Understanding the brain imaging signal

Bruss Lima - The Italian Academy at Columbia University
Laboratory of Aniruddha Das

1 - INTRODUCTION

Our perceptions, sensations, emotions and memories rely on the approximately 80 billion neurons we have in our brain. You may wonder, nevertheless, if this number is actually too large, too small or barely enough. Is it the "billions" that makes us special? Or is there more to it than just a number, some neural mechanism that we have not yet envisioned? Cognitive Neuroscience has greatly advanced our understanding of the brain over the past decades. Even so, one may be utterly frustrated by the lack of any unifying theory explaining how we perceive the world around us, let alone understanding complex processes such as consciousness.

Some neurophysiologists (like myself) are interested in understanding the "language" of the brain. Putting it in terms of a jargon commonly used, we are interested in understanding the "neural code". That is not a given. I have heard prominent neuroscientists arguing that the brain has no such thing as a "code". Consequently, there is no code to be "deciphered".

Let us get rid of the jargon and simply assume that the brain employs certain mechanisms in order to perceive the world around us, eventually generating "relevant" behavioral responses. For example, suppose you are at home and someone knocks at your door. The knock on the door might be a complete surprise, specially if you are engaged reading a book. You will nevertheless process the auditory stimulus, interpreting it as someone waiting outside. As behavioral response you will close your book and walk towards the door. On your way you may begin to

build some expectation of whom it might be. Yes, you think to yourself, it must be the postman. As you pull the door open you orient your head upwards since you are expecting to see a tall man dressed in blue. But there is no one at the door! In the split of a second you reorient your assumptions and look downwards. There is a package for you left on the doormat. The postman knocked but did not wait.

This may be the description of a mundane event, but no one would deny that the neural processes behind it can be quite complex. Actually, it is quite challenging to even simulate a behavioral task of this complexity in a laboratory setting. We therefore simplify our tasks in the lab to an astonishing level so that we are still able to technically measure brain activity while the task is being executed. And, of course, still be able of interpreting the results. Technical problems aside, I would like to highlight some important aspects of the behavioral task described above that will be important for our discussion. Perception is usually depicted as a passive phenomena. In textbooks, for example, the eye is commonly represented as a camera collecting information about the world. Hearing and being startled by the knock on the door, as described on the anecdote above, might be the closest you get to passive perception. Otherwise, we are constantly making predictions about the world and about what we will see. Notably, these predictions can strongly influence what we will experience and when we will experience it. Back to our anecdote, your expectations will surely not be expressed in the same way as you approach the door. Rather, it will probably escalate to a maximum as you swing the door open.

We can describe perception as having roughly two components. The first and most obvious one is the external component, the stimuli coming from the world around us that drive our senses.

The second component of perception relies on internal processes, such as anticipation, expectation and attention. Both these components act together to enact our experiences.

Classical neuroscience has been strongly influenced by behavioral psychology. The latter focuses on observable behavioral responses to controlled external stimulation, not on "unaccessible" neural phenomena taking place in the "mind". This approach has been powerful and has taken us great lengths in understanding how the world is represented in our brain, as well as how our brain represents future motor actions upon the world. It has brought great promises of making blind people see and paralyzed patients walk using prosthetic devices, as long as we sufficiently understand the correct neural representation. In fact, the term representation here is so central that we will revisit its implications further below.

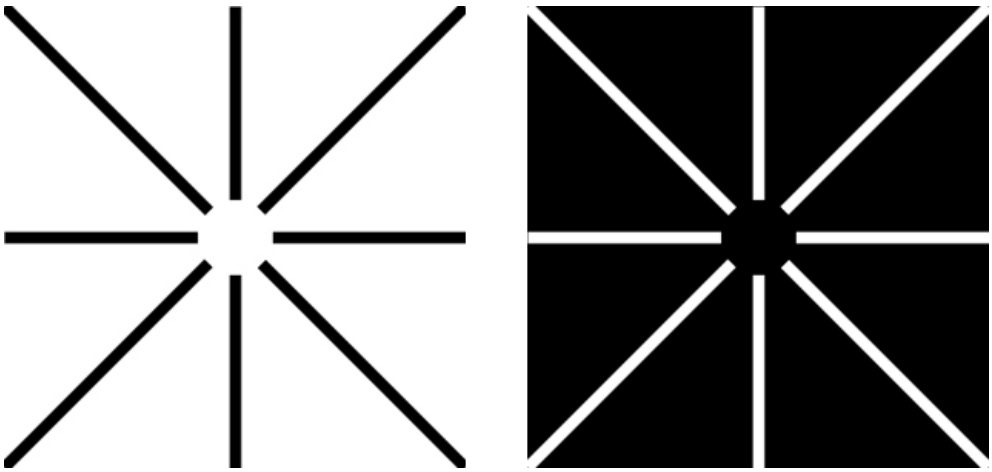

Figure 1. Ehrenstein illusion. The alignment of the line segments generate the perception of a circle in the middle of each figure. Note that you even perceive a difference in luminance inside the illusionary circles relative to their outside region.

The emergence of Gestalt psychology around the beginning of the 20th century offered an alternative approach to treating the mind as a simple "black box", and has inspired several other lines of thought centered on the active role of perception. The visual illusion in Fig. 1 supports

that notion that the brain constructs the visual image based on its internal structure and not necessarily on the physical characteristics of the image. Scientists began to consider the internal ongoing dynamics to the brain as playing a fundamental role in perception, the default-mode network being one of its most recent expressions. My general scientific interest lies in understanding the interplay between external and internal components of perception. What are their mechanisms and how do they interact?

2 - NEURAL SIGNALS

2.1 - Spikes

When I started working with neuroscience I was exposed to the most reductionist approach of all for understanding the brain: recording from a single neuron in the monkey cerebral cortex (Fig. 2).

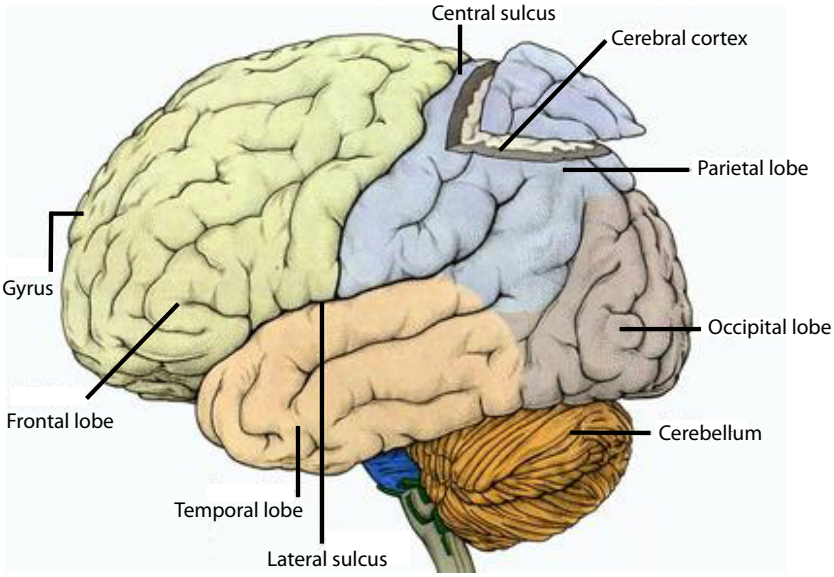

Figure 2: Illustration of a human brain. The cerebral cortex, with its typical grooves and folds (gyri), is the outermost tissue covering a great part of the brain. The excision on the upper part of the figure exposes the thickness of the cerebral cortex (approximately 2 mm). It contains 6 layers in most of its extent, which are disposed parallel to the surface (adapted from Wikispaces).

Despite the fact that the monkey's cerebral cortex has approximately 500 million neurons, the single cell approach has lent us a tremendous deal of knowledge concerning the brain. A neuron is the basic functional unit of the nervous system. In its most common form, it is composed of a cell body, several dendrites and an axon (Fig. 3).

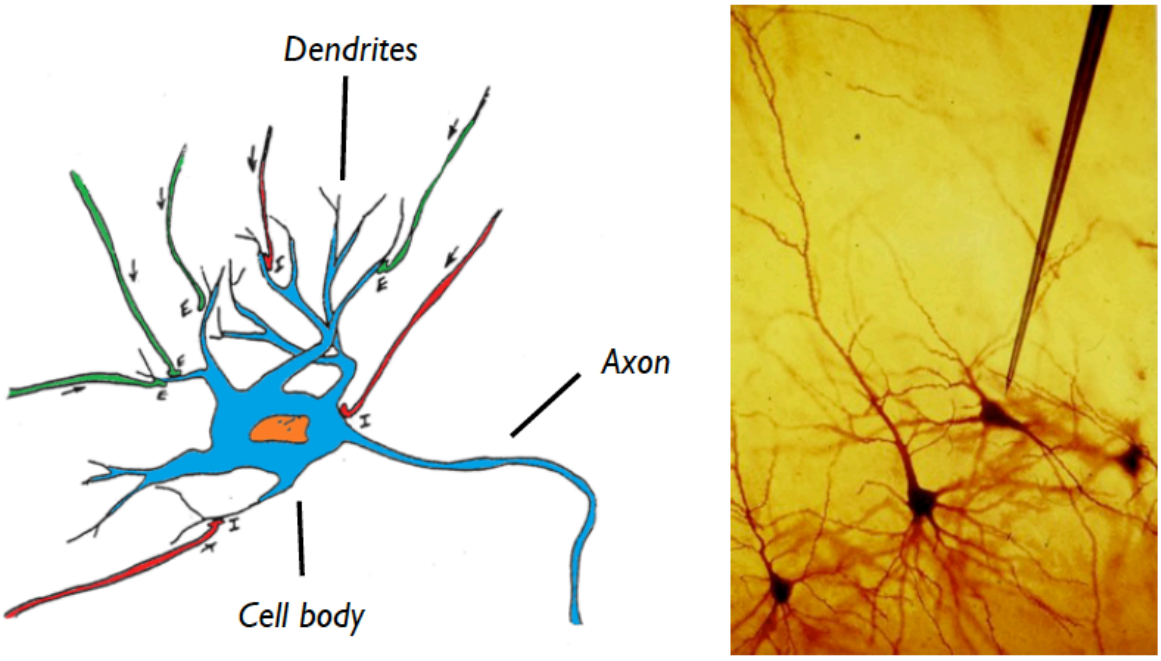

Figure 3: Left. Schematic illustration of a neuron with its cell body, axon and dendrites. The green and red segments represent axons arriving from other neurons. They usually target the dendrites and form specialized connections called synapses. Synapses are believed to be central to learning processes since they ultimately establish the connectivity between neurons. "E" and "I" stand for excitatory and inhibitory synapses, which means that they either increase or decrease the membrane potential of their target cell body, respectively, by producing excitatory or inhibitory post-synaptic potentials (EPSP or IPSP). Thousands of these potentials are being constantly generated in every cell. When the overall potential of the target membrane exceeds a certain voltage threshold, a all-or-non spike event is generated (see Fig. 4). Once generated, spikes propagate through the axon towards other neurons, sometimes meters away. Right: Image of real neurons with a nearby microelectrode recording their activity. Electrodes used to record electrical activity outside neurons (extracellular recordings) are usually made of metal. They are coated with non-conductive material on their surface leaving only a tiny tip of metal exposed. It is possible to record the activity of a single neuron if its tip is sufficiently close to a cell body. Activity of several nearby neurons can be recorded when the electrode tip has a larger surface area (multi-unit activity or MUA).

Figure 4: Trace of a single spike recorded by Hodgkin and Huxley in 1939. Spikes usually have this stereotyped shape and are recorded as such anywhere in the brain.

Neurons are specialized in transmitting information between different regions by generating electrical events called spikes (Fig. 4). Spiking activity are all-or-non events that once triggered they propagate through the axons until they arrive at their target, usually the dendrite of another neuron. Microelectrodes placed outside cell bodies measure the net electrical activity of the extracellular space. This signal has a typical magnitude of several hundred microvolts. The signal arriving from the electrode needs to be amplified thousands of times in order to be heard and recorded. Such tiny signals have the disadvantage of carrying noise with them, usually 60 Hz from the electric outlets. If you are in Europe, this number will be 50 Hz.

As I mentioned above, study of single neurons has very much shaped the way we understand the nervous system. Placing microelectrodes in the extracellular space, as the one shown in Fig. 3 (right panel), has enabled scientists to understand the response properties of neurons in different brain regions. The procedure consists in testing a variety of physical stimuli and observing in which situation the neuron responds best (i.e. when it generates the highest amount of spikes within a period of time). In other words, one tries to find the neuron's response characteristics and feature selectivity. Ironically, chance has guided many of these discoveries.

One of the first response characteristics to be described for sensory neurons was the concept of receptive field. For a visual neuron, it consists in finding the portion of space to which the neuron responds best. Retinal ganglion neurons, for example, respond to tiny portions of visual space (Harline, 1938). Additionally, these neurons are poorly sensitive to homogenous illumination. They rather prefer light spots surround by darkness, or dark spots surrounded by light (Kuffler, 1953). This property introduced an important concept in visual processing: relative contrast is far more important for information encoding than absolute luminance. In the primary visual cortex (V1), the region I am currently working on, neurons respond maximally to oriented contours (e.g. bars), and receptive fields size is several times larger than the ones in the retina. Another characteristic of V1 is that its receptive fields are organized in an orderly manner forming a map of the world (Fig. 5).

As you may notice, the response features of a V1 neuron are substantially more complex than those of the retina (response selectivity to oriented contours versus spots of light). We will continue to observe this trend as we ascend in cortical hierarchy. By the time we reach the temporal cortex (see Fig. 2), neurons there are capable of responding to complex features such as specific faces. It is unavoidable to think that the brain builds a representation of the world. It may start simple in the retina and in V1, but it will eventually get complex. Note that we are talking about response properties of single neurons. If a neuron selectively responds to an object, does that mean that a single neuron holds a "concept" of that object? Additionally, do we depend on concepts of single neurons in order to construct our "own" concepts?

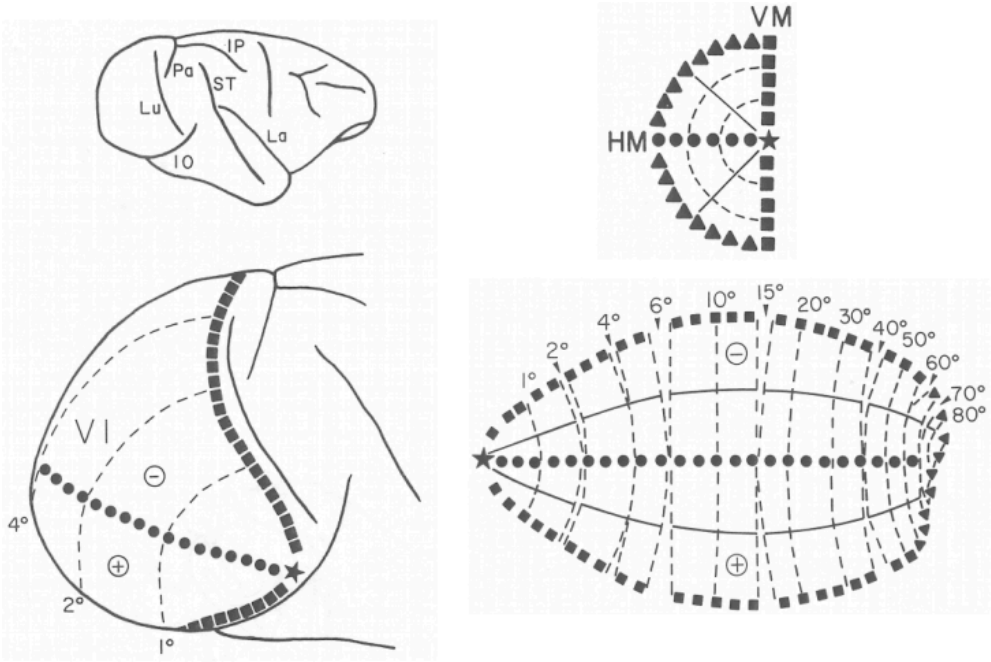

Figure 5: Topographic organization of the primary visual cortex (V1). Left: schematic representation of a monkey brain (top) with the delimitation of cortical area V1. Top right: representation of the field of view in the external world. The star symbol represents the point in the visual scene where you are directing your gaze to. VM and HM correspond to vertical and horizontal meridians, respectively. Bottom right: the representation of the external visual world in V1 using the symbology employed in the top right. Note that approximately half of V1 surface is devoted to the representation of the central 5 degrees of the visual field, implying that a huge proportion of cortical tissue is dedicated to the processing of stimuli presented at our center of gaze. Adapted from Gattass et al (1987).

When asked by his students how neurons represent objects, neuroscientist Jerome Lettvin came up in 1969 with a humorous and fictitious story that has gradually become a serious theme in neuroscience circles (Gross, 2002). The main personage of his tale was his second cousin, Dr. Akakhi Akakhievitch, a neurosurgeon living in the distant Ural Mountains. Dr. Akakhievitch was convinced that single neurons were capable of representing ideas. Ever more obsessed, he decided to search for those neurons representing the most primitive of all objects: one's mother. He was able to find some 18,000 neurons in the human brain that responded exclusively to a mother, which ever way she was presented. Not only was Dr. Akakhievitch very pleased with his

breakthrough, he also found immediate practical application for his discovery when a patient named Portney entered his office. Portney was psychologically tormented by problems involving his own mother, and Dr. Akakhievitch was keen to help. After being led to the operating table and having each one of his 18,000 mother neurons ablated, all his problems were finally solved: he lost the concept of his mother. After solving the enigma of the mother cells, Dr. Akakhievitch then turned to his next great challenge: the grandmother cells.

In 1972, Horace Barlow proposed the “neuron doctrine” which linked sensation directly to the activity of single units in the brain. In part, Barlow’s proposal offered a significant paradigm shift. At the time, the activity of individual cells was considered too unreliable to be examined singly. Therefore, many turned to more macroscopic manifestations of neuronal activity, such as the electroencephalogram (EEG), in order to understand behavior. Barlow, on the other hand, was able to capture and synthesize a thinking that would dominate neurophysiology for many years to come: that single neurons are the prime movers leading to our perceptions and sensations (Hubel, 1982). Not every cortical neuron’s activity had necessarily a simple perceptual correlate, but active high-level neurons directly and simply caused the elements of our perception. The mechanism by which this was achieved consisted in the frequency of neural impulses. Barlow argued that the rate of firing coded the certainty that the cause of a percept was present in the external world.

A pertinent criticism generally applied to the single unit doctrine is the one commonly referred to as “combinatorial problem” or “combinatorial explosion” (Singer and Gray, 1995). We are constantly confronted with unlimited combinations of elements and features, and yet we face no problem in recognizing patterns and objects presented in our visual scene. Additionally, we can

easily cope with patterns not previously exposed to us, and we can effortlessly recognize objects encountered in unusual orientations and sizes, under different illumination conditions, or partially occluded by other objects (Logothetis and Sheinberg, 1996). Contrary to what is usually assumed, however, Barlow did not believe in the existence of “grandmother cells” (Barlow, 1972). He acknowledged that there would not be enough “grandmother cells” or, as he called them, “pontifical cells” to account for the enormous variety of our perceptions. He believed that the activity of a single neuron would not be able to convey the richness of our sensations, which consist on the combination of our various percepts. Barlow, therefore, proposed the term “cardinal cells” instead of “pontifical cells” to explain how the brain represented the world. The interaction of various “cardinal” cells, which are lower in the hierarchy and more numerous than “popes”, would be responsible for the emergence of our perceptions and sensations.

2.2- The local field potential (LFP)

Extracellular recordings pick up two basic signals. One is obviously the spikes generated by single neurons, which we discussed previously. Spikes have a very short duration (around 1 milli-second) and can be heard as clicks after amplification. The second signal component, much slower than spikes, is the local field potential (LFP). Depending on the electrode properties, the LFP can express the average behavior of either tens, hundreds or thousands of neurons (Logothetis, 2003a, 2003b). The LFP has a similar nature to the one recorded with electroencephalography (EEG). Berger (1929), who performed pioneering studies in human EEG, first described the 10 Hz oscillation while recording electric potentials from the scalp of his 15 year old son Klaus. This rhythmic activity, initially named Berger waves and currently known as alpha waves, was typically observed in subjects under rest (sitting quietly with eyes closed).

Engagement into “active” behavior caused dramatic changes in the brain waves. Simply opening the eyes or solving math problems with the eyes shut shifted the oscillation frequency to the (higher) beta range. Neuronal oscillations show a wide range of frequencies, ranging from approximately 0.05 Hz to 500 Hz (reviewed in Buzsáki and Draguhn, 2004).

The LFP is a measure of synchronized synaptic activity originating from the population of neurons around the electrode tip. Since it reflects an average population activity, it is specially sensitive to coordinated fluctuations in membrane potential. Synchronous extracellular currents, resulting from EPSPs and IPSPs (see legend of Fig. 3) are believed to be the main components contributing to the generation of the LFP (Mitzdorf, 1987). The contribution of spiking activity to the LFP is believed to be negligible during normal brain activity. The reasons for this are twofold. Even though EPSPs and IPSPs have much lower amplitude as compared to spikes, they have a much longer duration (up to tens of milliseconds). This gives the postsynaptic potentials a higher chance of overlapping and summating with one another than is the case for spikes. Second, only a small number of neurons reach spike threshold at any given moment, while postsynaptic events are much more frequent. Therefore, features observed in the LFP are mainly the result of coordinated subthreshold postsynaptic processes.

The oscillatory cycle can be divided in two phases: the depolarizing and the hyperpolarizing phase. EPSPs arriving during the depolarizing phase are much more likely to elicit spikes in the postsynaptic neuron as compared to EPSPs arriving during the hyperpolarizing phase (Fries, 2009). Therefore, the main impact of oscillations on the membrane potential would be to coordinate the neuronal excitability of sub-populations of neurons.

Note that this network characteristic of the LFP places it in a special category compared to the activity of single neurons. Regardless if a single neuron is responding or not to someone's "grandmother", its activity can now be considered embedded within larger network fluctuations which are coordinated in both space and time. This is a completely different context to the one where a single unit, at the topmost of the information processing chain (e.g. a "grandmother cell"), gives rise to complex thoughts and concepts.

3 - THE HEMODYNAMIC SIGNAL

At the lab of Aniruddha Das, Columbia University, I have been recording the intrinsic optical imaging signal of the cerebral cortex. This technique allows you to investigate much larger portions of the brain, but with one important caveat: it measures brain activity indirectly by changes in local blood oxygenation and volume (hemodynamics). You may think this is a drawback, but recent data from the lab (Sirotin and Das, 2009) suggest that this signal carries information that is not directly available in the underlying neural activity. One may appropriately ask why that is the case. Information on hemodynamics would thereby compliment the information we gather using direct neuronal measurements, such as single neuron and LFP recordings. I believe the combination of recorging techniques can be a powerful tool in understanding the neuronal "code" mentioned earlier, or at least be in a better position to disprove its existence. There is another important reason why optical imaging is a relevant technique. It is the optical analog of Function Magnetic Resonance Imaging (fMRI), a widely used non-invasive technique to study the human brain. The signal measured with fMRI is commonly referred to as the BOLD signal (blood-oxygen-level dependent contrast imaging). But in essence

it is a hemodynamic signal of the same nature as the one recorded with optical imaging. There are still many open question regarding the BOLD signal, and optical imaging is a strategic technique to address them. We will later revisit the nature of the hemodynamic signal.

After introducing you to three techniques used in measuring brain activity I am in a better position to explain you the work we carry out in the lab. We simultaneously record these three signals (single neurons, LFP and hemodynamic recordings) from the primary visual cortex of an awake monkey trained to perform behavioral tasks. Our major goal is to understand the nature of the hemodynamic signal and its correlation with the underlying neural activity. The primary visual cortex (V1) is located on the posterior pole of the occipital lobe (rightmost portion of the brain illustration in Fig. 2). It is the first cortical stage of visual computation and therefore processes early visual information. We could have chosen another visual area. The advantage of V1 is that it is large and is located on the surface of the brain, not hidden in a sulci, lending it accessible to optical imaging.

As mentioned above, the lab has recently described a component of the hemodynamic signal which is poorly correlated to local neural activity (Sirotin and Das, 2009). This signal was named Trial-related Signal or Task-related Signal (TRS). In some ways this result was considered bad news by the fMRI community. The underlying assumption was that the BOLD signal was a indirect reflection of neural activity. The assumed link between BOLD and neural activity made the fMRI technique a very convenient non-invasive tool to study human cognition.

That is not to say that the hemodynamic signal has no correlation with underlying neural activity and external stimulation. We know it does since the initial discovery of the intrinsic optical imaging signal. The link between blood supply and neural activity lies in the fact that the latter is

a metabolically demanding process. To be sustained it requires an adequate supply of oxygen and nutrients (i.e. sugar in the form of glucose). Our muscles can function for prolonged periods of time with low oxygen. That is not at all the case for the brain, which is particularly sensitive to hypoxia (lack of oxygen). Increases in brain activity are tightly linked to increased blood supply, and thereby oxygen and sugar concentration. Importantly, blood supply changes can be directed to the specific brain regions that are metabolically active. Both optical imaging and fMRI explore this characteristic in order to infer local brain activation from hemodynamics. This is possible due to certain characteristics of the blood. Red blood cells get their color from the iron present in a protein (hemoglobin) they carry. Hemoglobin transports oxygen from the lungs to the rest of the body. Oxygen chemically binds to the iron present in the hemoglobin in order to be transported. Interestingly, hemoglobin undergoes a slight change in color when oxygenated. Therefore, by measuring the light absorption of blood vessels in the brain one is able to quantify blood volume and oxygenation level. That is how optical imaging is able to infer brain activity from hemodynamics. Hemoglobin also changes its magnetic properties when oxygenated. fMRI explores this characteristic in order to make the same inference as the one made with optical imaging.

The initial experiments using optical imaging we mostly conducted on animals under anesthesia. Internal brain processes present during conscious states were probably extinguished in such situation. Similar experiments, when performed in awake animals, reveal additional components of the hemodynamic signal not observed under anesthesia. When the monkey is performing a task which is predictable in time, part of the hemodynamic signal seems to follow the periodicity of the task, even when the animal is working in darkness, where no V1 response is expected

(Fig. 6). These results lead us to a tempting speculation (Moore and Cao, 2008). Neural activity requires blood supply in order to be maintained. If the brain "knew" it would engage in higher activity at a predictable moment in time, could it anticipate blood supply to the relevant region? The task-related signal, a possible correlate of internal brain states, could be doing something similar to that. At the moment it is still a speculation, but the lab is currently addressing this question.

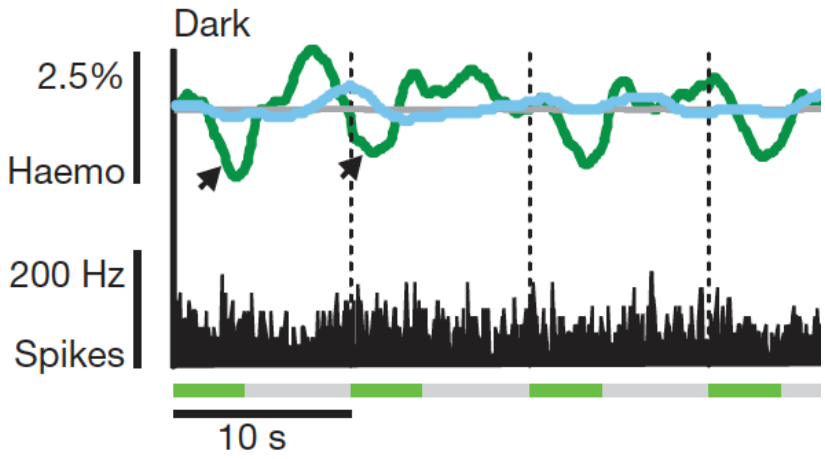

Figure 6: The task-related signal. Simultaneous recording of hemodynamic (dark green trace on the top) and spike activity (black trace on the below). The experiment was performed on cortical area V1 in the awake monkey. The task was very simple and consisted in periodically looking at a computer screen (during periods represented by light green bars shown on the bottom) in order to receive a juice reward. Note that this is a temporally predictable task since it has a constant length (10 seconds). Since external visual stimulation is mostly absent (the monkey is working in a dark room) no modulation of spike activity is observed (the black trace below is mostly flat). However, the hemodynamic signal fluctuates as a function of trial length. The black arrows indicate decreases in the optical imaging signal at the beginning of the trial. Adapted from Sirotin and Das (2009).

The hemodynamic signal seems to carry information both of the external world and of internal process. Do these two signals interact in some complex, non-tractable way? We recently described that these two components add to each other in a simple linear fashion (Cardoso et al, 2012). This is again good news to the fMRI community wishing to relate BOLD to underlying

neural activity. The only requirement is to be able to correctly predict the trial-related signal and subtract it away. Doing that you are left with the well-behaved stimulus-related signal, but which still has poses challenges. We know it is associated with neuronal activity, but which aspect of it?

A pioneering work addressing this question was performed by the group of Logothetis and collaborators (Logothetis et al, 2001). The most appropriate test is to simultaneously measure neural and hemodynamic activity and directly access their correlation. In a heroic experiment they were able to record neuronal activity from an anesthetized monkey inside a fMRI scanner. They observed that the LFP was the best correlate for the hemodynamic signal. This result has been particularly influential in interpreting fMRI data. The authors argued that the LFP, more than spikes, reflects local brain metabolism.

During my recent work in the lab I have revisited this issue. We performed the analogous experiment using optical imaging instead of fMRI. To our advantage, we recorded from the awake instead of the anesthetized monkey. Additionally, we decomposed the hemodynamic signal into stimulus-related and task-related components. We tested specifically the correlation between neural activity and the stimulus-related imaging component. Contrary to Logotheis and collaborators we observed that spikes, and not LFP, provide the best correlate to the hemodynamic signal.

4 - REFERENCES

Barlow HB. 1972. Single units and sensation: A neuron doctrine for perceptual psychology? *Perception.* 1: 371-394.

Berger H. 1929. Über das Elektrenkephalogramm des Menschen. *Arch. Psychiatr. Nervenkr.* 87: 527-570.

Buzsáki G, Draguhn A. 2004. Neuronal oscillations in cortical networks. *Science.* 304: 1926-1929.

Cardoso MM, Sirotin YB, Lima B, Glushenkova E, Das A. 2012. The neuroimaging signal is a linear sum of neurally distinct stimulus- and task-related components. *Nat Neurosci.* 15: 1298-1306.

Fries P. 2009. Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu. Rev. Neurosci.* 32: 209-224.

Gattass R, Sousa APB, Rosa MGP. 1987. Visual topography of V1 in the *Cebus* monkey. *J. Comp. Neurol.* 259: 529-548.

Gross CG. 2002. Genealogy of the "Grandmother Cell". *Neuroscientist.* 8: 512-518.

Hartline HK. 1938. The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *Amer. J. Physiol.* 121: 400-415.

Hubel DH. 1982a. Cortical neurobiology: a slanted historical perspective. *Ann. Rev. Neurosci.* 5: 363-370.

Kuffler SW. 1953. Discharge patterns and functional organization of mammalian retina. *J. Neurophysiol.* 16: 37-68.

Logothetis NK. 2003a. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J. Neurosci.* 23: 3963-3971.

Logothetis NK. 2003b. MR imaging in the non-human primate: studies of function and of dynamic connectivity. *Curr. Opin. Neurobiol.* 13: 630-642.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412: 150-157.

Logothetis NK, Sheinberg DL. 1996. Visual object recognition. *Ann. Rev. Neurosci.* 19: 577-621.

Mitzdorf U. 1987. Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. *Int. J. Neurosci.* 33: 33-59.

Moore CI, Cao R. 2008. The hemo-neural hypothesis: on the role of blood flow in information processing. *J. Neurophysiol.* 99: 2035-2047.

Singer W, Gray CM. 1995. Visual feature integration and the temporal correlation hypothesis. *Ann. Rev. Neurosci.* 18: 555-586.

Sirotin YB, Das A. 2009. Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. *Nature.* 457: 475-479.