MARK COHEN's training is equal parts engineering and neuroscience. His contributions include his critical role in the development of practical echo-planar scanning, ultra-fast MRI applications, contrast-based and BOLD functional MRI, applications of linear systems analysis to increase fMRI sensitiivity and resolution, and concurrent recordings of EEG and fMRI to better understand brain dynamics and distributed processing. He and his lab have contributed to an understanding of the power of pattern recognition and machine learning to both interpet/classify neural data and as a source of discovery of the processes that result in cognition, perception, emotion and pathology.

Cohen is passionate about graduate and post-graduate education. As the creator and director of the UCLA/Semel NeuroImaging Training Program he has pushed his students to an integrative understanding of the role of imaging in neuroscience: The use of images as hypothesis tests, the relationship between blurring, convolution, statistical error and inference from images, and an understanding of the structures common to neuroimages regardless of imaging modality.

His current focus now includes inquiry into the broader problems of images, beyond neuroscience, to encompass astronomy and nanoscale imaging, aesthetics to statistics, dimensional compression and dimensional expansion.

Cohen holds appointments in the UCLA Departments of Psychiatry, Neurology, Radiology, Biomedical Engineering, Psychology and Biomedical Physics and is a member of the California NanoSystems Institute.

WITH SPARSITY IN MIND

When we look at *Donna che dorme* we not only see the nominal subject: a woman sleeping. We see also her beauty, her calm and comfort...

MARK COHEN
NEUROSCIENTIST
university of california, los angeles

NEUROSCIENCE

...and we see the deep feelings of love that Pablo Picasso felt towards his subject (Figure 01). There is tenderness in the lines and in the intensity of focus that the artist lavishes on his model.

All of this is the more remarkable in that Picasso made this drawing using (by my count) just 29 lines on paper. The sparsity of this work is astonishing, not just as demonstration of Picasso's towering skill, but also as an insight into seeing itself.

01

01 donna che dorme
Pablo Picasso, 1952

our SENSES

To a neuroscientist driven to understand, or at least gain insight into, how the physical reality of the brain interacts with the subjective reality of perception and consciousness, works like *Donna che dorme* offer challenge. How is it that such a seemingly simply form becomes a vehicle for such depth of expression?

as FILTERS

07

A portion of the answer comes from relatively well-understood physiology of the most peripheral portions of the visual system: the eyes themselves. Working on the visual system of the horseshoe crab some sixty years ago, Haldan Hartline and colleagues showed that light reaching individual photoreceptors (light detecting neurons) reduced the responsiveness of their neighbors, a phenomenon known as lateral inhibition (Hartline, Wagner and Ratliff). In brief, each individual photoreceptor will detect light from a single fixed location in the external world. Through a simple neural circuit located within both the eye of the horseshoe crab, and in the eye of the human (Figure 02A), that single photoreceptor will suppress the output of its spatial neighbors. One effect of lateral inhibition is to enhance the contrast at intensity boundaries. Thus, when we look for example, at adjacent grey rectangles whose darkness differs, the edges appear particularly prominent (Figure 02B). In effect, our eyes preferentially present our brain with edges rather than areas. We are tuned, by evolution, to attend to edges, and those edges are enough for us to draw conclusions about our world. Certainly this forms part of the reason that Picasso's line art satisfies us. Part of the reason, but far from enough; this is not quite so trivial a problem.

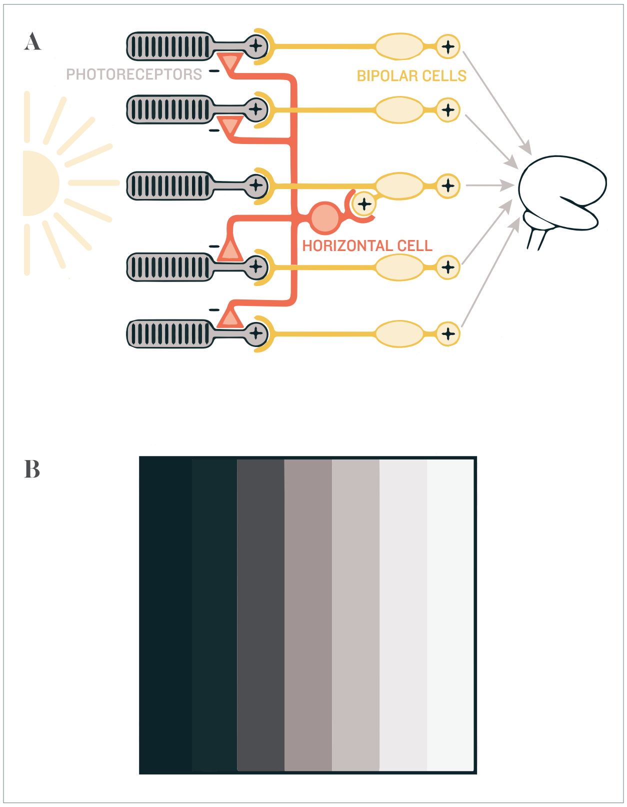

02

02 a. simplified writing diagram of the eye

Light enters our eyes through the lens where it is projected to the retina, a complex neural structure specialized for vision. Across the retina are more than 100 million photodetectors (grey) that respond to light from individual locations in the external world. Rather than sending their signals directly to the brain, the output of each photoreceptor is conveyed through neurons known as bipolar cells (yellow). The (+) sign indicates that output from the photoreceptors increases the output of its bipolar cell. A remarkable feature of our eyes is that the bipolar cells send signals backwards to the photoreceptors through "horizontal cells" (red) that reduce the output of their neighbors, as indicated by the (–) sign. This lateral inhibition has the effect of enhancing the contrast and salience of edges before the signals reach our brains.

02 b. perceived effects of lateral inhibition

The figure shows a set of gray bars of varying density. Perceptually, however, we seem to see the darker bars becoming darker at borders with their lighter neighbor, and vice versa.

08

We now know that the visual system of our brains separates multiple components of the scenes that we gaze at, and processes these separately. Specialized brain regions are known for color, texture, object location, speed and the direction of motion of the things we observe (Figure 03). Observing each of these features of the visual world has proven useful to our ancestors in the past, and evolution has gradually enhanced the ability of our brains to detect them.

At the same time, however, there are a host of physical factors we cannot detect or perceive. For example, we cannot detect the weak magnetic and electrical fields that are present in virtually everything around us. We cannot see the infrared, or ultraviolet reflections, and we cannot see the polarization of light, yet we know that other living organisms are apparently able to detect and respond to each of these. There are surely a huge number of other such physical attributes that we know nothing of.

Life is difficult to sustain and each ability that we have comes at a cost.

09

Each new sensory or physical ability of our bodies requires energy in the form of food, and re-adaptation by some other part of our body to support it. If our ears didn't convey our species with survival advantage they would long ago have been omitted, along with the large portion of our brain dedicated to hearing. As a result, we would require less energy to survive and more able to reproduce. However, hearing has proven an enormous net benefit to our species' survival.

Our minds are resident within brains whose sensory capabilities are the result of relentless pruning by evolution. When we make a decision that the object we pluck from a tree is an orange, we do so based only upon the attributes we can detect: its color, shape, texture, size, smell, location, and so on. All of the remaining characteristics of an orange, such as its chemical makeup, its nutritional value, its low toxicity and the like must be inferred from the sharply limited information that we have.

09

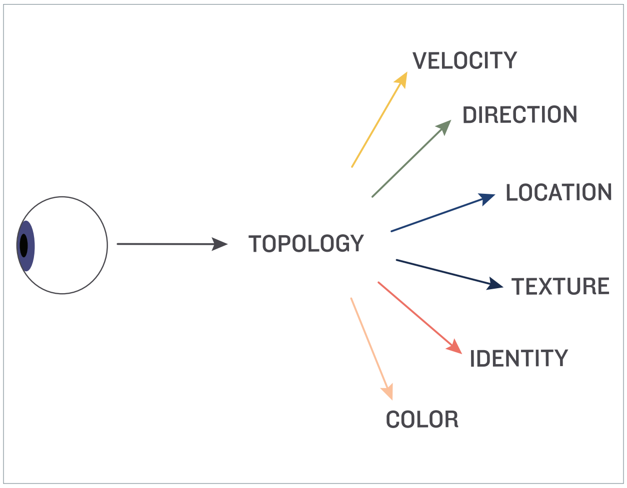

03

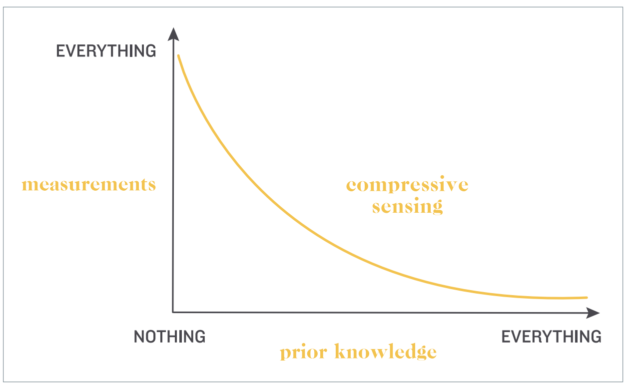

04

03 Once visual information enters our brains, various features of the visual scene are processed separately in different brain regions. As suggested in the diagram, specialized brain regions have been identified for speed and direction of motion, location, texture, color and the identity of visual objects.

04 compressive sensing

(from an idea of Michael Lustig)
In general, the more we know about something a priori the fewer measurements are needed to identify it.

10

the SPARSITY of IMAGES

The cameras available in cell phones currently boast sensors that contain on the order of 16 million pixels—each pixel being roughly analogous to a single photoreceptor in the retina. The electronics in the sensor device represent each color with about 12 bits of data.

11

Thus, the 16-megapixel camera produces on the order of 24 million bytes of data. The actual picture files saved by the camera are vastly smaller, typically about 10% as large yet, when reproduced, they appear to us to be complete. Each time we take a picture, the data from the camera sensor is run through a compression algorithm so that we are not obliged to save enormous files consuming correspondingly enormous amounts of memory or storage space.

As a general rule, image compression algorithms work by removing redundancy while preserving the portions of the picture that we find informative. When the data are uncompressed for viewing, the missing information is filled in with an educated guess at the missing features. Following on the example of our eye's preference for edges, image compression algorithms preserve the edge data, but are less accurate in preserving regions of near uniform color. It becomes the job of our brain to fill in the missing pieces.

Not surprisingly, standard cameras do not capture infrared or ultraviolet light, nor do they capture x-rays or gamma rays. Why should they? These are things our eyes don't see, so the pictures we take look perfectly natural to us without them—even though a large portion of the light energy coming from the objects in our photos is actually in those invisible ranges.

This might beg an intriguing question:

If we need scarcely 10% of the information that our cameras acquired in order to create a perfectly satisfying image, why did we collect the other 90% in the first place? At first, the question itself might sound absurd. Surely the answer is that we simply didn't know which of those 16-odd million pixels we needed until we had them in hand. Amazingly, the correct answer is rather different. A set of methods collectively known as Compressive Sensing (reviewed in Baraniuk) make it possible to acquire the compressed image directly. Roughly speaking, this can happen because even before we snap the photo we know something about what is in it. After all, it is a picture of something and not a picture of nothing. But we know much more. Natural images have a large number of statistical regularities that we can exploit, and the more we know about what an image contains, the less densely we have to sample it. For example, if we know in advance that we will be handed a banana or a grape, we hardly need anything more than the weight to determine what object we've been given. This is described diagrammatically in Figure 04 (itself adapted from an idea of Michael Lustig).

12

The efficacy of compressive sensing in digital imaging has been proven repeatedly and, over the next few years, many if not most consumer imaging devices will take advantage of the method. By using compressive sensing the costs of such gadgets will be reduced greatly as fewer components are needed. In many ways this process of engineering efficiency is much like evolution: the art of engineering is, in large part, to remove that which doesn't contribute and to preserve that which does. Brilliant engineering almost always results in simplicity. Brilliant design does so as well, and our aesthetic preference for the simple likely shares a deep link, but this might be the topic of a future essay.

The most efficient image sensors are specifically those that collect only (but all) of the information that is salient to us as observers. They are matched to our physiology, detecting the gamut of colors, light intensities, and spatial details that our eyes and brains can see,

and nothing more.

OUR BRAINS ARE

Since the early 1990's it has become possible to use tools developed for medical diagnostics, particularly functional magnetic resonance imaging (fMRI), to observe the actual activity of the human brain in tremendous detail (Cohen and Bookheimer). We can observe just which parts of the brain are involved when people perform tasks from observing flashing lights to considering racial identity, or to creating dreams. These patterns of brain activity are sufficiently reliable that we can, for example distinguish if someone is thinking about a game of tennis, or imagining themselves walking about their house (Monti, Coleman and Owen).

13

In a technology known now as "brain reading," we collect a catalog of fMRI images showing the brains of people performing a variety of different tasks. Following this, computers are used to provide statistical comparisons between new pictures and those in the catalog. If a match is found, we can conclude that the new image is of a person performing the same task as observed previously in the catalog image, much as computers can be trained to detect the identity of faces in photographic images.

The individual fMRI images are made up of tens of thousands of data points ("voxels") covering different locations within the brain. It has been known for some time, however, that the statistical tools of computational pattern detection fail if they are asked to operate on too many input points at once, a problem known as "the curse of dimensionality." One means to reduce the number of input points is to detect any redundancy among the various locations, and to operate on only non-redundant information. A convenient tool for doing so is called "principal component analysis" or PCA. Following PCA, the data become more sparse, and therefore more amenable to statistical pattern analysis. While PCA has certain attractive features however, there are many other means that sparsify the data either differently, or to a greater extent.

SPARSE

The human brain is organized into (at least) several hundreds of local regions that are highly specialized by task, as noted previously. In general, these regions interact with one another as part of a functional system or network. For example, all of the regions described in Figure 03 are parts of a visual network. Each individual region can participate in multiple networks, depending on the ongoing cognitive or behavioral task at hand. In our lab, we have been interested in the use of a sparsifying tool known as "independent components analysis" or ICA (Bell and Sejnowski). We, and many others, have observed that ICA has the apparent power to detect latent functional networks of the brain. For example, following ICA the brain images might select patterns of activity in a set of regions involved in language, a set of regions involved in visual thinking, or a network of regions implicated in emotional responsiveness (Figure 05).

14

We are attracted to the idea that even very complex human behaviors might be identified by the contributions, at any moment in time, of each of a set functional brain networks. Reading, for example, might be expected to engage a visual system, a network for controlling the eyes, and a network for parsing grammar (as well as others, of course). To address the problem of "the curse of dimensionality" perhaps we could create a decoding dictionary whose words (the behaviors we would like to identify) are made up of letters (the elements that are composed into those words) that themselves are the networks nominated for us through ICA. Graphically, we can image a cognitive state as a location in a space whose coordinates are the relative activity in each of these functional networks (Figure 06).

Remarkably, we can make very subtle distinctions among cognitive states in this manner, by observing the activity of only a very small number (e.g., 40) of networks Figure 07, for example, shows that examining as few as 19 networks made it possible to determine with better than 80% accuracy whether or not subjects believed or disbelieved in a series of truth propositions (Douglas et al.).

is the WORLD

We can represent complex ideas, emotions and thoughts with just a few lines of ink. Our brains are satisfied to draw life and death conclusions from just a few senses. A camera can take a high resolution picture from only 10% of the data. What does all of this tell us about the nature of the world, and about reality?

SPARSE?

15

Perhaps something about the world itself is sparse? Perhaps all data, all measurements, all observations are linked in a deep network that can permit reality to occupy only a small fraction of theoretical states. This is a tantalizing proposition, and is one for which a geometry and mathematics might well be developed. In fact, a related concept was put forward by David Layzer in a 1975 article in the popular magazine, Scientific American (Layzer).

Alternatively, I would propose that it not the world that has this sparsity of states, it is our minds, and it is our minds specifically as a result of the limited trajectory of evolution in endowing us with only very specific brain abilities. The concepts and ideas that we can hold are constrained heavily to the simple sensory and motor processes that in the past served us to escape being eaten or to help us feed or reproduce. As humans we have enormous hubris, and at times imagine that our minds have limitless potential to hold and to consider ideas unbounded by the mere physical stuff of which we are made and in which we live. Instead, I feel increasingly that the content of our thoughts only rarely can extend beyond the contents of our senses.

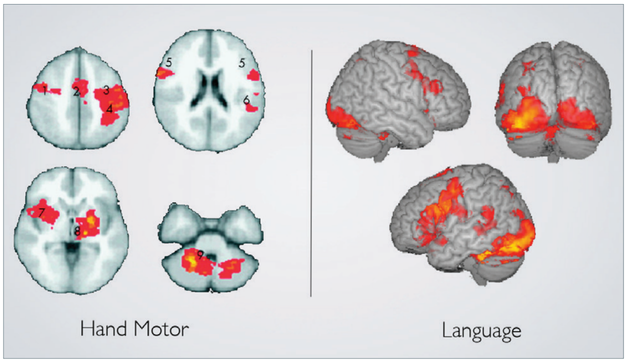

05

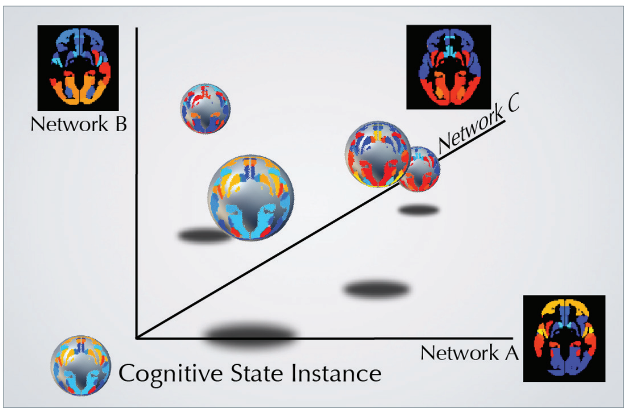

06

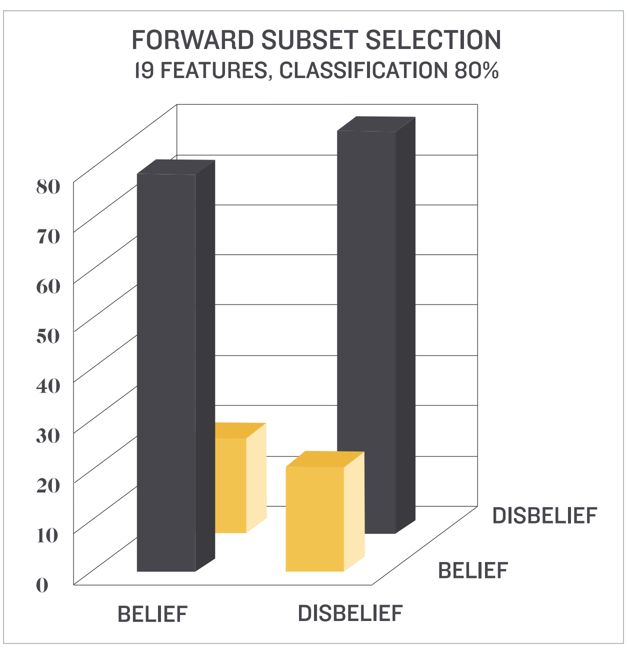

07

05 brain networks

Neuroscientists believe that specialized regions of the brain participate together in functional networks to perform complex tasks. The organization of these networks is dynamic and depends our current behavior or thought processes.

06 cognitive states in 3D

For the purposes of brain reading—identifying the cognitive state of an individual by observing their brain activity, we can conceive of each different cognitive state as having a location in a space whose dimensions are the relative activities of different functional networks in the brain. If these patterns are sufficiently unique, it is necessary only to measure the activity within each network to determine the cognitive state.

07 belief detector

A brain reading method, based on the observations of the activity in just 19 brain networks, was able to distinguish with better than 80% accuracy where an individual believed, or disbelieved, in a propositional statement, such as "Sugar is Sweet" or "I own a toaster oven."

16

references

Baraniuk, R. G. "Compressive Sensing." Ieee Signal Processing Magazine 24.4 (2007): 118-+. Print.

Bell, A. J., and T. J. Sejnowski. "An Information-Maximization Approach to Blind Separation and Blind Deconvolution." Neural Comput 7.6 (1995): 1129-59. Print.

Cohen, M. S., and S. Y. Bookheimer. "Localization of Brain Function Using Magnetic Resonance Imaging." Trends Neurosci 17.7 (1994): 268-77. Print.

Douglas, P. K., et al. "Performance Comparison of Machine Learning Algorithms and Number of Independent Components Used in fMRI Decoding of Belief Vs. Disbelief." NeuroImage 56.2 (2011): 544-53. Print.

Hartline, H. K., H. G. Wagner, and F. Ratliff. "Inhibition in the Eye of Limulus." J Gen Physiol 39.5 (1956): 651-73. Print.

Layzer, David. "The Arrow of Time." Scientific American 233.6 (1975): 56-69. Print.

Monti, M. M., M. R. Coleman, and A. M. Owen. "Neuroimaging and the Vegetative State: Resolving the Behavioral Assessment Dilemma?" Ann N Y Acad Sci 1157 (2009): 81-9. Print.