CS/NEUR125 Brains, Minds, and Machines

Assignment 5: Neural mechanisms of object-based attention

Due: Friday, April 14

This Assignment is a guided reading of the 2014 paper, "Neural Mechanisms of Object-Based Attention," by Daniel Baldauf and Robert Desimone at MIT. Reading this article will prepare us to discuss this work during our fifth Journal Club in class on Tuesday, April 18.

To begin, create a copy of this Google document and modify the title of the copy to include your name. Questions that you should submit answers to are shown in blue. As with labs, you'll turn in this Assignment by sharing your copy of this Google document with Ellen and Mike.

This paper by Baldauf and Desimone is again a primary research article. As is typical of reports in the high-impact journals *Science* and *Nature*, the article is brief. Partly for this reason the authors are likely to assume knowledge in the reader, or leave out details that are familiar to experts. We'll try to fill in some of the assumed knowledge with this document, but we again have to accept that we may not be able to digest and understand every line in this technical paper. Our goal is to appreciate the basic results of the authors' MEG and fMRI experiments aimed at revealing how our brains pay attention to specific objects when they overlap with other objects in the same part of the visual field.

As usual, if you use phrases from the paper in answering the questions, you must put them in quotation marks, and you should try to reformulate the idea in your own words.

Because it's easy to get bogged down in technical details in a paper, we first want to understand what is the question or hypothesis the authors are trying to address with their study. That way you can try to relate everything else you read to answering that question--and if it doesn't help address the main question, you might be able to safely ignore it.

There are no subsections in a Science report aside from the Abstract. Let's call the Abstract the first paragraph and refer to sections of the paper by paragraph number. There are just 15 paragraphs total, plus the three figure captions.

Abstract and second paragraph

You might feel immediately "bogged down" in unfamiliar technical terms in these first couple paragraphs, like "gamma synchrony," "gamma phases," and "coupled oscillations" in the Abstract and "covert" and "top-down" attention in the second paragraph. Still, even before understanding those terms we can hang on to their statement of the basic question in the first sentence of the Abstract and the last sentence of the second paragraph.

Q1. What is the main question being addressed in this study?

Now let's start to get a basic handle on the unfamiliar terms before considering their specific experiment and their results.

First, **attention** usually refers to allocating more of some "limited cognitive resource" to processing certain types of sensory stimuli over others. So we can think of attention as a "bias" in favor of perceiving certain stimuli or types of stimuli over others. **Top-down attention** is basically another name for **voluntary attention**, where you direct attention to specific targets based on your own goals or a task you've learned. That's as opposed to **bottom-up attention**, a.k.a. automatic or involuntary attention, which refers to situations where properties of the stimulus itself "grab" your attention--like a loud noise or flashing light.

The "top-down" in "*top-down signals*" refers to neural signals from anatomically "higher" regions like the frontal eye fields (FEF) in a processing hierarchy down to lower level regions like sensory cortex. The top-down signals do appear to be involved in top-down attention.

"**Covertly attending** to a location in the periphery" means paying attention to that spot without moving your body, head, or eyes to look at that spot. So those are some terms and concepts related to the psychological notion of attention.

"Gamma synchrony," on the other hand, is a concept from neuroscience. "Synchrony" in this context is referring to the electrical activities of different neurons being synchronized so that they tend to fire or not fire together. In general, neurons that fire synchronously will have a larger impact on common target neurons than if they were to fire at random times. This is illustrated in the figure below, taken from a study of the neural correlates of *tactile* attention in monkeys by Niebur, Hsiao and Johnson (2002). The point is that when the presynaptic spikes are synchronized, the postsynaptic response will tend to be larger because the excitatory postsynaptic potentials (EPSPs) add up.

.

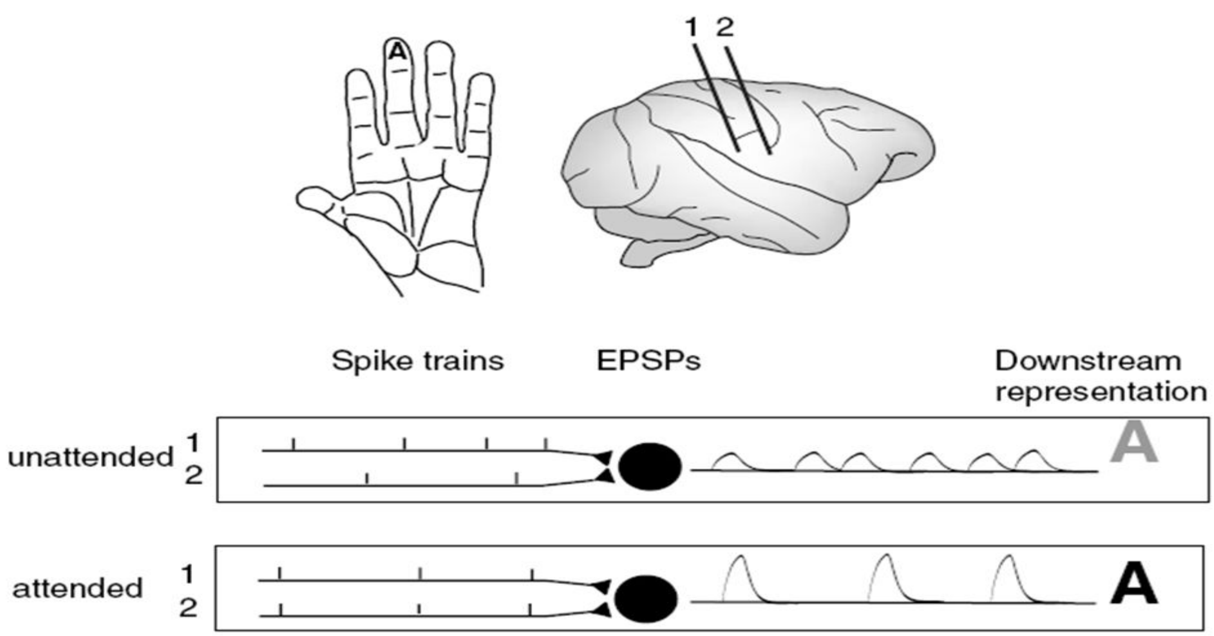

Gamma synchrony refers to groups of neurons that not only fire together, but do so *rhythmically*, in an *oscillation* with a relatively high frequency above 30 cycles per second (Hz). If gamma oscillations in two parts of the brain are “in synch,” we refer to “coupled oscillations” in the two areas and gamma synchrony between those areas. If the peaks of oscillation in one area consistently come earlier than in the other area, we say the “**phases were advanced**” in the earlier area[1].

There have been suggestions based on empirical evidence and theoretical considerations, that gamma synchrony (and/or neural synchronization at other frequencies) might play a role in routing neural sensory signals among brain regions. For example, gamma synchronization has been proposed as a neural mechanism of attention. The basic idea is that if neurons in a target region have rhythmic windows of excitability, inputs from another region will be more effective in causing the target neurons to respond if the input neurons are rhythmically active during the windows of increased excitability in the targets--that is, if the two populations are rhythmically synchronized, i.e. **coherent**. This idea is illustrated in the schematic below (from https://image.slidesharecdn.com/abnormalneuraloscillationsandsynchronyinschizophrenia-100308232949-phpapp02/95/abnormal-neural-oscillations-and-synchrony-in-schizophrenia-2-728.jpg?cb=1268091000).

The timing of rhythmic activity in cortical networks influences communication between neuronal populations.

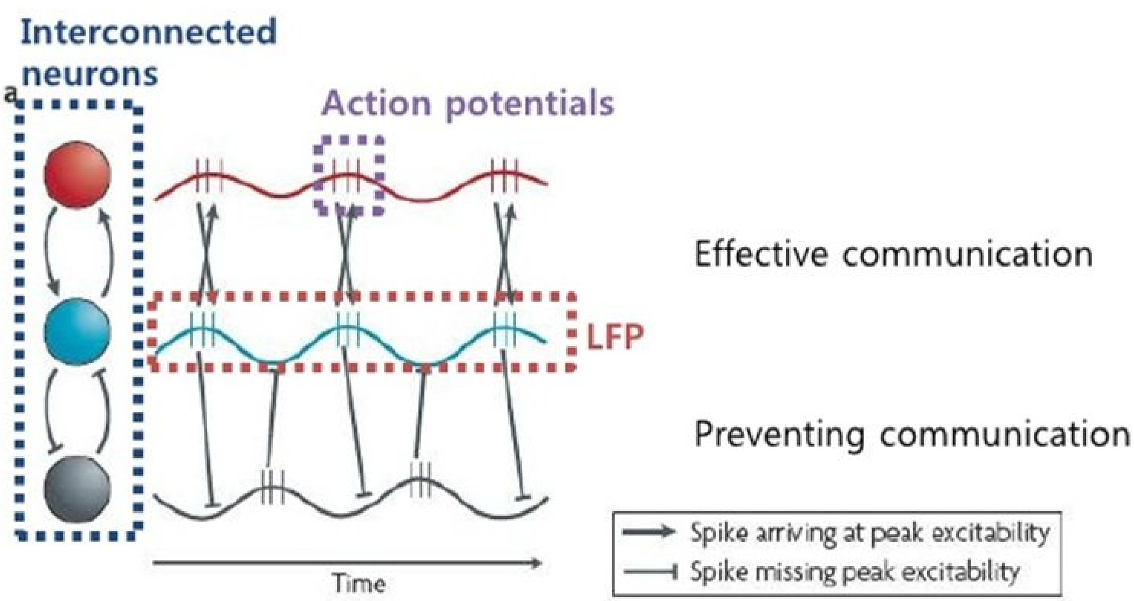

Q2. Match each area with its role as described in the Abstract.

___FFA A. Putative attentional area that directs visual information flow
___IFJ B. Area specialized for representing places
___PPA C. Area specialized for representing faces

[1] This is what they mean in the Abstract. Unfortunately the authors use “advancement” in the opposite sense in the caption of Figure 2J.

Paragraphs 3-4 and Figure 1: stimuli and attention

Q3. The authors say they used MEG and fMRI methods "to optimize both spatial and temporal resolution." Which of the two neuroimaging methods provides the better temporal resolution? (You can learn about MEG in this short (~ 3 minutes) video by Dr. John Gabrieli at MIT.)

Next the authors describe the stream of face and house images simultaneously presented to human subjects in their MEG experiments. They say the stream of faces and the stream of houses were "**tagged** at different presentation frequencies (1.5 and 2.0 Hz)." This means the faces are fading in and out 1.5 times per second and the houses are fading in and out 2 times per second as illustrated in Figure 1B. You could say the two image streams *flicker* at these two frequencies. This trick is known as frequency *tagging* because the neural response in different brain areas is tagged according to which stimulus is driving the response in each area: *if the response in area X is flickering at 1.5 Hz we infer that area X is responding to the 1.5 Hz stimulus.* Here is a YouTube video of a short segment of the stimulus.

Q4. Figure 1B illustrates the "attend FACE" condition in blue and the "attend HOUSE" condition in red. How did the experimenters ensure that subjects were attending to the desired stimulus category (i.e. to the faces or houses)? In other words, what task are the subjects performing during the MEG recordings?

Now if we're going to use this frequency tagging technique we need a way to analyze our neural (MEG) data to see which signals are "flickering" rhythmically, and at what frequencies. These are the "frequency analyses" referred to in paragraph 3, which are also called Fourier analysis or spectral analysis. The **power spectrum** is a standard way to represent the contribution (i.e. the power) of each frequency in a mixed signal like the signal from one MEG sensor.

Q5. In the caption for Figure 1C, we don't know what the "Fourier-transform of the minimum norm estimate" is, but based on the text of paragraph 4 and the labels in the figure itself, what is illustrated in Figure 1C?

Paragraphs 5-7 and Figure 2: Measures of attention localization

Figure 2 has lots of panels, and they are not cited in order in the text. To help make sense of them, note that the *panels on the left* (A-F) are all illustrating where three regions are, according to MEG and fMRI measurements. The three regions (ROIs) are two object-related areas shown in Figures 2 B, D and F, and one attention-related area shown in Figures 2 A, C and E. Then the *panels on the right* (G-J) are all results of the frequency analyses, in terms of power spectra of the signals from different brain regions, to see which areas are "tagged" by which stimulus.

Also notice that the caption for Figure 2D does not tell us the meaning of red and blue colors in 2D; however they seem to correspond to the same meanings of red and blue as in 2F.

Q6. The "high resolution fMRI localizers for FFA and PPA" used to generate Figures 2 B and D are not described in the paper, but based on your experience with fMRI approaches to identifying face- and language-specific brain regions, suggest two stimulus conditions that could define an activation contrast that should work as an "fMRI localizer for FFA."

Q7. Based on the power spectra shown in Figure 2G, can one record measurable MEG responses to *houses* from the Fusiform Face Area (FFA)? Briefly explain.

You may skip paragraph 8 and Figure 2H about lateralization, and skim paragraph 9 and Figure 2 I and J about response latencies. The upshot is that the sensory signals arrive at the frontal area (IFJ) later than the temporal perceptual areas (PPA and FFA).

Paragraphs 10-14 and Figure 3: Coherence measures of attention

Now the authors shift from using the frequency analysis just to tag which part of a region's MEG response is due to which stimulus (as in Figure 2G), to using a different kind of spectral analysis that quantifies the level of rhythmic synchronization, or **coherence**, between the signals in two separate brain locations.

Whereas the power spectrum plots the contribution of each frequency in a *single* mixed signal, the **coherence spectrum** plots how synchronized *two* signals are *at each frequency*, as a number between zero (no synchronization) and one (perfect synchrony).

Q8. Figure 3B shows increased coherence between IFJ and PPA when subjects attend to houses. At what two frequency ranges is this increased coherence observed? Which of the two coherence increases is driven by the stimulus, and which appears to be generated internally by the brain?

Q9. From paragraph 14, what are two parallels between the brain's control system for object-based attention and the previously studied neural correlates of spatial attention?

The authors are not completely consistent in their use of the terms "phase advance" and "phase lag." Both terms are referring to a difference in timing between the MEG signals in two brain areas. Figure 2I shows that MEG sensory responses to houses or faces occurred later in IFJ than in PPA or FFA. In this (low frequency) context the PPA and FFA responses "lead" the IFJ responses.

Q10. Which area "leads" when considering coherent *gamma* (high frequency) oscillations? The authors address this issue in the Abstract and in paragraphs 12 and 14. Just skim paragraph 12 to find a relevant statement that is consistent with the Abstract and paragraph 14.

Big picture: We began with a question about how the brain is able to pay attention to just one of two objects that overlap in space.

The answer suggested in this paper is that the frontal area, **IFJ, synchronizes its activity with the perceptual area representing whichever type of object is being attended** (FFA for faces and PPA for houses), as a way to *amplify* the neural representation of the attended object and "direct the flow of visual

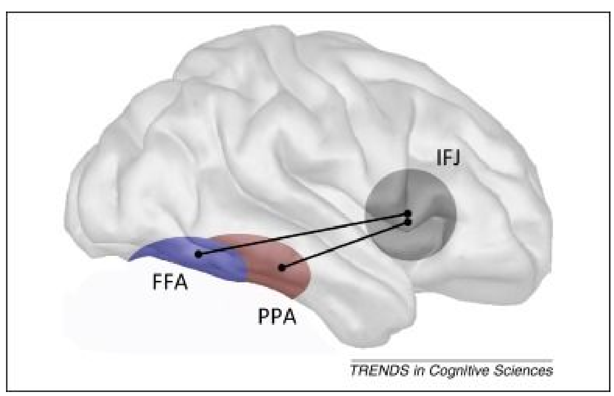

processing" among these and other brain areas.

The *timing* of the gamma oscillations in the frontal area as compared to the posterior visual areas suggested a picture in which, for example, a (top-down) desire to attend to a particular object category originates in IFJ, and signals from IFJ down to PPA or FFA then further activate the PPA/FFA neurons, and the impact of that excitation or target groups of other neurons is enhanced by the rhythmic synchronization.

Q11. Please submit two questions you have about terms, figures, concepts or anything in this article that confused you or that you'd like to pursue further during our Journal Club discussion. For example, one question might be related to a technical detail, and another might be broader (e.g. related to assumptions, methods, interpretation, or open questions for future research).