, while the theory of neural networks isn't going to change the way systems are built anytime soon, the blueprints are being drafted for a new theory of how computers learn — one that's poised to take humanity on a ride with even greater repercussions than a trip to the moon.

Philosophy can make the previously unthinkable thinkable

Detail from Woman at a Window (1822) by Caspar David Friedrich
- Courtesy Alte Nationalgalerie, Berlin

In the mid-1990s, Joseph Overton, a researcher at the US think tank the Mackinac Center for Public Policy, proposed the idea of a '*window*' of socially acceptable policies within any given domain. This came to be known as the Overton window of political possibilities. The job of think tanks, Overton proposed, was not directly to advocate particular policies, but to shift the window of possibilities so that previously unthinkable policy ideas – those shocking to the sensibilities of the time – become mainstream and part of the debate. Overton's insight was that there's little point advocating policies that are publicly unacceptable, since (almost) no politician will support them. Efforts are better spent, he argued, in shifting the debate so that such policies seem less radical and become more likely to receive support from sympathetic politicians. For instance, working to increase awareness of climate change might make future proposals to restrict the use of diesel cars more palatable, and ultimately more effective, than directly lobbying for a ban on such vehicles. Overton was concerned with the activities of think tanks, but philosophers and practical ethicists might gain something from considering the Overton window. By its nature, practical ethics typically addresses controversial, politically sensitive topics. It is the job of philosophers to engage in 'conceptual hygiene' or, as the late British philosopher Mary Midgley described it, '*philosophical plumbing*': clarifying and streamlining, diagnosing unjustified assertions and pointing out circularities.

Hence, philosophers can be eager to apply their skills to new subjects. This can provoke frustration from those embedded within a particular subject. Sometimes, this is deserved: philosophers can be naive in contributing their thoughts to complex areas with which they lack the kind of familiarity that requires time and immersion. But such an outside perspective can also be useful. Although such contributions will rarely get everything right, the standard is too demanding in areas of great division and debate (such as practical ethics). Instead, we should expect philosophers to offer a counterpoint to received wisdom, established norms and doctrinal prejudice.

Ethicists, at least within their academic work, are encouraged to be skeptical of intuition and the naturalistic fallacy (the idea that values can be derived simply from facts). Philosophers are also familiar with tools such as thought experiments: hypothetical and contrived descriptions of events that can be useful for clarifying particular intuitions or the implications of a philosophical claim. These two factors make it unsurprising that philosophers often publicly adopt positions that are unintuitive and outside mainstream thought, and that they might not personally endorse. This can serve to shift, and perhaps widen, the Overton window. Is this a good thing? Sometimes philosophers argue for conclusions far outside the domain of '*respectable*' positions; conclusions that could be hijacked by those with intolerant, racist, sexist or fundamentalist beliefs to support their stance. It is understandable that those who are threatened by such beliefs want any argument that might conceivably support them to be absent from the debate, off the table, and ignored.

However, the freedom to test the limits of argumentation and intuition is vital to philosophical practice. There are sufficient and familiar examples of historical orthodoxies that have been overturned – women's right to vote; the abolition of slavery; the decriminalization of same-sex relationships – to establish that strength and pervasiveness of a belief indicate neither truth nor immutability. It can be tedious to repeatedly debate women's role in the workforce, abortion, animals' capacity to feel pain and so on, but to silence discussion would be far worse. Genuine attempts to resolve difficult ethical dilemmas must recognize that understanding develops by getting things wrong and having this pointed out. Most (arguably, all) science fails to describe or predict how the world works with perfect accuracy. But as a collective enterprise, it can identify errors and gradually approximate '*truth*'. Ethical truths are less easy to come by, and a different methodology is required in seeking out satisfactory approximations. But part of this model requires allowing plenty of room to get things wrong.

It is unfortunate but true that bad ideas are sometimes undermined by bad reasoning, and also that sometimes those who espouse offensive and largely false views can say true things. Consider the 'born this way' argument, which endorses the flawed assumption that a genetic basis for homosexuality indicates the permissibility of same-sex relationships. While this might win over some individuals, it could cause problems down the line if it turns out that homosexuality isn't genetically determined. Debates relating to the 'culture wars' on college campuses have attracted many *ad hominem* criticisms that set out to discredit the authors' position by pointing to the fact that they fit a certain demographic (white, middle-class, male) or share some view with a villainous figure, and thus are not fit to contribute. The point of philosophy is to identify such illegitimate moves, and to keep the argument on topic; sometimes, this requires coming to the defense of bad ideas or villainous characters. Participation in this process can be daunting. Defending an unpopular position can make one a target both for well-directed, thoughtful criticisms, and for emotional, sweeping attacks. Controversial positions on contentious topics attract far more scrutiny than abstract philosophical contributions to niche subjects. This means that, in effect, the former are required to be more rigorous than the latter, and to foresee and head off more potential misappropriations, misinterpretations and misunderstandings – all while contributing to an interdisciplinary area, which requires some understanding not only of philosophical theory but perhaps also medicine, law, natural and social science, politics and various other disciplines.

This can be challenging, though one does not mean to be an apologist for thoughtless, sensationalist provocation and controversy-courting, whether delivered by philosophers or others. We should see one important social function of practical ethicists as widening the Overton window and pushing the public and political debate towards reasoned deliberation and respectful disagreement. Widening the Overton window can yield opportunities for ideas that many find offensive, and straightforwardly mistaken, as well as for ideas that are well-defended and reasonable. It is understandable that those with deep personal involvement in these debates often want to narrow the window and push it in the direction of those views they find unthreatening. But philosophers have a professional duty, as conceptual plumbers, to keep the whole system in good working order. This depends upon philosophical contributors upholding the disciplinary standards of academic rigor and intellectual honesty that are essential to ethical reflection, and trusting that this will gradually, collectively lead us in the right direction.

Can Intelligence Buy You Happiness?

New research suggests that IQ leads to greater well-being by enabling one to acquire the financial and educational means necessary to live a better life.

Courtesy of Getty Images

In his classic 1923 essay, "*Intelligence as the Tests Test it*, Edwin Boring wrote "*Intelligence is what the tests test.*" Almost a century of research later, we know that this definition is far too narrow. As long as a test is sufficiently cognitively complex and taps into enough diverse content, you can get a rough snapshot of a person's general cognitive ability -and general cognitive ability predicts a wide range of important outcomes in life, including academic achievement, occupational performance, health and longevity.

But what about happiness? Prior studies have been mixed about this, with some studies showing no relationship between individual IQ and happiness, and other studies showing that those in the lowest IQ range report the lowest levels of happiness compared to those in the highest IQ group. In one study, however, the unhappiness of the lowest IQ range was reduced by 50 percent once income and mental health issues were taken into account. The authors concluded that "*interventions that target modifiable variables such as income (e.g., through enhancing education and employment opportunities) and neurotic symptoms (e.g., through better detection of mental health problems) may improve levels of happiness in the lower IQ groups.*" One major limitation of these prior studies, however, is that they all rely on a single measure of happiness, notably life satisfaction. Modern day researchers now have measures to assess a much wider array of indicators of well-being, including autonomy, personal growth, positive relationships, self-acceptance, mastery, and purpose and meaning in life. Enter a new study conducted by Ana Dimitrijevic and colleagues, in which they attempted to assess the relationship between multiple indicators of intelligence and multiple indicators of well-being. They relied on the following definition of intelligence: "*the ability to understand complex ideas, to adapt effectively to the environment, to learn from experience, to engage in various forms of reasoning, and to overcome obstacles by taking thought*." This definition covers several more specific notions of intelligence, such as emotional intelligence. The researchers administered a battery of intelligence and well-being measures to 288 adults employed within various departments of a large dairy production company in Belgrade. What did they find?

Intelligence and Well-Being

The researchers found that both IQ and emotional intelligence were independently correlated with well-being. * IQ was positively correlated with personal relationships, self-acceptance, personal growth, mastery and purpose in life. Emotional intelligence was correlated with the same well-being measures but was additionally related to a sense of autonomy in life.

Zooming in on the IQ test, the most predictive subscale for well-being was a measure of *non-verbal fluid reasoning*, which requires pattern detection and abstract reasoning (constructing generalizable principles from minimal information). Some people argue that this form of reasoning is strongly related to general intelligence.

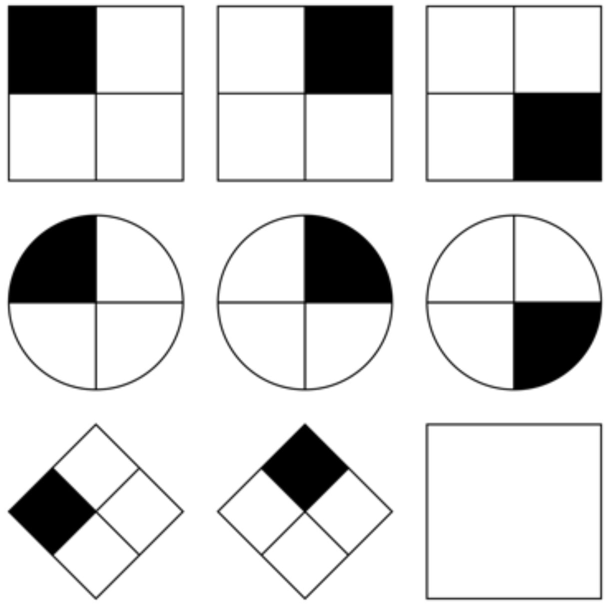

Credit: Life of Riley. Wikimedia (CC BY-SA 3.0)

Once socioeconomic status (SES) was taken into account (reflecting higher education and income), however, *there was no relationship between IQ and well-being*. According to the researchers, this suggests that IQ leads "*to greater contentment with oneself and life primarily by enabling one to acquire the social status and financial means which ensure better opportunities and quality of life*." Of course, this does *not* mean that IQ is simply a measure of SES; IQ was positively correlated with well-being. However, it does suggest that the extent to which IQ is related to happiness depends to a large extent on the opportunities (e.g., financial, educational) you must utilize your IQ.

What about emotional intelligence? The emotional intelligence tests that were most predictive of well-being were the two higher, more "*strategic*" branches - *Understanding* and *Managing Emotions*. The person who scores higher in these facets of emotional intelligence are better able to comprehend the emotional signals coming from others, and to regulate and manage their own and others' emotions to further their own and others' personal and social goals. Emotional intelligence had a *direct* effect on well-being, and this association remained strong even after controlling for SES. What's more, of the two measures of intelligence -IQ and emotional intelligence- *emotional intelligence was the strongest predictor of well-*

being, outweighing not only IQ, but also a person's SES and age. This finding suggests that emotional intelligence -particularly the capacity to manage one's emotions toward optimal personal goal attainment- is a form of intelligence that can help people live a more fulfilled life regardless of their economic circumstances.

Why Is Intelligence Associated with Well-Being?

Scott B. Kaufman thinks intelligence matters for a fulfilling life for a number of reasons. For one, a higher IQ is a gateway to better education. Those with higher IQ scores are much more likely to score well on standardized tests of achievement, and academic performance is often the first hurdle necessary to continue up the ladder of occupational opportunities. Also, relevant here is the association between IQ and openness to experience. Those with a higher IQ tend to score higher in several facets of openness to experience, including intellectual engagement, intellectual creativity, introspection, ingenuity, intellectual depth and imagination. This tendency for deeper cognitive processing is critical for dealing with a lot of life's up and downs. While trauma is inevitable in life, research shows that we can grow from our traumas if we have a healthy form of rumination in which we reflect on the deeper meaning of the event and can use that cognitive processing to perceive greater opportunities for ourselves and others. Regarding emotional intelligence, since having a fulfilling life often requires accomplishing the goals you have set out for yourself, it makes sense that being able to manage your emotions in the service of a larger goal will be associated with well-being and self-actualization. Perhaps the most important analysis will turn out to be how IQ and emotional intelligence interact. There is some evidence that in certain contexts, emotional intelligence can amplify the effectiveness of a high IQ, and high emotional intelligence can even compensate for a lower IQ. Future research should definitely look more closely at the interaction between these two important aspects of human intelligence. Of course, it's possible that the findings operate in reverse causation, and being happier increases intellectual skills. Most likely, both directions are at play in the correlations found in the study. Clearly more research will need to look at the association between intelligence and well-being *over time*.

NB: It should be noted that IQ and emotional intelligence were moderately correlated with each other. This suggests that both tests are tapping into a common set of processes (e.g., executive functioning, working memory, etc.), even though IQ and emotional intelligence also involve a partially different set of skills.

How the Body and Mind Talk to Each Other

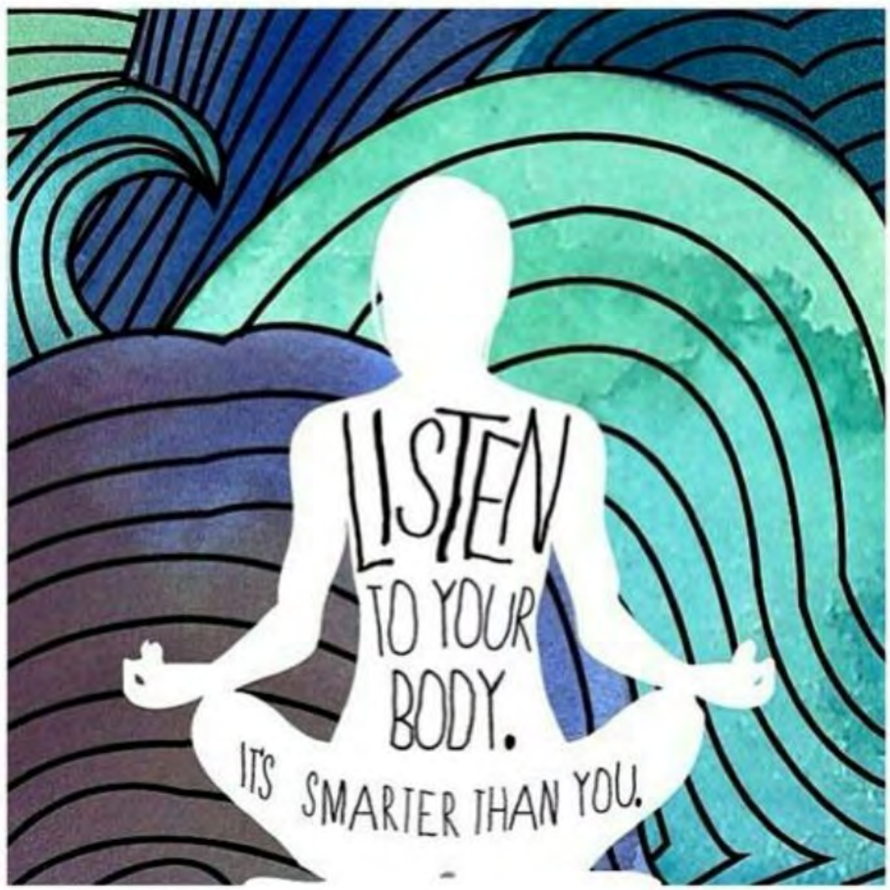

Courtesy of Primal Coding

Have you ever been startled by someone suddenly talking to you when you thought you were alone? Even when they apologize for surprising you, your heart goes on pounding in your chest. You are very aware of this sensation. But what kind of experience is it, and what can it tell us about relations between the heart and the brain?

When considering the senses, we tend to think of sight and sound, taste, touch and smell. However, these are classified as exteroceptive senses, that is, they tell us something about the *outside world.* In contrast, interoception is a sense that informs us about our *internal* bodily sensations, such as the pounding of our heart, the flutter of butterflies in our stomach or feelings of hunger. The brain represents, integrates and prioritizes interoceptive information from the internal body. These are communicated through a set of distinct neural and humoral (i.e. blood-borne) pathways. This sensing of internal states of the body is part of the interplay between body and brain: it maintains homeostasis, the physiological stability necessary for survival; it provides key motivational drivers such as hunger and thirst; it explicitly represents bodily sensations, such as bladder distension. But that is not all, and herein lies the beauty of interoception, as our feelings, thoughts and perceptions are also influenced by the dynamic interaction between body and brain.

The shaping of emotional experience through the body's internal physiology has long been recognized. The American philosopher William James argued in 1892 that the mental aspects of emotion, the '*feeling states*', are a product of physiology. He reversed our intuitive causality, arguing that the physiological changes themselves give rise to the emotional state: our heart does not pound because we are afraid; fear arises from our pounding heart. Contemporary experiments demonstrate the neural and mental representation of internal bodily sensations as integral for the experience of emotions; those individuals with heightened interoception tend to experience emotions with greater intensity. The anterior insula is a key brain area, processing both emotions and internal visceral signals, supporting the idea that this area is key in processing internal bodily sensations as a means to inform emotional experience. Individuals with enhanced interoception also have greater activation of the insula during interoceptive processing and enhanced grey-matter density of this area.

So, what is enhanced interoception? Some people are more accurate than others at sensing their own internal bodily sensations. While most of us are perhaps aware of our pounding heart when we are startled, or have just run for the bus, not everyone can accurately sense their heartbeats when at rest. Interoceptive accuracy can be tested in the lab; we monitor physiological signals and measure how accurately these can be detected. Historically, research has focused on the heart, as these are discrete signals that can easily be quantified. For example, a typical experiment might involve the presentation of a periodic external stimulus (e.g. an auditory tone) that is time-locked to the heartbeat, such that each tone ('*beep*') occurs when the heart is beating, or in between heartbeats. Participants state whether this external stimulus is synchronous or asynchronous with their own heart. An individual's interoceptive accuracy is an index of how well they can do this.

It is also possible to measure subjective indices of how accurate people *think* they are at detecting internal bodily sensations, ascertained via questionnaires and other self-report measures. Sarah Garfinkel's work (University of Essex) shows that individuals can be interoceptively *accurate* (i.e. good at these heartbeat-perception tests) without being *aware* that they are. In this way, interoceptive signals can guide and inform without fully penetrating conscious awareness. Individual differences in interoception can also be investigated using brain-imaging methods, such as through brain representation of afferent signals (e.g. heartbeat-evoked potentials expressed in a neural EEG signal). Functional neuroimaging (fMRI) can also be used to investigate which areas of the brain are more active when focusing on an interoceptive signal (e.g. the heart) relative to an exteroceptive signal (e.g. an auditory tone).

Our hearts do not beat regularly and, while we can identify that our hearts race with fear or exercise, we might not fully appreciate the complexity of the temporal structure underlying our heartbeats. For example, cardiac signatures are also associated with states such as anticipation. Waiting for something to happen can cause our heartrate to slow down: this will happen at traffic lights, when waiting for them to go green. These effects of anticipation, potentially facilitating the body and mind to adopt an action-ready-state, highlight the meaningful composition of internal bodily signals. Internal bodily signals can be deeply informative, which is why sensing them can provide an extra channel of information to influence decision-making. Gut instinct or intuition during a card game can also be guided by interoception. Bodily signatures (heart rate, skin-conductance response) can signal which cards are good (i.e. more likely to be associated with a positive outcome) even in the absence of conscious knowledge that a card is good. Thus, the heart 'knows' what the mind does not yet realize, and access to this bodily signature can guide intuitive decision-making to a better outcome. In a real-world extrapolation of this, Garfinkel visited the London Stock Exchange to work with high-frequency traders. These traders claimed that their decisions were often driven by gut instinct, when faced with fast-coming information that the conscious brain could not yet fully process. Garfinkel and her colleagues demonstrated that interoceptive accuracy was enhanced in those traders who were most adept at trading, potentially grounding their intuitive instincts in a capacity to sense informative changes in internal bodily signals. An appreciation that bodily signals can guide emotion and cognition provides potential interoceptive mechanisms through which these processes can be disrupted. Alexithymia, defined as an impaired ability to detect and identify emotions, is associated with reduced interoceptive accuracy. Autistic individuals, who often have difficulty in understanding emotions, have also been shown to have impaired interoceptive accuracy. Neural representation of bodily signatures are altered in borderline personality disorder (also known as emotionally unstable personality disorder), and interventions designed to focus on the body, such as mindfulness, have been shown to reduce anxiety. Insight into the nature of these embodied mechanisms opens up potential avenues for further understanding and targeted intervention.

As well as telling us about our own emotions, our bodies respond to the joy, pain and sadness of others. Our hearts can race as loved ones experience fear, and our pupils can adopt a physiological signature of sadness in response to the sadness of others. If you pay attention to your heart and bodily responses, they can tell you how you are feeling, and allow you to share in the emotions of others. Interoception can enhance the depth of our own emotions, emotionally bind us to those around us, and guide our intuitive instincts. We are now learning just how much the way we think and feel is shaped by this dynamic interaction between body and brain.

Laughter Is What Made Us Humans

A Jeep full of the Daughters of Charity in St Louis, Missouri in 1964 -Photo by Bert Glinn/Magnum

'All the acts of the drama of world history were performed before a chorus of the laughing people.'
From "Rabelais and his World" (1965) by Mikhail Bakhtin

The central question that anthropologists ask can be stated simply: '*What does it mean to be human?*' In search of answers, we learn from people around the world – from city-dwellers to those who live by hunting and gathering. Some of us study fossil hominins such as Homo erectus or the Neanderthals; others look at related species, such as apes and monkeys.

Something that sets us apart from these ancestors and primate relatives, and should be of special interest to anthropology, is our unique propensity to laugh. Laughter is a paradox. We all know it's good for us; we experience it as one of life's pleasures and a form of emotional release. Yet to be able to laugh, we must somehow cut ourselves off from feelings of love, hate, fear or any other powerful emotion. The fall of a pompous fool slipping on a banana skin is the cliché of comic routines; we laugh at his misfortune because we don't really care.

A helpful way to get a handle on laughter is to place it in evolutionary context. Other animals play, and their playful antics can prompt vocal sounds. But human laughter remains unique. For one, it is contagious. When a group of us get the giggles, we soon become unmanageable. The evolutionary psychologist Steven Pinker notes that this might be what allowed laughter to be pressed into the service of humor. In *How the Mind Works* (1997), he writes:

> *No government has the might to control an entire population ... When scattered titters swell into a chorus of hilarity like a nuclear chain reaction, people are acknowledging that they have all noticed the same infirmity in an exalted target. A lone insulter would have risked the reprisals of the target, but a mob of them, unambiguously in cahoots in recognizing the target's foibles, is safe.*

Besides contagion, laughter also leaves us peculiarly helpless and vulnerable. We can be doubled up with laughter, or laugh until we weep. Physiologically, it can come close to crying. Nearly every aspect of the body – voice, eyes, skin, heart, breathing, digestion – can be powerfully affected. What we find funny might vary by culture, but people across the world make essentially the same sounds.

When we apply Darwinian theory to laughter, it's tempting to look for a plausible precursor among our ape-like ancestors. The primatologist Jane Goodall, for example, points out that young chimpanzees often engage in tickling games, making huffing and puffing noises all the while. Maybe, then, human laughter is best viewed as an evolutionary extension of certain playful vocalizations already found among apes. The objection to this theory is that ape tickle-play vocalizations don't sound like human laughter at all – they are more like heavy breathing, with inhalations and exhalations equally audible. Another problem is that the apes' sounds are not socially contagious, and don't bond the group together in quite the same way. No chimpanzee will laugh just because others are doing so – each animal must itself be tickled. By contrast, when humans meet up on social occasions, the most frequent sounds you're likely to hear are not grunts and screams but ripples of laughter. Those sounds convey a certain level of relaxed happiness in the company of others. Although monkeys and apes can be friendly, their face-to-face social dynamics are typically competitive and despotic in ways that humans tend to find intolerable. Everyday encounters between nonhuman great apes oscillate between dominance and submission, with facial expressions and instinctive vocalizations to match. There is nothing egalitarian about their encounters.

Building on these insights, scores of theorists have attempted to explain why humans evolved to be the species that laughs. One classic idea is the Superiority Theory, according to which the loudest laughs were originally cries of triumph made at the expense of the enemy. Another is the Relief Theory, in which laughter is thought to have evolved long before words or grammar, as an instinctive way of signaling that danger had passed and everyone could relax. Finally, the Ambivalence Theory holds that laughter erupts as a means of escape from contradictory emotions or perceptions. What these ideas have in common is their focus on individual psychology. In each case, the thinking is that tension is released with the sudden realization that there is nothing to fear. For supporters of the Superiority Theory, the initial threat comes from other people who are suddenly exposed as harmless. The Relief Theory agrees that we laugh upon realizing we are safe. The Ambivalence theory also proposes that laughter arises when a mental or physical challenge or paradox suddenly dissolves.

The shared insight can be expressed in a single word: *reversal*. The evolution of the human smile neatly illustrates the idea. When we smile, we stretch out the corners of our mouth and show our teeth. If other animals were to bare their teeth in this way it would be threatening. In the case of nonhuman primates, baring the teeth can be more ambivalent - as in the chimpanzee '*fear grin*', which simultaneously shows resistance and submission to a more dominant animal. Although humans, too, sometimes behave in this way, we can all spot the difference between a nervous grin and a genuinely warm smile. So, it seems likely that the happy smile is probably a fear-grin that has adapted to relaxed social conditions, its significance reversed because there is no longer anything to worry about.

Against existing theories, however, I view laughter as a more profound social and collective endeavor – though still tied to reversal. Smiling, after all, can easily become laughter, so it's worth exploring whether reversal might explain this behaviour too. When animals collectively mob an enemy, they sometimes bare their teeth and make threatening sounds. Typically, there is something rhythmic, contagious and emotionally bonding about those intimidating screams and cries. The primatologist Frans de Waal describes how an angry coalition of female chimpanzees can sometimes direct a chorus of '*woaow*' barks at a misbehaving male, maintaining the noise until he finally gets the message. If human laughter evolved through progressive modification of ancestral primate signals, then – as proposed by the late ethologist Irenäus Eibl-Eibesfeldt – similar '*mobbing*' cries constitute a likely candidate.

Mobbing, then, might be the behavioral precursor to laughter. Taking a step further, it might even help to account for the broader architecture of the human mind. Over evolutionary time, our psychology has been shaped by the demands of face-to-face relationships based on mutual respect; we have adapted to reflect a much more egalitarian socio-political order than anything known among monkeys and apes. The break is so sharp that there must have been some kind of radical regime-change – a *human revolution*, as Chris Knight and some of his colleagues call it– to accomplish the transition from ape-like politics to hunter-gatherer-style egalitarianism. The evolutionary anthropologist Christopher Boehm has proposed an influential theory about the emergence of human society that he terms Reverse Dominance. According to Boehm, great-ape society is like a pyramid, with one despotic leader – the alpha male – at the apex and the rank-and-file underneath.

By contrast, Boehm notes that our hunter-gatherer forebears were profoundly egalitarian. He argues that this was established not simply via incremental change, but in the final stages, through an upheaval so profound that political relationships went into reverse. By this he means that certain rebel coalitions, formed to resist the dominant males, eventually became all-embracing and powerful enough to overthrow the former regime. In its place, a political system was established that still prevails among many hunter-gatherers to this day: Reverse Dominance or community-wide rule from below. What Boehm terms Reverse Dominance is an upturned pyramid, with the rank-and-file dominant over any would-be alpha male. While Boehm himself doesn't mention laughter, it seems likely that such a profound political revolution would trigger a great sense of relief. When the threat posed by the fear-inducing alpha-male was defied, we can imagine the rejoicing and laughter that must have accompanied such a reversal of fortune. For our evolving species, perhaps laughter is a marker of our irrevocable departure from the psychology of apes.

From then on, society was decisively egalitarian, with power – now socially accountable – in the hands of the community as a whole. One consequence was that no-one could simply follow their instincts or pursue their own selfish agenda. You needed to consider what everyone else thought, on pain of being laughed out of town. Collective laughter, then, might have served as a social levelling device helping to keep everyone in line. The outcome was not only a social and political reversal but also a cognitive one: a transition that every child re-enacts as it develops into a self-aware, smiling, laughing, fully human being. Because of the human revolution, whenever we engage with one another informally, we find it natural to put one another at ease, and to establish at least the appearance of equality. This has become so habitual that our instinctive social signals, inherited from our primate ancestors, have been largely repurposed: the tense primate fear-grin has given way to the relaxed human smile, while the angry mobbing cry has transformed into uproarious laughter. The emotional significance of the signal might be reversed, but remnants of its original form and meaning have been preserved.

For most of the time since the emergence of our species some 300,000 years ago, we have been hunter-gatherers. To answer the anthropologist's question about what it means to be human, then, modern hunter-gather societies remain particularly important. Every aspect of our minds and bodies has evolved in response to this long-lasting and immensely stable way of life. It's true that as a species we have evolved to be flexible, a kind of second-order adaptation – but when we find adapting to power inequalities stressful, as many of us do today, it damages both our physical and our mental health. Our need for companionship, for relaxed playfulness, for opportunities to sing and laugh together – all these things have their roots in the hunter-gatherer way of life.

It was once imagined that people in these societies must have always struggled to survive, teetering on the edge of starvation, their relentless quest for food leaving them with no time for leisure or play. It's hard to know where this strange idea came from, because it is utterly wrong. The prejudice was refuted by the anthropologist Marshall Sahlins in *Stone Age Economics* (1972), which describes 'the original affluent society'. Today's hunting and gathering peoples, Sahlins explained, have a far more healthy and varied diet than people who farm or live in cities. Theirs is an economy of abundance, even super-abundance. Hunter-gatherers typically enjoy hours of leisure time for creative activities such as art, dancing and singing. A striking feature of these societies is their profound egalitarianism. As an anthropologist, Chris Knight can report that in any hunter-gatherer camp, equality is maintained by almost nonstop laughter aimed at anyone who is getting above themselves. Everywhere you look, there is a palpable atmosphere of playfulness and fun. It's no coincidence that the gods of hunter-gatherers are not solemn guardians of morality, but mischievous tricksters whose antics provoke helpless mirth in listener and storyteller alike.

Anthropologists who have studied hunter-gatherers have contrasted '*immediate return'* societies – those who do not store food supplies – with '*delayed return'* or storage-based societies. In immediate-return societies, resources such as hunted game are meticulously shared out from the moment the hunters return. No one is allowed to acquire power by storing wealth of any kind. However, once food begins to be stored, certain individuals acquire more power than others, and egalitarianism is progressively undermined. Today, it's a settled consensus that Africa is where our species evolved, and where large game animals such as elephants and giraffes have managed to survive into the modern age. For this reason, with few exceptions, African hunter-gatherers have traditionally enjoyed economic abundance and fall into the '*immediate return'* category. More than city-dwellers or farmers, these people can inform us about how to use laughter as a way of maintaining egalitarianism. Grandmothers and other senior females demonstrate how, by derisively laughing at those who throw their weight around or put on airs and graces, people can be persuaded to respect egalitarian norms. Nothing in those psychologically individualistic theories about mocking, or about experiencing relief from fear or tension, implies anything specific about the social conditions required for laughter to flourish. But social anthropologists all agree that, among hunter-gatherers, laughter functions as a levelling device, bringing people down to size. The major figure here is Jerome Lewis at University College London, who has been studying the Mbendjele people in the Republic of Congo for many years. This is a population of immediate-return hunter-gatherers who live in the forests of Central Africa. They are sometimes referred to as '*Pygmy people'*, because of their short stature. Among the Mbendjele, Lewis is able to pin-point exactly how laughter maintains egalitarianism in practice. He explains that it would be risky for a young person to make fun of an older one, no matter how foolish the elder's behavior. But senior women exercise a special privilege, seeing it as their enjoyable role to bring down anyone who seems to be getting above themselves.

By way of example, Lewis relates how a woman who is upset with her husband's behavior – he might be chasing another woman, or not providing enough to eat, or not having sex often enough with her – will go to sit with other women in a prominent place. In loud, exaggerated tones, she talks about her problems with her husband, while her listeners enthusiastically take up her gestures as she mimes his actions and expressions. This is a terrible situation for the hapless husband as he hears the women, children and other men laughing boisterously at his expense. One of Jerome's students in the field, Daša Bombjaková, has given us a detailed picture of how women's laughter manages to keep men in line. In any dispute with a man, a woman can expect support from other women, regardless of the rights and wrongs. Should a male threaten or attempt violence, other women will immediately support the victim and give her shelter and protection. They will collectively insult the man and might give him a beating, typically with their cooking utensils. But their main weapon is derisive laughter.

Male sexual behavior is the constant butt of jokes. Women have a powerful ritual called *Ngoku,* designed to re-assert female values. The entire female community comes together to sing and dance, seizing control of the public space and making graphic, ribald fun of male sexuality. It is also during the build-up to *Ngoku* and immediately afterwards that women will take part in *mòádzò* – female, public, mocking re-enactments of stereotypically male behavior. Bombjaková went to the trouble of listing the faults likely to be ridiculed. These included greediness, selfishness, dishonesty, cheating, laziness, arrogance, boastfulness, carelessness, cowardice, intolerance, moodiness, impulsiveness, aggression and possessiveness. A feature of *mòádzò* is that few words are used. Instead, performers utter expletives and onomatopoeic sounds that are amplified by the audience and stretched out musically. A senior woman might start the ball rolling by silently imitating some characteristic mannerism of her target. One or two others immediately grasp whom she means. They begin to laugh and, because laughter is so contagious, soon everyone is laughing and uproariously pantomiming the behaviour being mocked. After a while, the only person still not laughing is the man himself. But the laughter goes on until, at last, even he gets the joke. The chorus subsides only as he finally joins in, laughing at his own expense. He now sees the funny side of things, at last viewing himself as others see him. His behavior was comical because of its incongruity with what is deemed socially acceptable. The culprit might have imagined he could get away with such outrageous actions. But hunter-gatherer women adopt a collective perspective on badly behaved males and will do everything possible to bring each culprit back into line.

Although it can seem cruel, the truth is that women's laughter is generous and inclusive. Despite its hurtfulness, the target is invited to save face by joining in. A good *mòádzò* performance will succeed in calming the atmosphere by allowing everyone to laugh and forget their anger.

Looking at laughter from the perspective of an anthropologist, it's possible to claim that all humor is essentially political. That insight transcends comedic forms such as satire; my point here is that humor in general, whatever its content, is political by nature. Down to the smallest details of our lives, our relationships and encounters involve exercises and exchanges of power. In the face of these dynamics, laughter is an equalizing gesture, a restoration of a rightful order in the face of an unjust hierarchy. Similarly, when we find something funny, it's often because of some incongruity between mind and body, the ideal and the real. That division is political to the core. The humor of medieval carnival, according to Bakhtin, relied on the way that the body makes a mockery of the lofty purposes of the mind. Buttocks, thighs, coughs, splutters, farts, 'the bodily lower stratum' – all mock the spiritual solemnities of humorless bishops and other supposed guardians of morality. Comedy is about exposing the gap between our supposedly noble intentions, and the grimier truths about our condition. For this reason, the amusing features of life are never far away; if something seems funny, it's because it's uncomfortably close to home.

Perhaps this explains why nothing in nature can be truly comical. A strange rock formation or a pattern in the clouds might seem weird or intriguing, but it can't be amusing; rocks and clouds don't have human-style intentions and motivations, so they can't be tripped up. Animals might seem funny, but only because we anthropomorphize them. They can't be brought down to earth, because that's where they are already. For a situation to provoke genuine laughter, it must form a pattern that we recognize from our own mental and social lives. Laughing, then, appears to be intimately tied to our ability to reflect back on ourselves. When we chuckle at our own foibles, we show that we are no longer trapped inside our individual egos, but can see ourselves through one another's eyes. Likewise, when speaking, we separate ourselves from those around us by using words such as '*I*' or '*me*', drawing attention to ourselves as one person among others, as if from outside. Language would be impossible without the ability to adopt such a reverse-egocentric standpoint.

Humans are instinctive egalitarians, who work best with one another when no one has absolute authority, when teasing is good-natured, when there is sufficient affection and trust for shared tasks to constitute their own reward. Laughter is a vital part of this picture – not simply a psychological relief valve, but a collective guard against despotism. When moved to laugh by those around us, we reveal ourselves to be truly human.

The Universe Is a Hologram

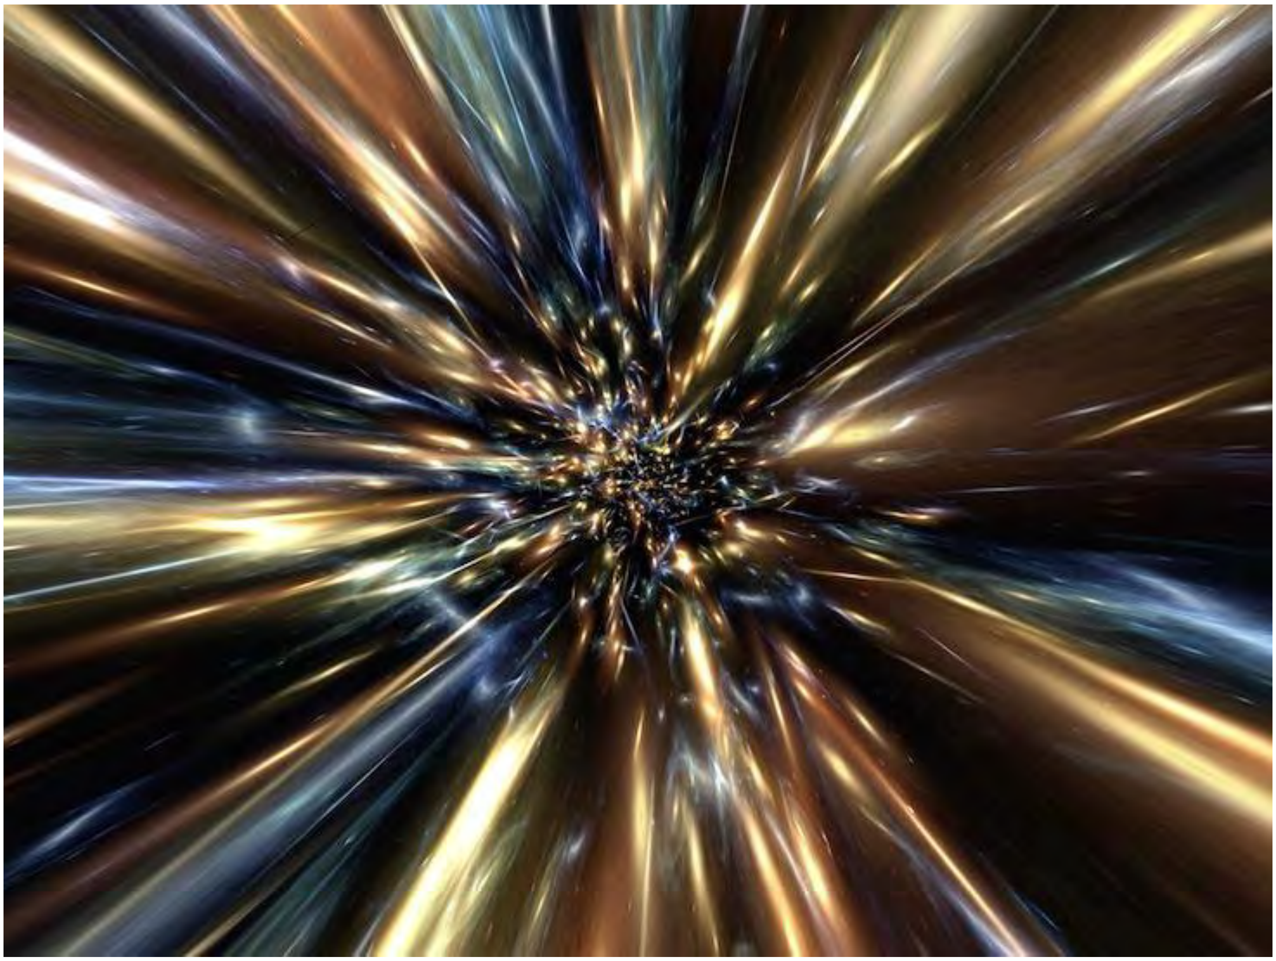

Photo illustration by L. Lacertae/Flickr

One of the great mysteries of modern cosmology is how our universe can be so thermally uniform -the vast cosmos is filled with the lingering heat of the Big Bang. Over time, it has cooled to a few degrees above absolute zero, but it can still be seen in the faint glow of microwave radiation, known as the cosmic microwave background. In any direction we look, the temperature of this cosmic background is basically the same, varying by only tiny amounts. But according to the standard "*cold dark matter*" model of cosmology, there wasn't enough time for hotter and cooler regions of the early universe to even out. Even today we would expect parts of the cosmic background to be much warmer than others, but that isn't what we observe.

One solution to this cosmological problem is known as early inflation. If the observable universe was extremely tiny in its earliest moments, it could have reached a uniform temperature very quickly. Afterwards, the theory says, the universe underwent a brief period of rapid expansion, eventually leading to the universe we observe today. We don't have any direct evidence for early cosmic inflation, but because it would solve several issues in cosmology, it is a widely supported idea.

Recently, a team of astronomers looked at data from the Planck satellite, which gathered the most accurate measurements of the cosmic background thus far. They wanted to compare fluctuations across vast regions of the sky, known as low multipole moments, with the predictions of the standard cosmological model and a model that's somewhat stranger, a holographic one. What if everything around you, from the distant stars to your very hands, were a hologram? Like Plato's cave, our world of solid objects and three-dimensional space would simply be a shadow of a two-dimensional reality. On the human scale a holographic universe would be indistinguishable from the reality we expect, but on a cosmic scale there could be subtle differences we might be able to detect. In the holographic view of cosmology, early inflation is driven by interactions of the quantum field, which would slightly change the appearance of the cosmic microwave background. This is particularly true for low multipole moments, and this difference makes it possible, at least in principle, to prove that the holographic principle is true. In their paper, published in January 2017 in *Physical Review Letters,* the team report the holographic model fitting the Planck satellite data slightly better than the standard model. The results don't prove the universe is holographic, but they are consistent with a holographic model. The idea that our universe might be holographic comes from string theory. Although string theory hasn't been proven experimentally, its mathematical structure has an elegance and power that makes it appealing as a theoretical model. The holographic principle in string theory is just such an example. In its broadest form, the holographic principle states that anything you can know about a particular volume of space can be learned by looking at the surface enclosing the volume. Just as a hologram can contain a three-dimensional image within a sheet of glass or plastic, the universe could contain its vast volume within a surface.

For example, imagine a road 10 miles long that is "*contained*" by a start line and a finish line. Suppose the speed limit on this road is 60 miles per hour, and we want to know if a car has been speeding. One way to do this is to watch a car travel the whole length of the road, measuring its speed the whole time. But another way is to simply measure when a car crosses the start line and finish line. At a speed of 60 miles per hour, a car travels a mile a minute, so if the time between start and finish is less than 10 minutes, we know the car was speeding. If the holographic principle is true, then the universe can be viewed in two different ways: one of space and volume as we intuitively experience it, and one of a "*surface*" with one less dimension. This holographic duality is mathematically powerful because some laws of physics can be much easier to work with in one view than the other.

The structure of our universe is driven by the constant pull of gravity between stars and galaxies. In the present era, gravity is weak compared to other forces, and is described as a gravitational field in general relativity. In the dual holographic view, gravity is described as a quantum field that can interact strongly with mass. Since it is easier to calculate weak interactions than strong ones, the general relativity approach is more useful. However, in the early moments of cosmic time, when the universe was hot and dense, the gravitational fields of relativity were strong, so quantum fields of the holographic view might be easier to deal with. The fact that both the standard and holographic models can account for early inflation supports the idea that the holographic principle applies to our universe. Cosmic inflation remains a mystery, but by viewing the universe as a hologram we might just be able to solve it.

The Case for Transmissible Alzheimer's Grows

What separates a lethal prion from a dementia-associated amyloid plaque? Maybe not much.

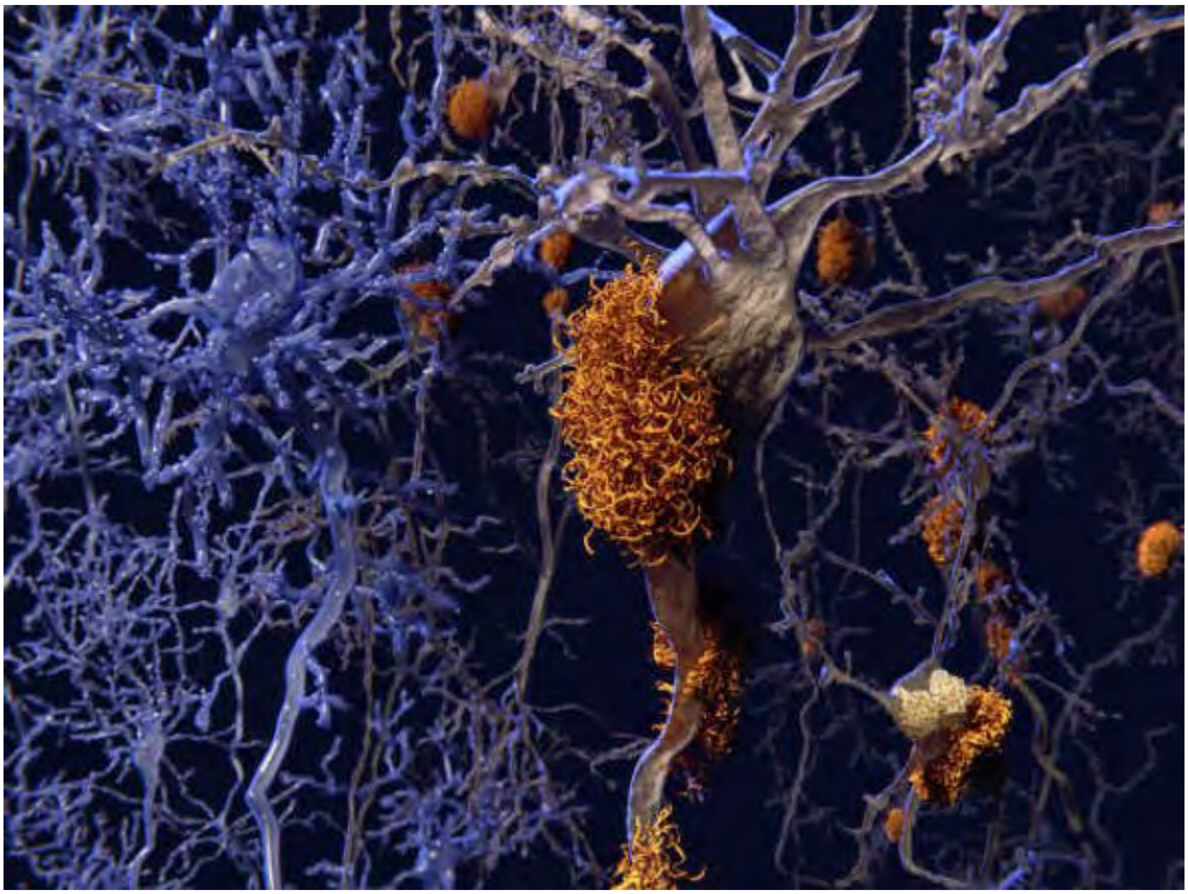

Orange amyloid beta plaques accumulate on a neuron in this computer-generated image illustrating the pathology of Alzheimer's Disease - Credit: Juan Gaertner Getty Images

The unsettling evidence that Alzheimer's Disease may be transmissible under limited - but definitely nonzero - circumstances keeps growing.

A study suggested that peptide aggregates – essentially sticky, self-propagating clumps of misfolded protein bits collectively referred to as amyloid- found in the brains of Alzheimer's patients may be transmissible in the same ways that prions are.

Then, a new paper appeared in *Nature* that seemed to take the evidence for the transmissibility of Alzheimer's peptides from "*circumstantial*" to "*experimentally* *produced*". It is unsettling, news, that further blurs the line between amyloid and prions.

Human prion diseases are rare. Prions usually form spontaneously or are inherited via faulty genes, but sometimes find their way into humans through consumption of contaminated brain or spinal cord tissue. In the case of Mad Cow Disease, it happened via contaminated beef.

In rare cases (so far as we know), human prion transmission has happened when surgical instruments used on an infected patient were cleaned and reused on an uninfected one. Prions stick to steel like glue, are stable for decades at room temperature, and survive a bombardment of chemical and physical cleaning assaults that are more than sufficient to obliterate other pathogens. Prions are survivors. In the original Alzheimer's transmissibility study, scientists examined the brains of eight patients treated with prion-contaminated human growth hormone as children who decades later died from prion disease (out of over 30,000 people so treated, more than 200 died this way). The hormone had become contaminated with prions because it had been extracted from cadavers - one or a few of whom presumably died of prion disease - and processed in such a way that the prions remained. Of course, prions are not the only misfolded proteins that potentially lurk in the brains of cadavers.

The researchers discovered the brains of seven of the eight contained, in addition to prions, peptide aggregates called Amyloid beta (Aβ for short). Aβ is a collection of misfolded peptides whose correctly folded versions are present in the human brain and perform a variety of mid-level tasks. When the misfolded versions form, they behave like prions, catalyzing the conversion of healthy forms into diseased ones and accumulating in clumps called plaques. Indeed, past experiments have shown that injecting small amounts of human Aβ into the brains of primates or of mice bred to express a humanized form of the Aβ precursor protein generates Aβ plaques in these animals.

Plaques are characteristic of and possibly the instigators of Alzheimer's Disease when they accumulate around neurons in the brain. However, the seven brains did not have plaques. The Aβ in these brains had built up in the walls of blood vessels, where such accumulations can cause bleeding and dementia. This condition is called cerebral amyloid angiopathy, and it co-occurs with most Alzheimer's Disease but can also strike on its own. The eight victims had all still been young enough that their brains would not be expected to show any signs of Alzheimer's or cerebral amyloid angiopathy unless they had genetic risk factors. Understandably, given the implications, the scientists who studied their brains were concerned.

The December *Nature* study was authored by this same team. In it, they revealed that they had managed to get their hands on original vials of prion-contaminated growth hormone that had been helpfully squirreled away for decades by Public Health England. They tested the samples for both Aβ peptides and tau, another protein that builds up in the brains of Alzheimer's patients and causes its other brain pathology: tangles. Indeed, two types of Aβ and tau were still present in the vials, even after more than *three decades of room temperature storage*. Aβ and tau, at least, are survivors too. This team took their study a step further by injecting a tiny sample of these vintage vials into the brains of mice engineered to be susceptible to human Alzheimer's. The mice developed both Aβ plaques and cerebral amyloid angiopathy, although they showed no signs of tau. Aβ peptides had not only managed to survive decades of room-temperature storage, they were also still transmissible. This is concerning.

It is imperative to emphasize that transmissible does not equal contagious. There is absolutely no evidence that people with dementia can spread their disease casually to people around them. Even donated blood appears to be safe, as no association with blood transfusions and Alzheimer's Disease has ever been detected. Rather, in the cou2019rse of some neurological surgeries – and perhaps certain kinds of medical exams – prions may become lodged on equipment. And there is a chance this equipment could transmit the disease. Organ donation protocols may also warrant some review. It was already known that donations of dura mater, a tough brain covering, have transmitted Aβ to young people in the past.

Since Alzheimer's Disease is so common, and we have not been looking for Alzheimer's caused by surgical or other medical procedures that access eye or neural tissue - particularly in patients for whom the appearance of Alzheimer's would not be surprising - is it possible that we are underestimating the transmission potential of this disease, and that such events are less rare than we would guess? Alzheimer's is not the only neurodegenerative disease in which aggregating misfolded host proteins – a class referred to as amyloid-- seem to propagate and wreak havoc either.

In Parkinson's Disease, misfolded alpha-synuclein proteins spread through the brain, and in Amyotrophic Lateral Sclerosis (Lou Gehrig's Disease), the misfolded, accumulating protein is TDP-43. We should investigate the transmission potential of these diseases as well. The only thing that seemed to separate these conditions from classic prion diseases was transmissibility. But now that that barrier has been breached for at least one: What is the difference between amyloid and prions? Are they part of a spectrum? Are they one and the same? If not, what is the difference? Can what we've learned about the biology of prions help our efforts to fight amyloid dementias? Of course, since we still can't cure prion diseases, it may not be much help even if so. The realization that the peptides involved in some of the most common and feared dementias on Earth may be transmissible under even limited conditions is a sobering and humbling reminder of how very little we still understand about them. Given what we know about prions, we would be wise not to underestimate their abilities.

Acknowledgements

My sources were many, as can be expected from the mosaic of *verbatim* copying. This selection is a beginning: the future updates on Neurosciences and related areas, will be emailed regularly to readers interested. An announcement will be displayed on the website shortly. Yves P. Huin was the skilled reviewer, editor and formatter.

References

- https://www.scientificamerican.com/author/nature-magazine/
- news@email.scientificamerican.com
- https://www.scientificamerican.com/article/for-alzheimers-sufferers-brain-inflammation-ignites-a-neuron-killing-forest-fire/?utm_source=newsletter&utm_medium=email&utm_campaign=daily-digest&utm_content=link&utm_term=2019-03-04_featured-this-week&spMailingID=58640392&spUserID=Mzg2Mzk5NTQ1NDA0S0&spJobID=1600527785&spReportId=MTYwMDUyNzc4NQS2
- https://www.quantamagazine.org/the-brain-maps-out-ideas-and-memories-like-spaces-20190114/?utm_source=Quanta+Magazine&utm_campaign=ee876b780d-RSS_Daily_Biology&utm_medium=email&utm_term=0_f0cb61321c-ee876b780d-390008349&mc_cid=ee876b780d&mc_eid=13587a4bcf
- https://www.scientificamerican.com/article/there-is-no-such-thing-as-conscious-thought/
- https://www.scientificamerican.com/article/computers-determine-states-of-consciousness/?utm_source=newsletter&utm_medium=email&utm_campaign=mind&utm_content=link&utm_term=2018-12-19_featured-this-week&spMailingID=58075562&spUserID=Mzg2Mzk5NTQ1NDA0S0&spJobID=1542660458&spReportId=MTU0MjY2MDQ1OAS2
- https://www.quantamagazine.org/new-clues-to-how-the-brain-maps-time-20160126/
- https://nautil.us/blog/new.evidence.for.the.geometry.of.thought
- https://aeon.co/essays/how-experimental-psychology-can-help-us-understand-art?utm_source=Aeon+Newsletter&utm_campaign=06367fdc04-

EMAIL_CAMPAIGN_2019_01_16_05_40&utm_medium=email&utm_term=0_411a82e59d-06367fdc04-70468141

- https://onezero.medium.com/the-brain-that-remade-itself-bcc7b3a43cff
- https://www.quantamagazine.org/foundations-built-for-a-general-theory-of-neural-networks-20190131/?mc_cid=4be2a6d6a7&mc_eid=13587a4bcf
- https://www.basicbooks.com/titles/lance-esplund/the-art-of-looking/9780465094677/
- https://www.theatlantic.com/entertainment/archive/2019/01/art-looking-living-organism/580177/
- https://aeon.co/ideas/philosophy-can-make-the-previously-unthinkable-thinkable?utm_source=Aeon+Newsletter&utm_campaign=06367fdc04-EMAIL_CAMPAIGN_2019_01_16_05_40&utm_medium=email&utm_term=0_411a82e59d-06367fdc04-70468141
- https://blogs.scientificamerican.com/beautiful-minds/can-intelligence-buy-you-happiness/
- https://aeon.co/ideas/how-the-body-and-mind-talk-to-one-another-to-understand-the-world?utm_source=Aeon+Newsletter&utm_campaign=76470f634a-EMAIL_CAMPAIGN_2019_02_13_01_42&utm_medium=email&utm_term=0_411a82e59d-76470f634a-70468141
- https://aeon.co/essays/does-laughter-hold-the-key-to-human-consciousness?utm_source=Aeon+Newsletter&utm_campaign=76470f634a-EMAIL_CAMPAIGN_2019_02_13_01_42&utm_medium=email&utm_term=0_411a82e59d-76470f634a-70468141
- https://getpocket.com/explore/item/new-evidence-for-the-strange-idea-that-the-universe-is-a-hologram
- https://blogs.scientificamerican.com/artful-amoeba/the-case-for-transmissible-alzheimers-grows/?utm_source=newsletter&utm_medium=email&utm_campaign=weekly-review&utm_content=link&utm_term=2019-02-13_top-stories