CRISSP Lecture Series

A program for experimental syntax: data, theory, and biology

Jon Sprouse
University of Connecticut

Biology 2: Beyond acceptability judgments

KU Leuven - Brussels 03.18.15

From Cognitive Theories to Neural Systems

	phonology	**morphology**	**syntax**	**semantics**
representations:	[+/- feature]	V + ed	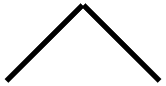	λx.λy
computations:	conversion	fusion/fission	merge/move	ƒ(x)/∧

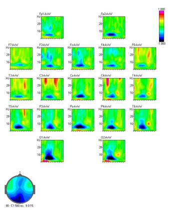

Use electrophysiology to identify large-scale neural correlates of representations and/or computations

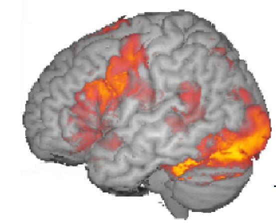

Use neuroimaging to identify networks underlying electrophysiological correlates

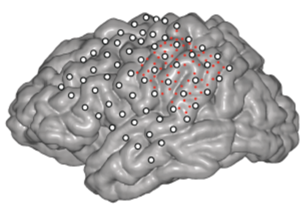

Use ECoG to identify single-unit activation in the relevant networks

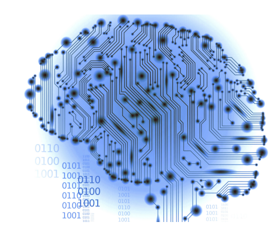

Use theoretical/computational neuroscience to postulate plausible encodings of representations and computations.

The brain is bloody and electric

There are two primary approaches to neurolinguistics, each based on a different property of the brain:

Hemodynamic approaches:

The brain is an organ like any other in the body, and therefore requires oxygenated blood to function. Hemodynamic approaches measure changes in (oxygenated) blood flow.

Electrophysiological approaches:

The cells of the brain (neurons), like the other components of the nervous system, use electricity to communicate with other cells. Electrophysiological approaches measure changes in the flow of electricity in the brain.

For most people, neuroimaging means hemodynamic approaches

When most people think of neuroimaging, they think about actually “seeing” the tissues of the brain, or seeing which areas of the brain are active during a cognitive task:

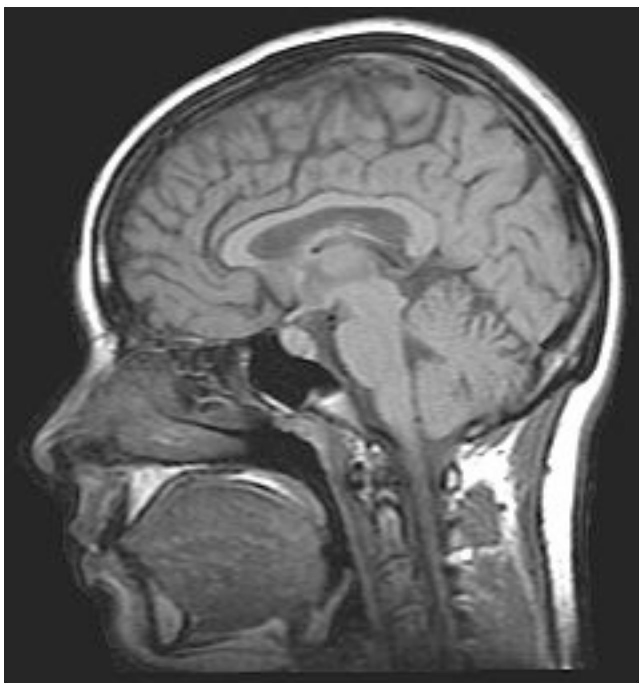

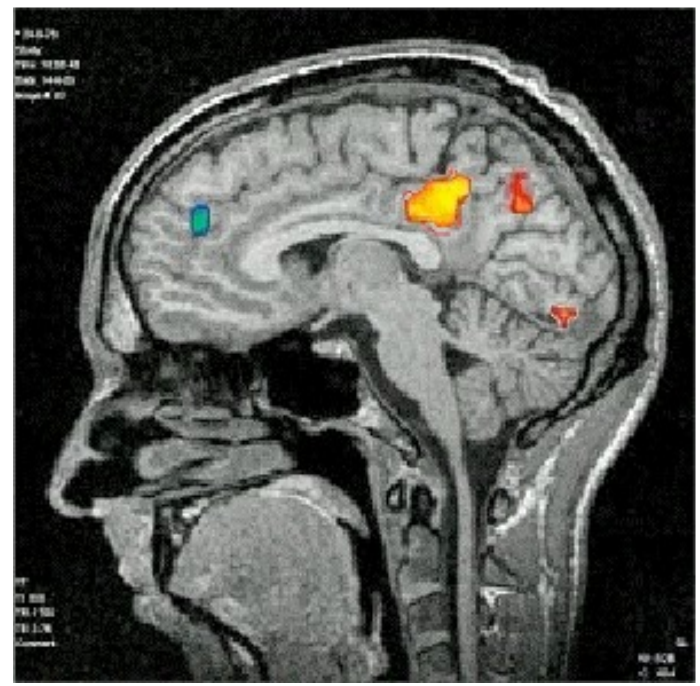

Hemodynamic approaches, in particular **functional Magnetic Resonance Imaging (fMRI)**, are designed to allow us to see these tissues and differences in activity (blood flow) between tissues.

Magnetic Resonance Imaging

fMRI is a brain imaging technique that allows us to measure the hemodynamic response to cognitive tasks.

This response is called the **Blood Oxygen Level Dependent** signal… the BOLD signal.

The idea behind this is that brain areas that are recruited for a **cognitive task will require more oxygen**, therefore there will be an increase in the BOLD signal for those areas, but not others.

fMRI is good for "where" questions

fMRI has excellent spatial resolution: it is great at telling us where activity is happening in the brain.

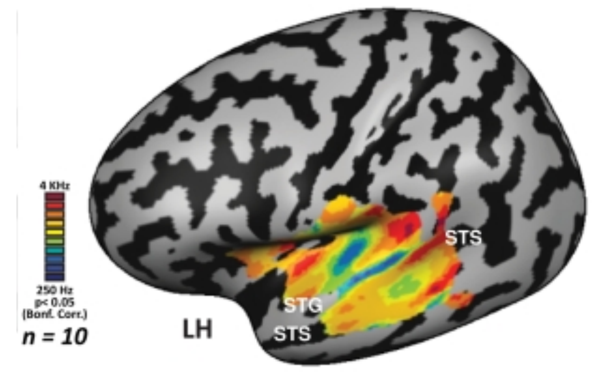

But fMRI has poor temporal resolution: the vascular system is slow (think about heartbeats), so the hemodynamic response to cognitive activity is 1-2 seconds or more.

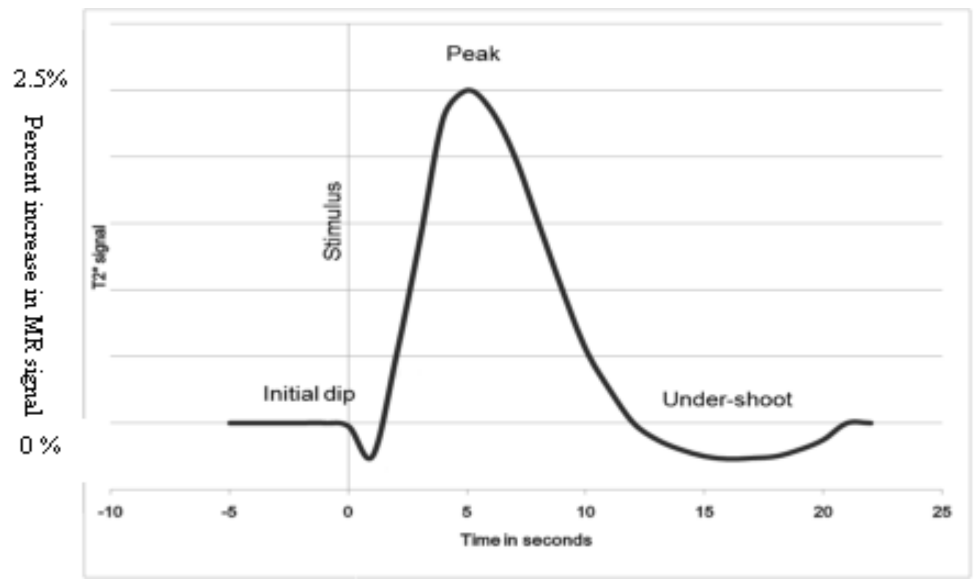

Santi and Grodzinsky 2007:

length of wh-dependency

... the woman **who** Kate burnt ___.

... the woman **who** the mailman and Kate burnt ___.

... the woman **who** the mailman and the mother of Jim burnt ___.

length of binding

... the **mailman** who burnt **himself**.

... the **mailman** who loves Anne burnt **himself**.

... the **mailman** who loves the sister of Kim burnt **himself**.

Broca's area

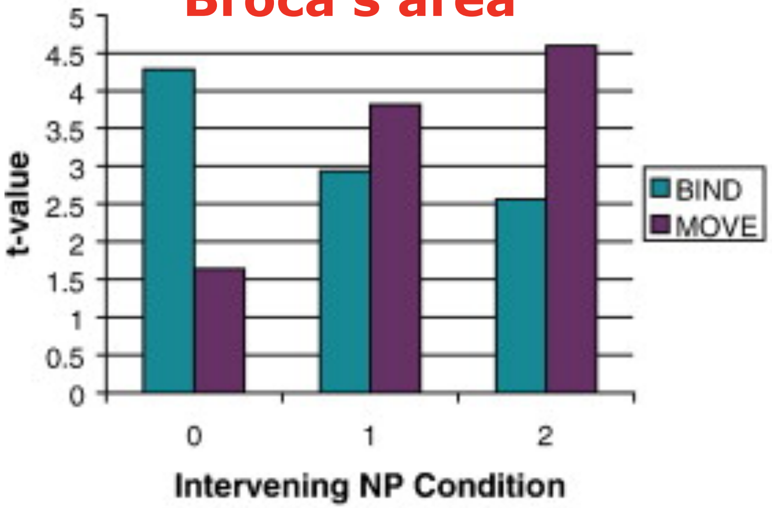

This could be taken to suggest that Broca's area is selectively sensitive to movement.

The problem is that there is also a difference in the **parsing** of these dependencies.

We can eliminate this confound by using **backward anaphora**, which share some aspects of the processing profile of wh-questions, such as the forward search for the second half of the dependency:

Because **she** decorated the wedding cake, the **baker** wowed the customer that made the long order.

Because **she** decorated the wedding cake that was six layers tall, the **baker** wowed the customer.

Which song did the band play ___ poorly and unenthusiastically at the concert that ended early?

Which song did the band that won the contest play ___ poorly and unenthusiastically at the concert?

We found a significant effect of long backwards anaphora (compared to short backwards anaphora) in Broca's area.

This suggests that Broca's area is not specific to wh-movement, but instead is responding to something about the way these dependencies are processed.

As a syntactician, this is disappointing, but seems like a necessary step toward figuring out if there are any brain areas that are sensitive to (truly) syntactic processes.

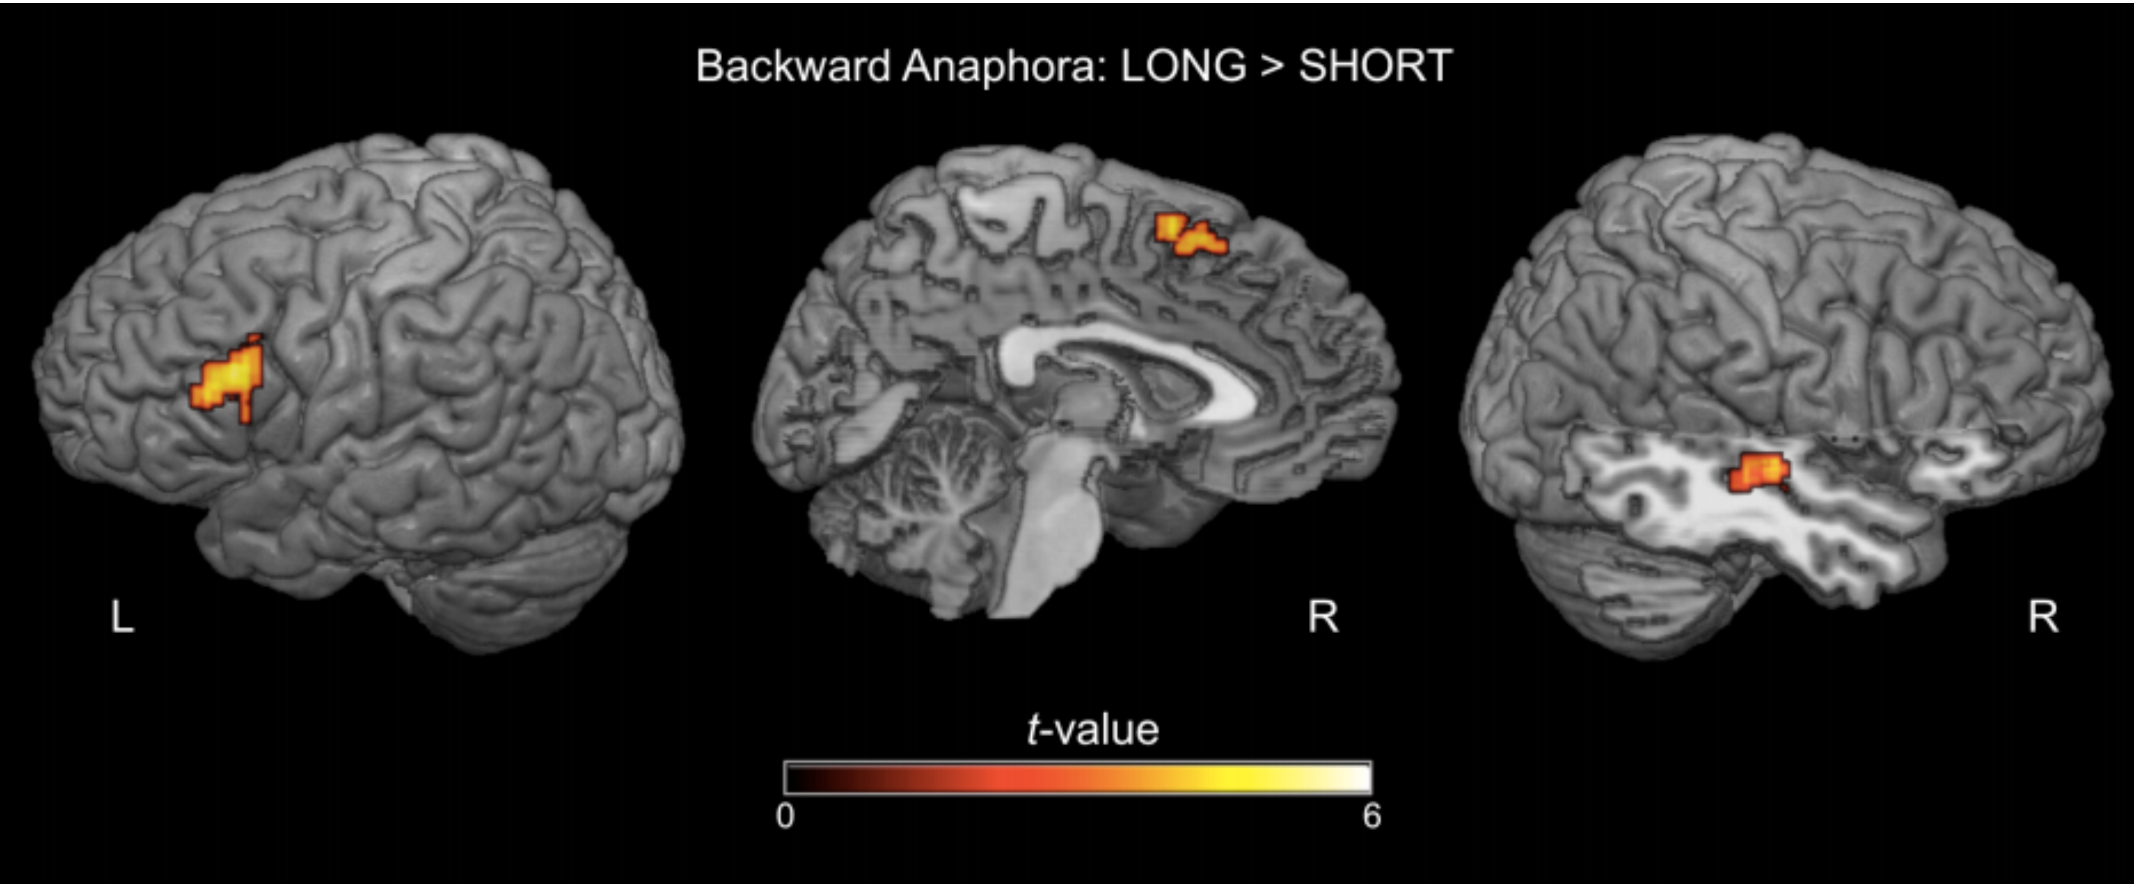

Electro**e**ncephalo**g**raphy

EEG means electric head writing, and that really is what it is:

EEG measures changes in electrical potentials that occur on the scalp.

The underlying idea is that these scalp potentials are generated by the electrical activity of the cortex.

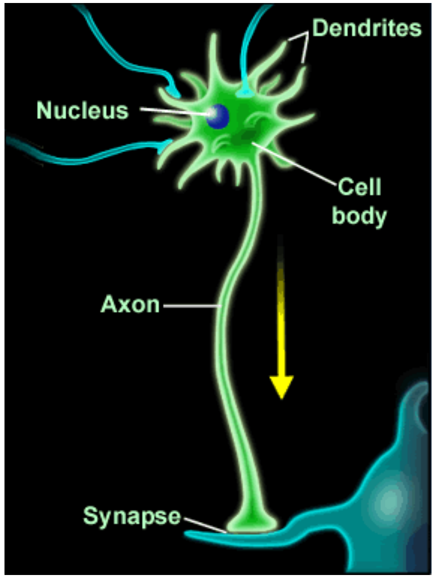

Where do scalp potential come from?

If a population of neurons discharges in the same direction, the current will become large enough to be detected on the surface of the scalp.

One nice property of EEG is that the orientation of the population of neurons doesn't matter -- the current can be oriented in any direction and still be (potentially) detectable on the scalp.

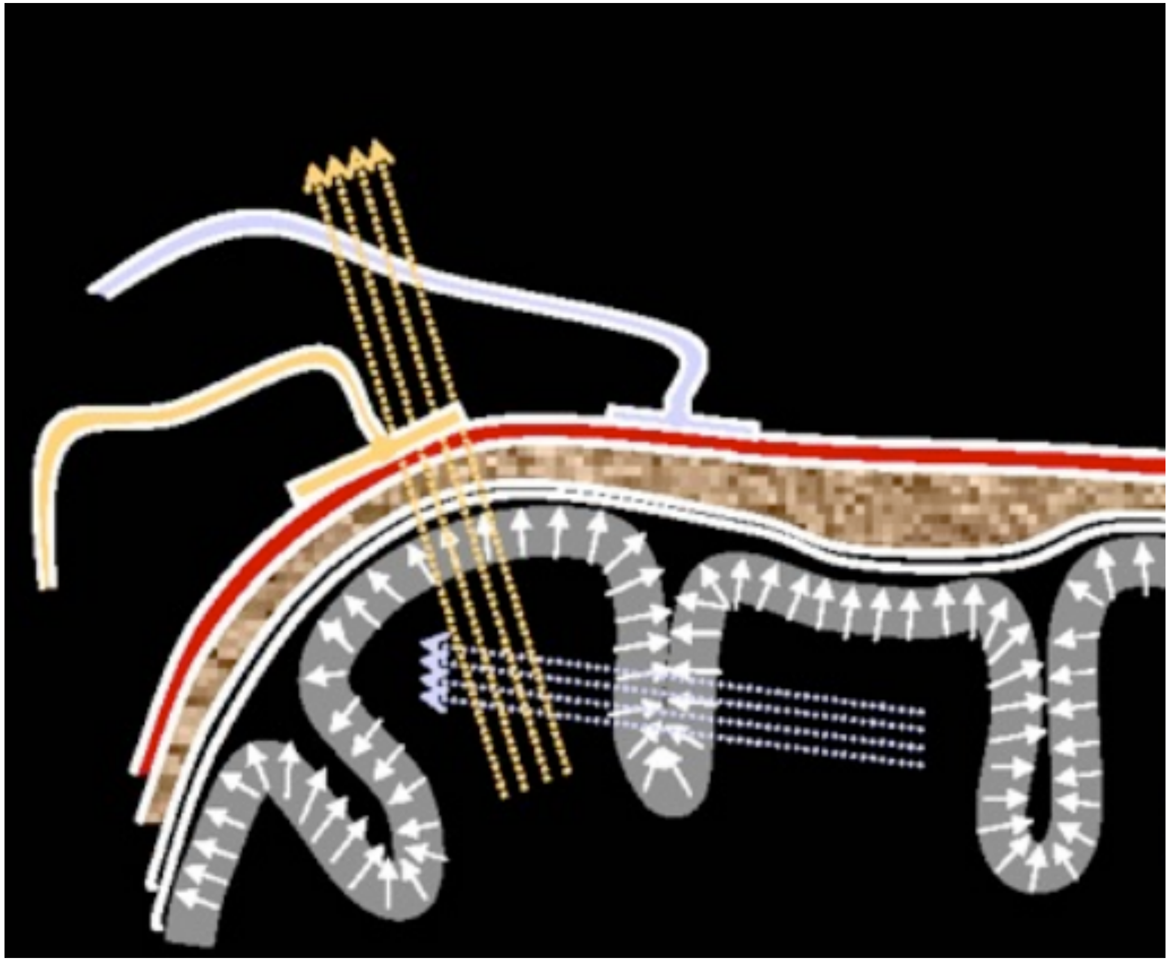

Excellent temporal resolution

EEG has excellent temporal resolution for two reasons:

1. The electrical activity of the brain (action potentials) travels very quickly: anywhere from 20m/s up to 120 m/s depending on the neuron (nerve) type.

   This means that the information reaching the scalp electrodes is nearly instantaneous -- much faster than the 1-2 second lag with hemodynamic responses.

2. EEG systems can measure scalp potentials thousands of times per second.

   We call the number of measurements that the EEG system makes per second the **sampling rate**: the electrical changes on the scalp are a continuous stream of information, and the EEG system samples from that information a certain number of times per second.

   Sampling rates are measured in Hertz (samples per second). A sampling rate of 1000 Hz would be 1000 measurements per second, or 1 measurement every 1 millisecond.

Lots of electrodes, but very poor spatial resolution :-(

You might think that EEG has good spatial resolution because of the evenly space electrodes around the scalp.

But in fact, EEG has very poor spatial resolution. The reason for this is that what we want to localize is the neural generator of the scalp potential inside the cortex, but there are several layers of “stuff” between the cortex and scalp.

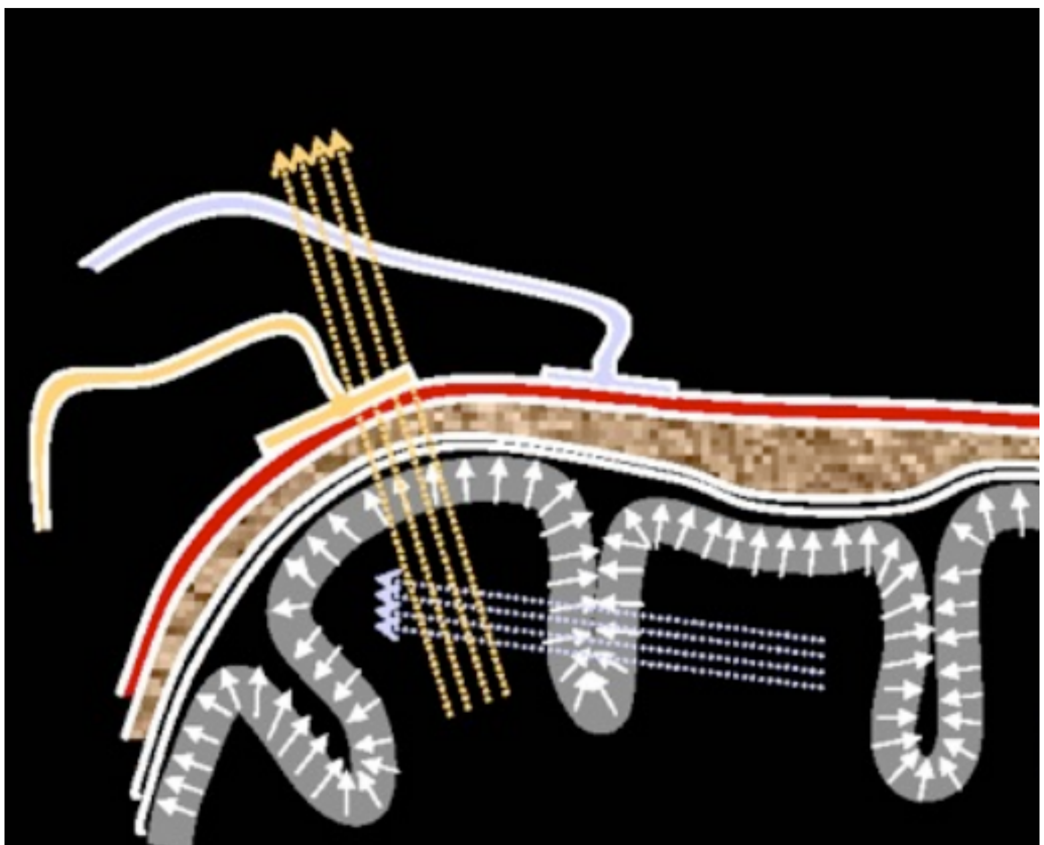

This stuff smears the electrical signal making it impossible to tell where it originated.

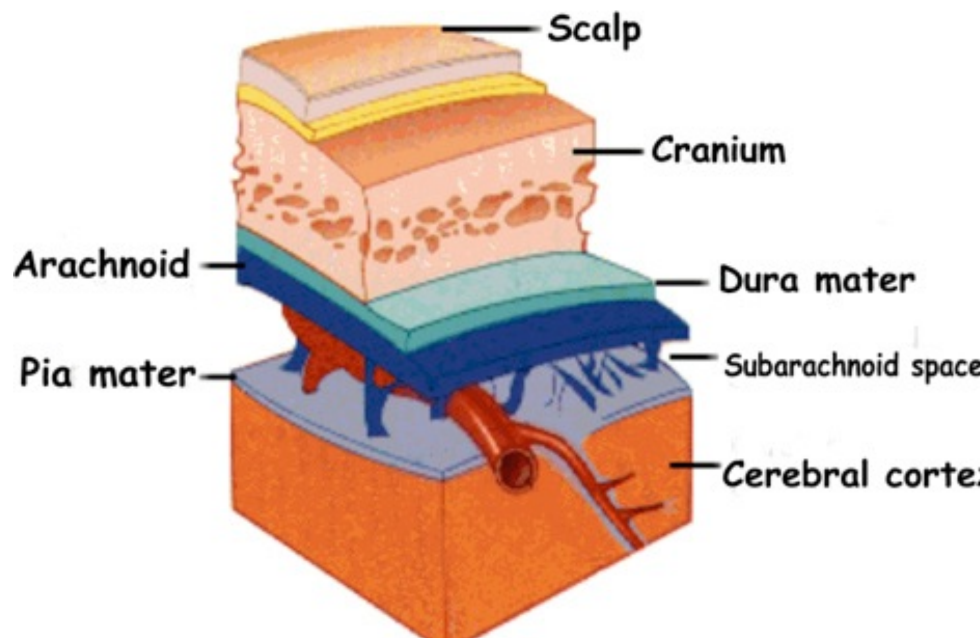

How are scalp potentials represented?

The most common visual representation of scalp potentials is as a **wave**.

In this waveform, the x-axis represents time, and the y-axis represents the polarity and amplitude of the potential (in microvolts):

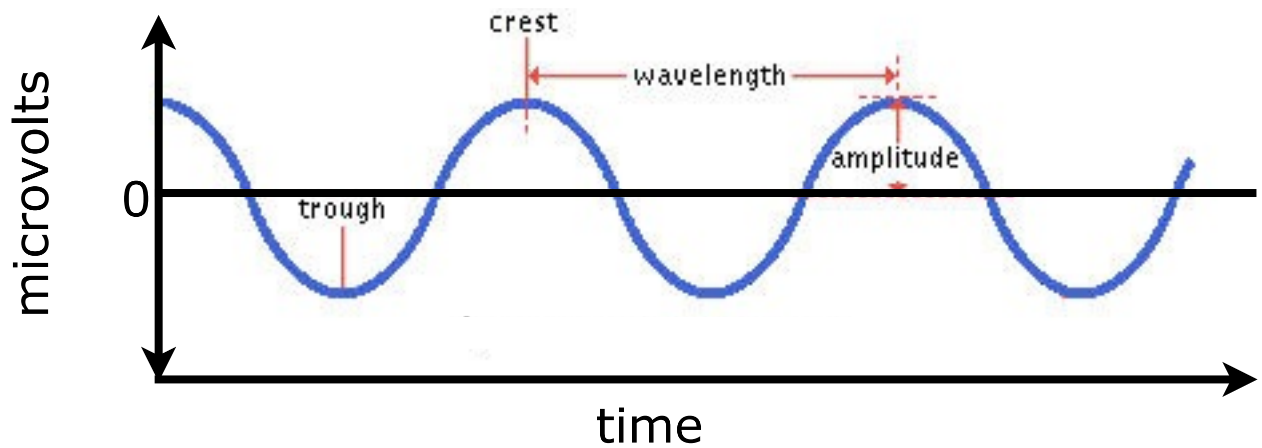

You will sometimes hear this described as the time-amplitude domain, which is simply a way of saying that this visual representation highlights changes in the amplitude of the scalp potentials over time.

The raw EEG trace isn't very useful

Here is what the (time-amplitude) representation of EEG looks like from multiple electrodes placed around the scalp.

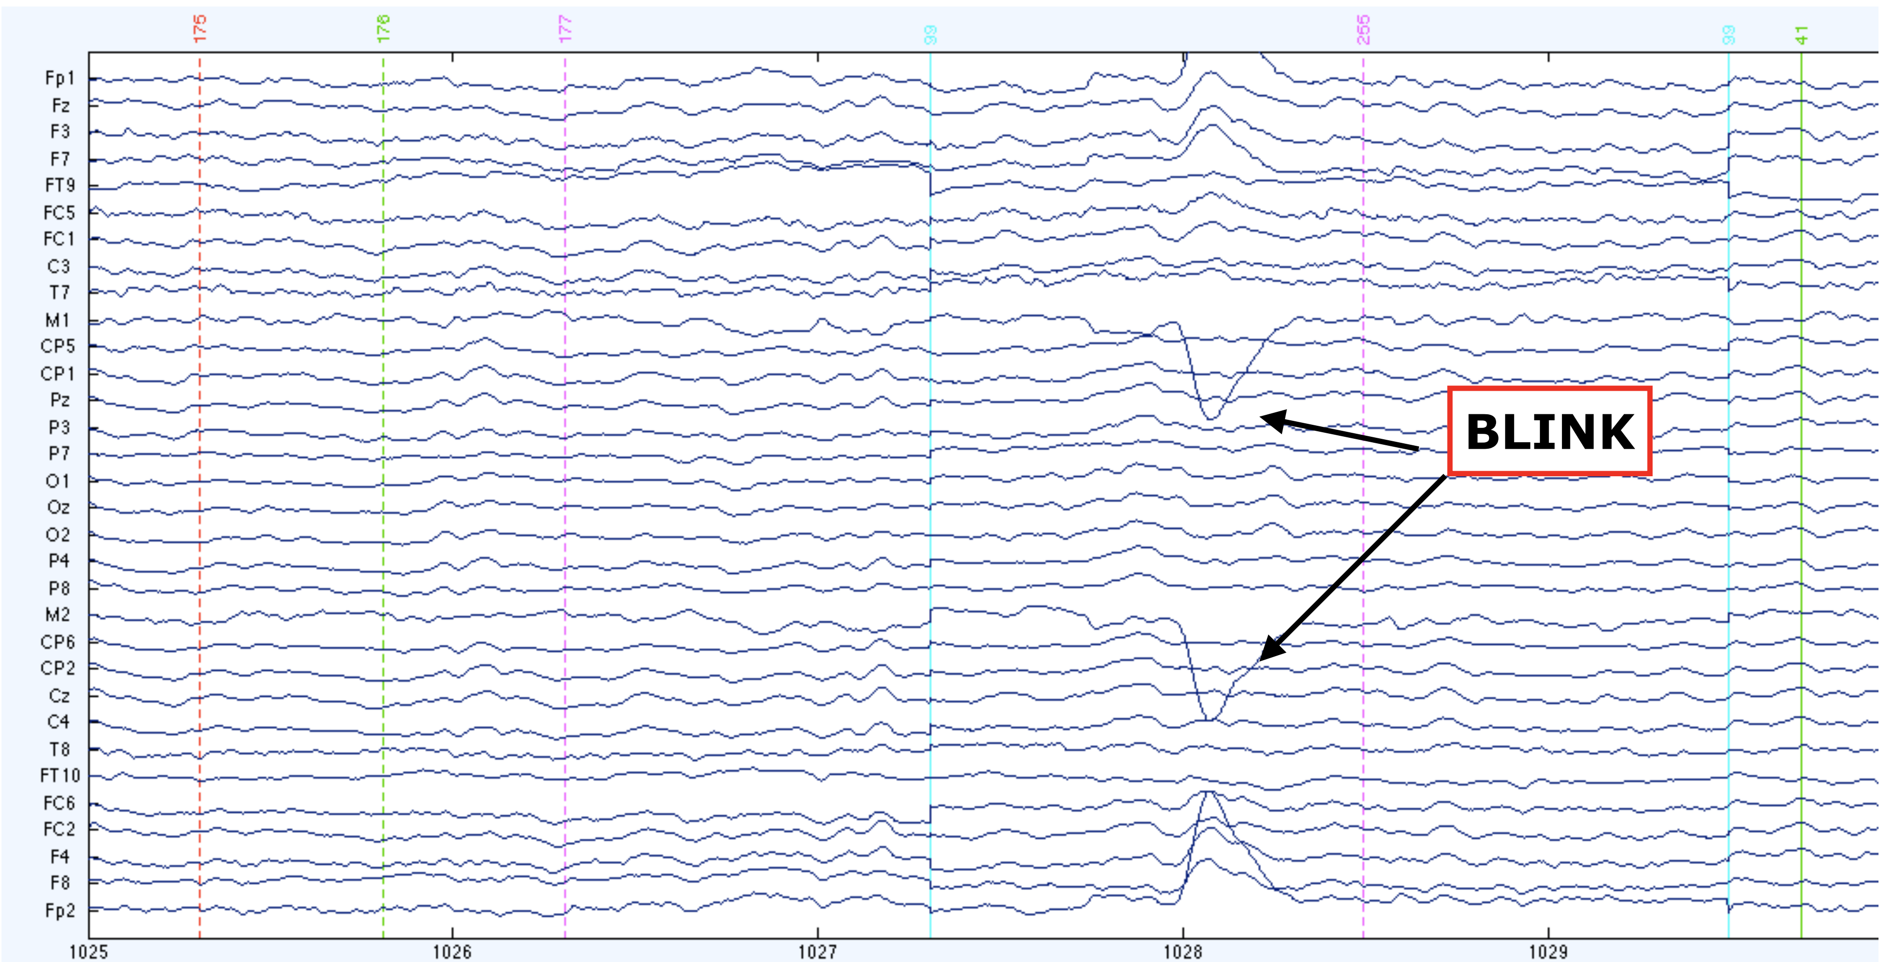

The problem with the raw EEG is that it contains both cognitive activity and lots of **other activity** (like all of the functions that keep you alive!)

Typical data processing to turn EEG into an Event-Related Potential (ERP)

If you look closely at the trace below, you will see colored vertical lines with numbers on them. Those are **event markers**. The experimenter puts those in to indicate times in the trace where an important event took place.

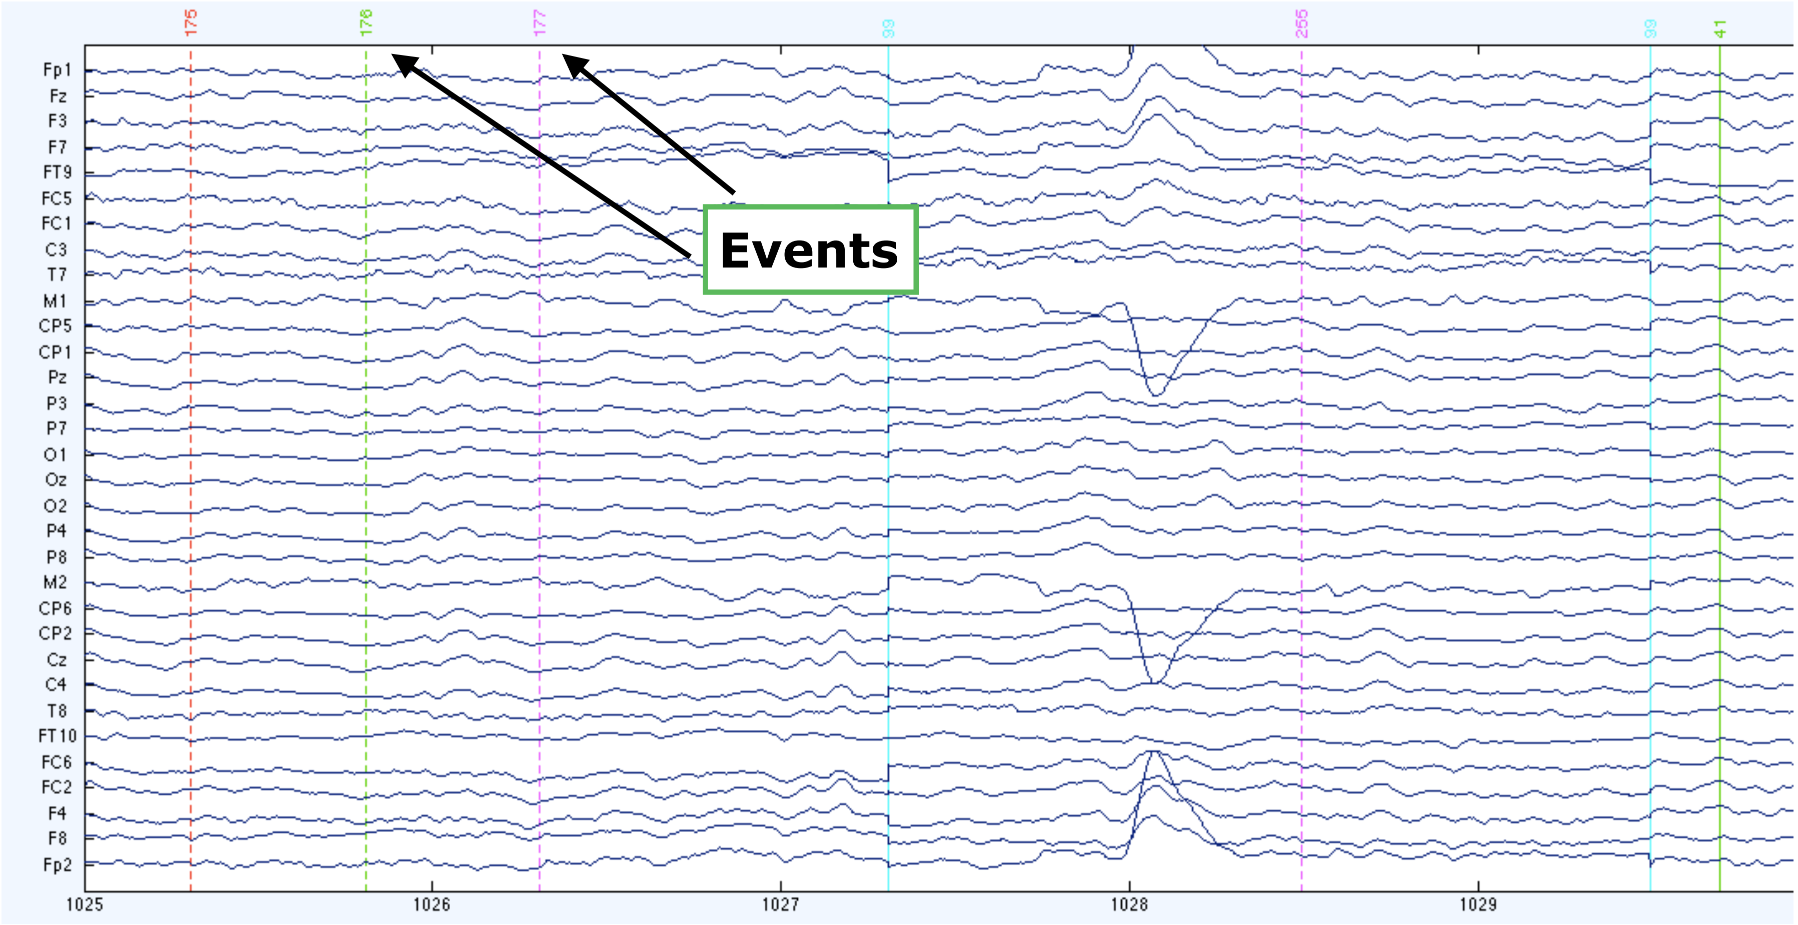

Typical data processing to turn EEG into an Event-Related Potential (ERP)

You can tell special computer programs (like Matlab toolbox ERPLAB) to cut out a slice of time around an event. Typically you cut out 100-200ms before the event, and 800-1000ms after the event. Here is a plot of a series of epochs after they've been cut out:

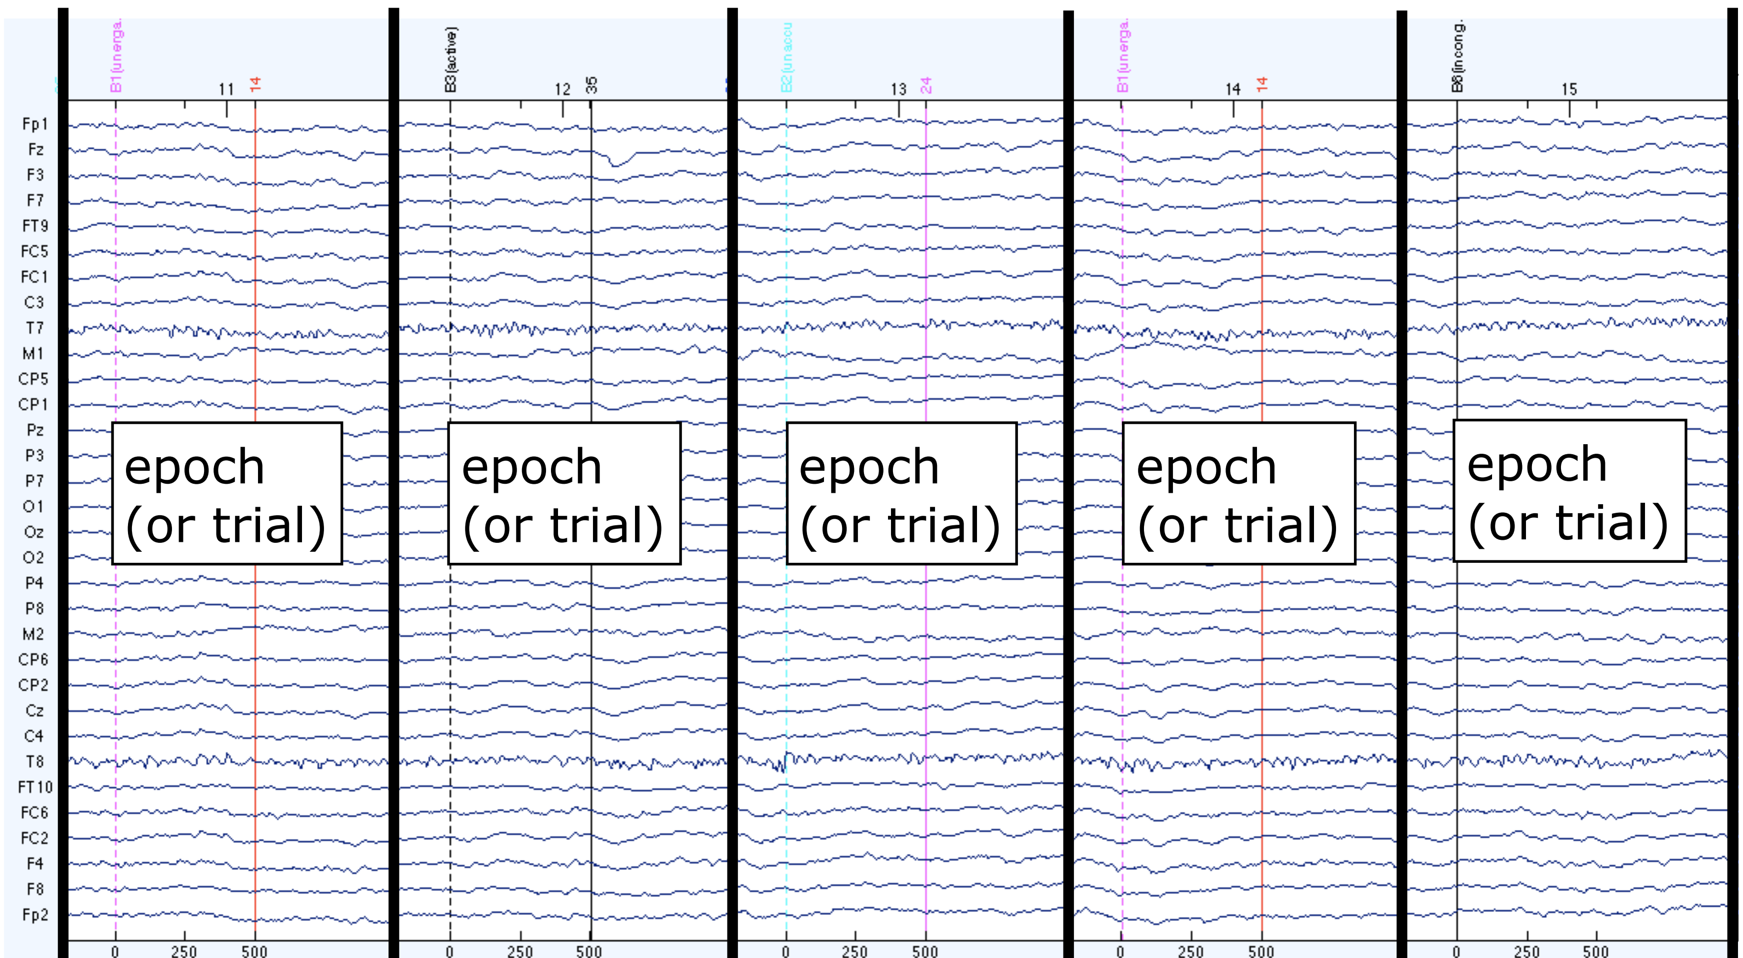

Typical data processing to turn EEG into an Event-Related Potential (ERP)

Then you can align all of the epochs (or trials) from one condition using the event as time-point 0, and average them together.

EEG epochs (or trials in your experiment)

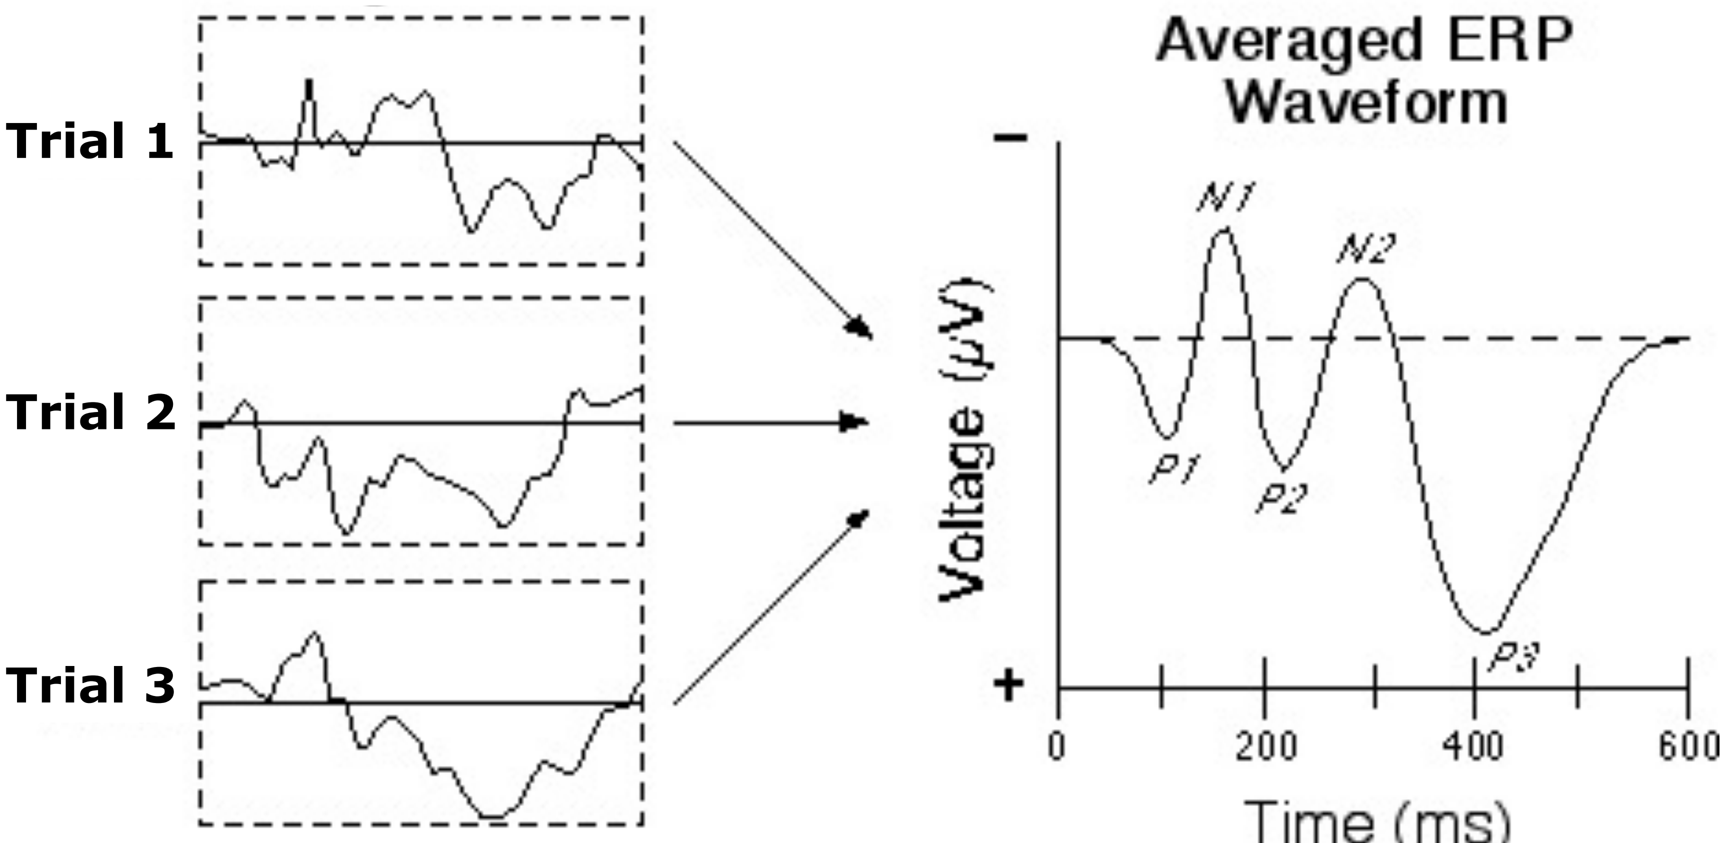

This process eliminates noise

The reason that we average over trials is that EEG data contains lots of noise. By noise, we mean electrical activity that is not related to the cognitive event we are studying.

The theory behind averaging is as follows:

Assumption 1: The signal from the cognitive event is roughly identical in latency and amplitude across trials.

Assumption 2: The signal from noise sources is roughly random in latency and amplitude across trials.

Therefore, when averaging across trials, the signal will remain roughly constant, and the noise will reduce to roughly zero.

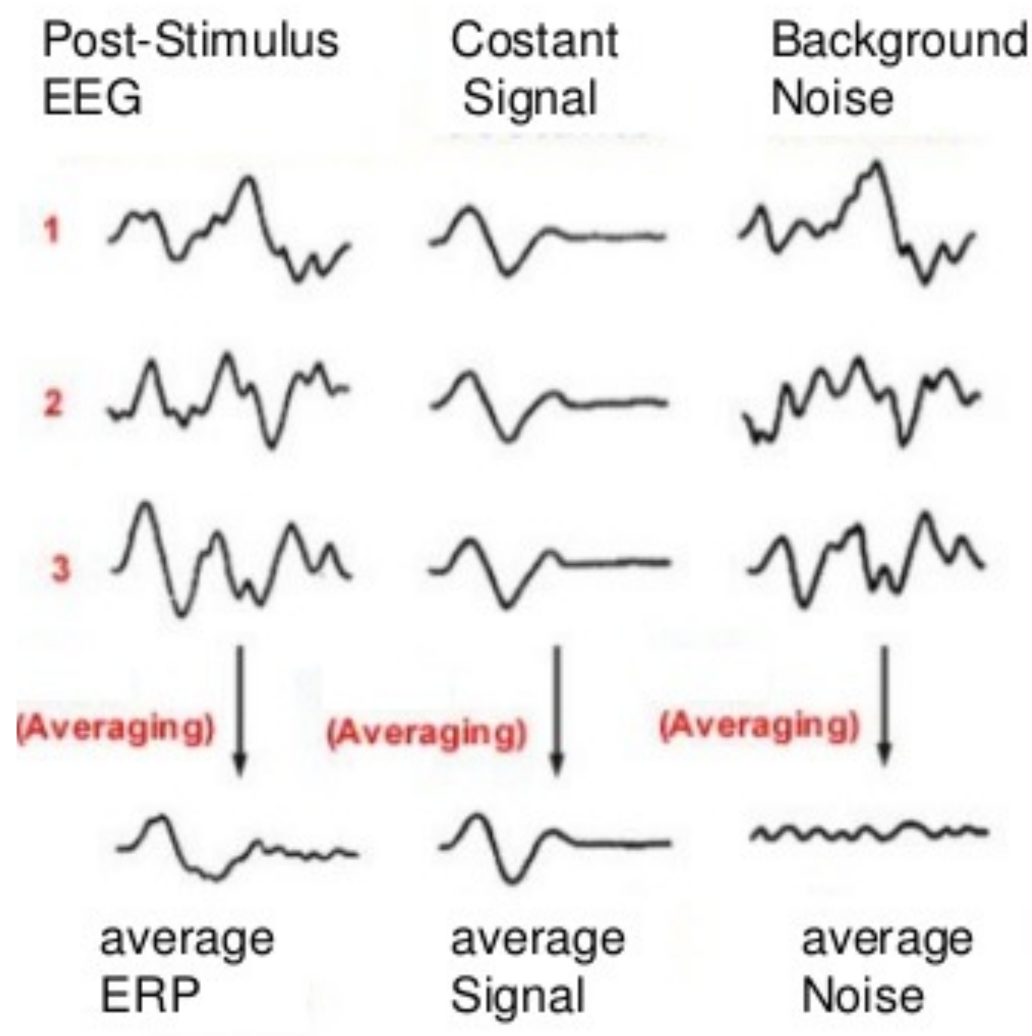

But it also eliminates a lot of information

ERPs only contain activity that is both **time locked** and **phase-locked** (because the simple averaging procedure creates destructive interference)

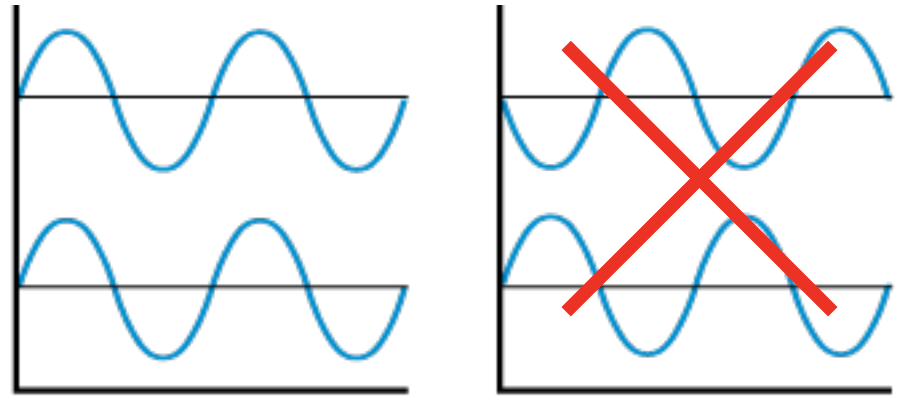

We don't know (yet) exactly what information has been lost. But it is probably no coincidence that there are relatively few ERPs that have been reliably identified in the language processing literature:

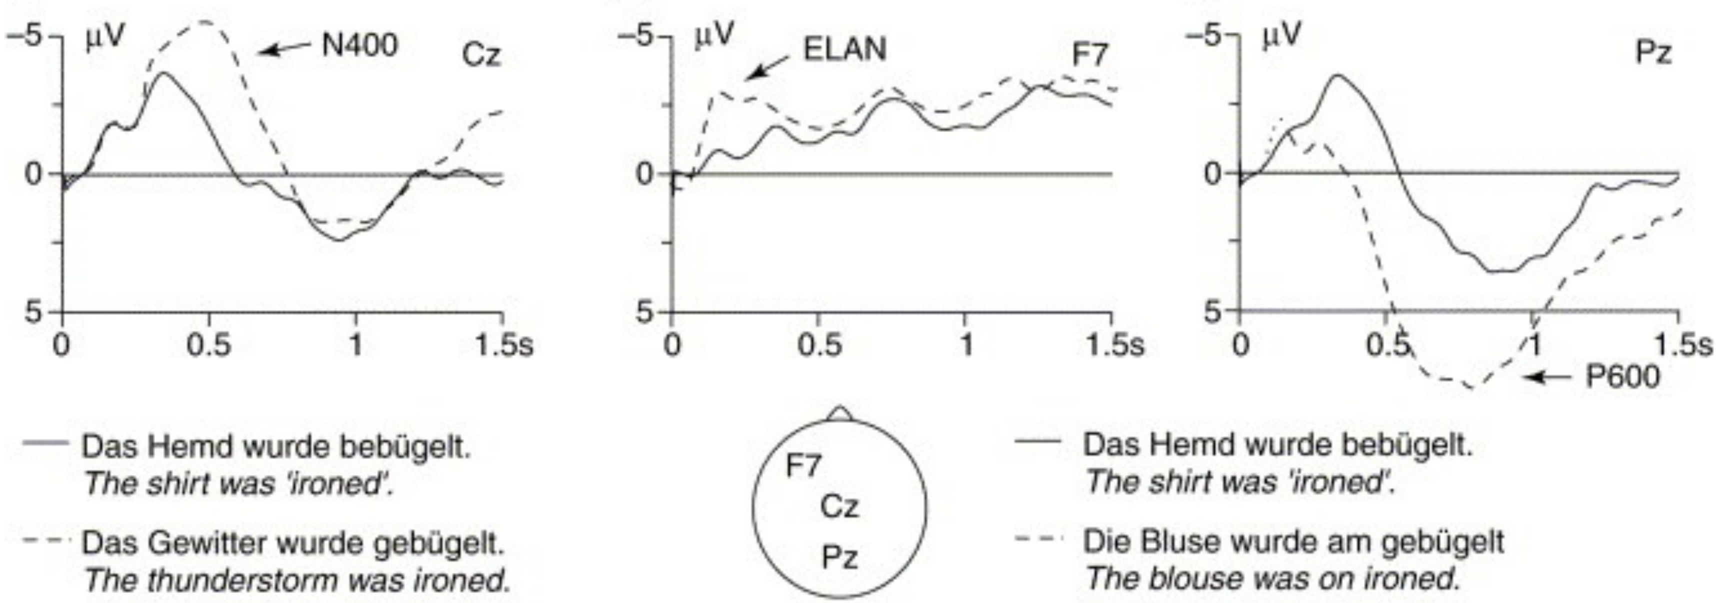

But there is a different way to analyze the data that sacrifices less information

Whereas ERPs are sometimes called a time-amplitude analysis because they focus on amplitude changes over time, there is also something called a **time-frequency analysis** that focuses on changes in the presence/absence of **neuronal oscillations** at a certain frequency over time.

What this means in practice is that time-frequency analyses allow us to see activity that is **time locked** and **out-of-phase**(but crucially the same frequency).

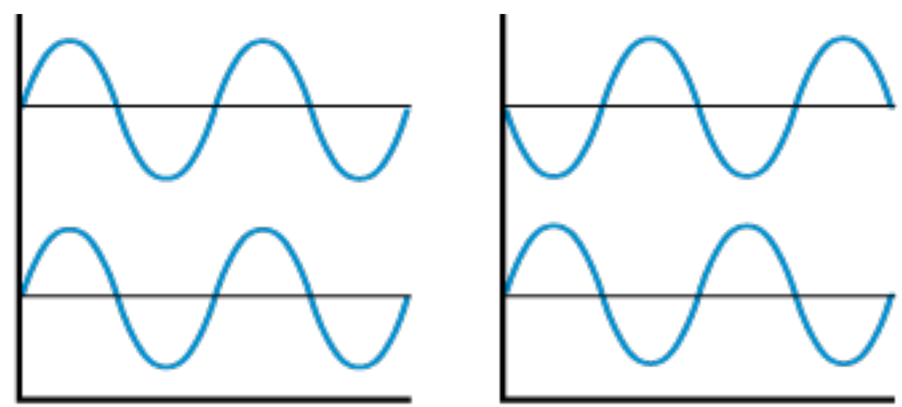

This is in contrast to ERPs, which only show activity that is both **time locked** and **phase-locked**

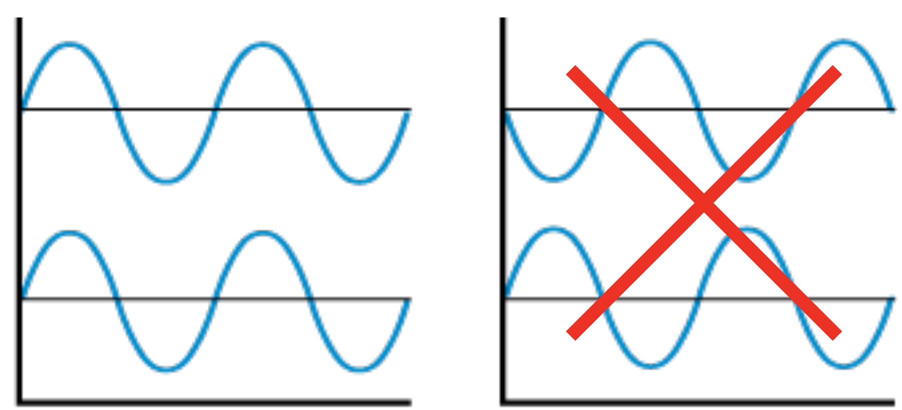

One way to think about it is as a measure of how much of each frequency is in the signal

One of the most important discoveries about waves is that any complex wave can be decomposed into a set of simple waves (i.e., a wave with only one frequency).

The process of decomposing a complex wave into its component waves is called a **Fourier Transform**.

EEG signals are just complex waves. So by doing a Fourier Transform, we can look at what component waves are in it.

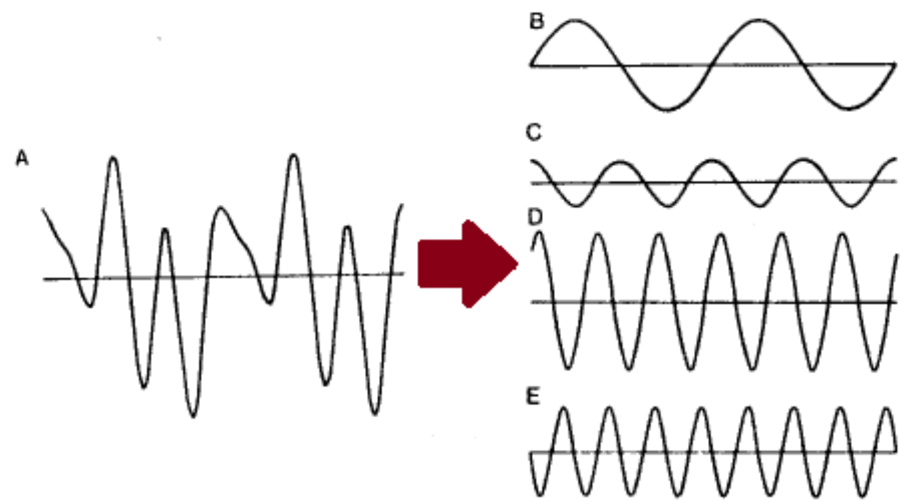

What we want to do is two-fold: first, we want to see what frequencies are in the complex wave. This (potentially) tells us which frequencies the neurons are oscillating at during an event.

Second, we want to see how much **power** is in each frequency so that we compare two (or more) conditions to see what changes.

To make our lives easier, we tend to chunk the frequency space into "bands"…

Here are four of the major bands (there is fifth called Gamma that isn't shown)

Part of the reason we chunk into these bands is that the frequencies in these bands tend to pattern together during different cognitive tasks, so this seems to be a good categorizations.

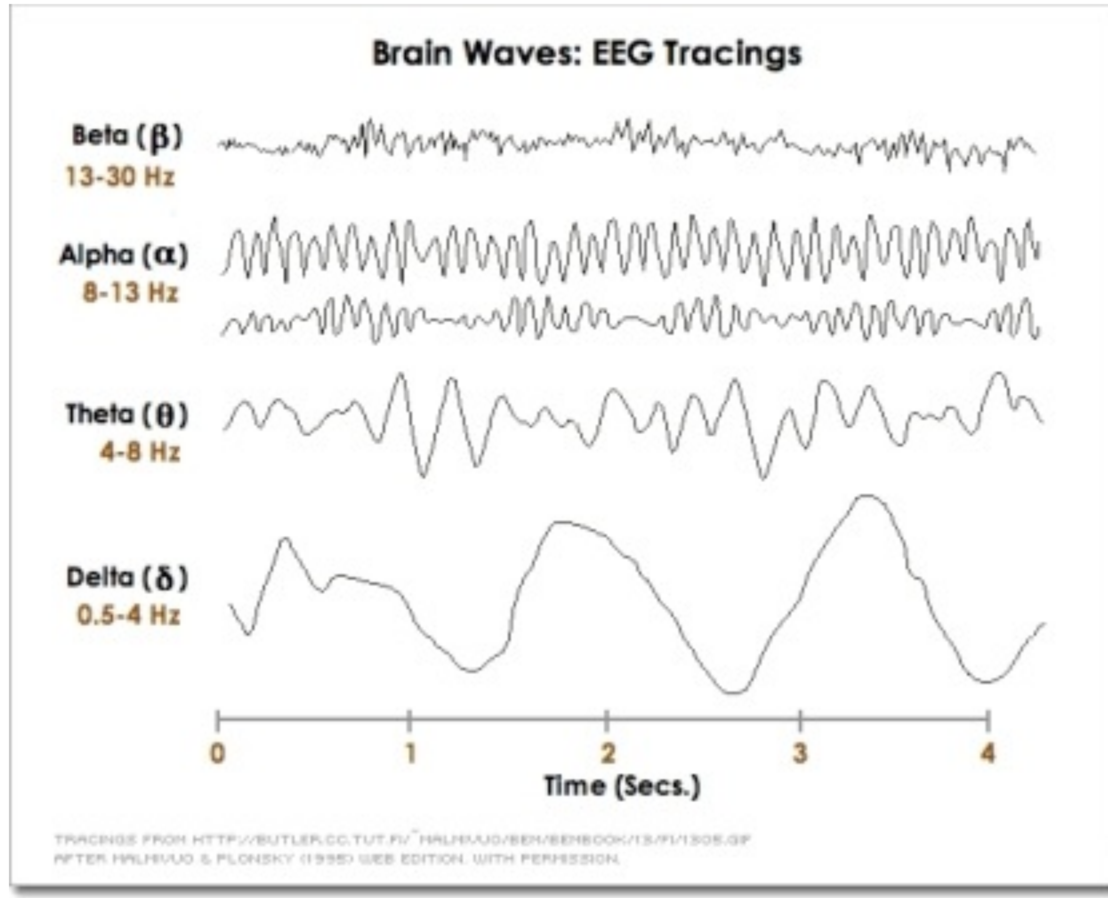

If you ever google these, you will read that lower frequencies tend to be associated with relaxation (theta) and sleep (delta), while higher frequencies tend to be associate with alertness (alpha) and concentration (beta). But we can also attempt to associate them with more specific cognitive processes.

Neuronal oscillations seem likely to be a better step along this path than ERPs

	phonology	**morphology**	**syntax**	**semantics**
representations:	[+/- feature]	V + ed	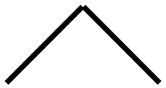	$\lambda x.\lambda y$
computations:	conversion	fusion/fission	merge/move	$f(x)/\wedge$

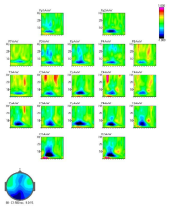

Use electrophysiology to identify large-scale neural correlates of representations and/or computations

Use neuroimaging to identify networks underlying electrophysiological correlates

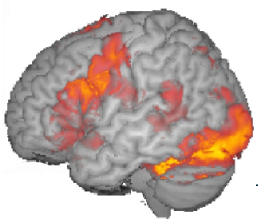

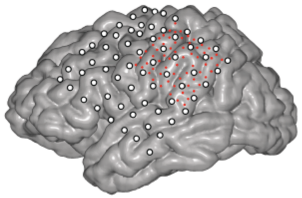

Use ECoG to identify single-unit activation in the relevant networks

Use theoretical/computational neuroscience to postulate plausible encodings of representations and computations.

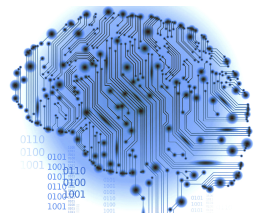

Study 1: Look at some violations

The cameraman knew that the mayor would honor **them**...

The cameraman knew that the mayor would honor **they**...

Case

The cameraman knew that the mayor would **honor** the soldiers...

The cameraman knew that the mayor would **honors** the soldiers...

Agreement

The cameraman knew that the mayor would honor **the soldiers**...

The cameraman knew that the mayor would honor __ **before** ...

Theta

For the record, all of these tend to show a LAN when analyzed using an ERP analysis.

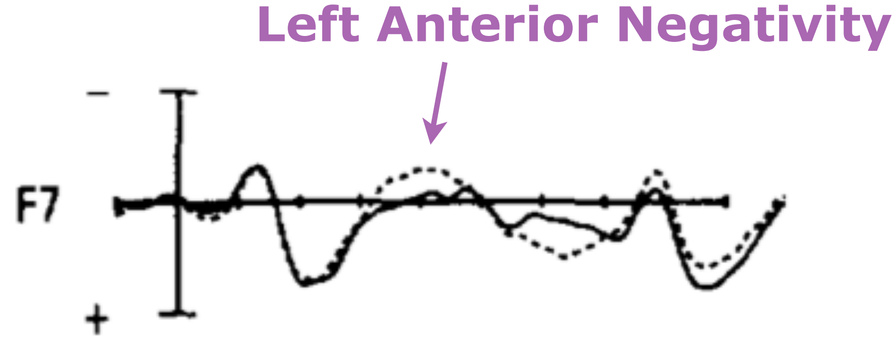

Case Violation (electrode Cz)

The cameraman knew that the mayor would honor **them**...

The cameraman knew that the mayor would honor **they**...

Case

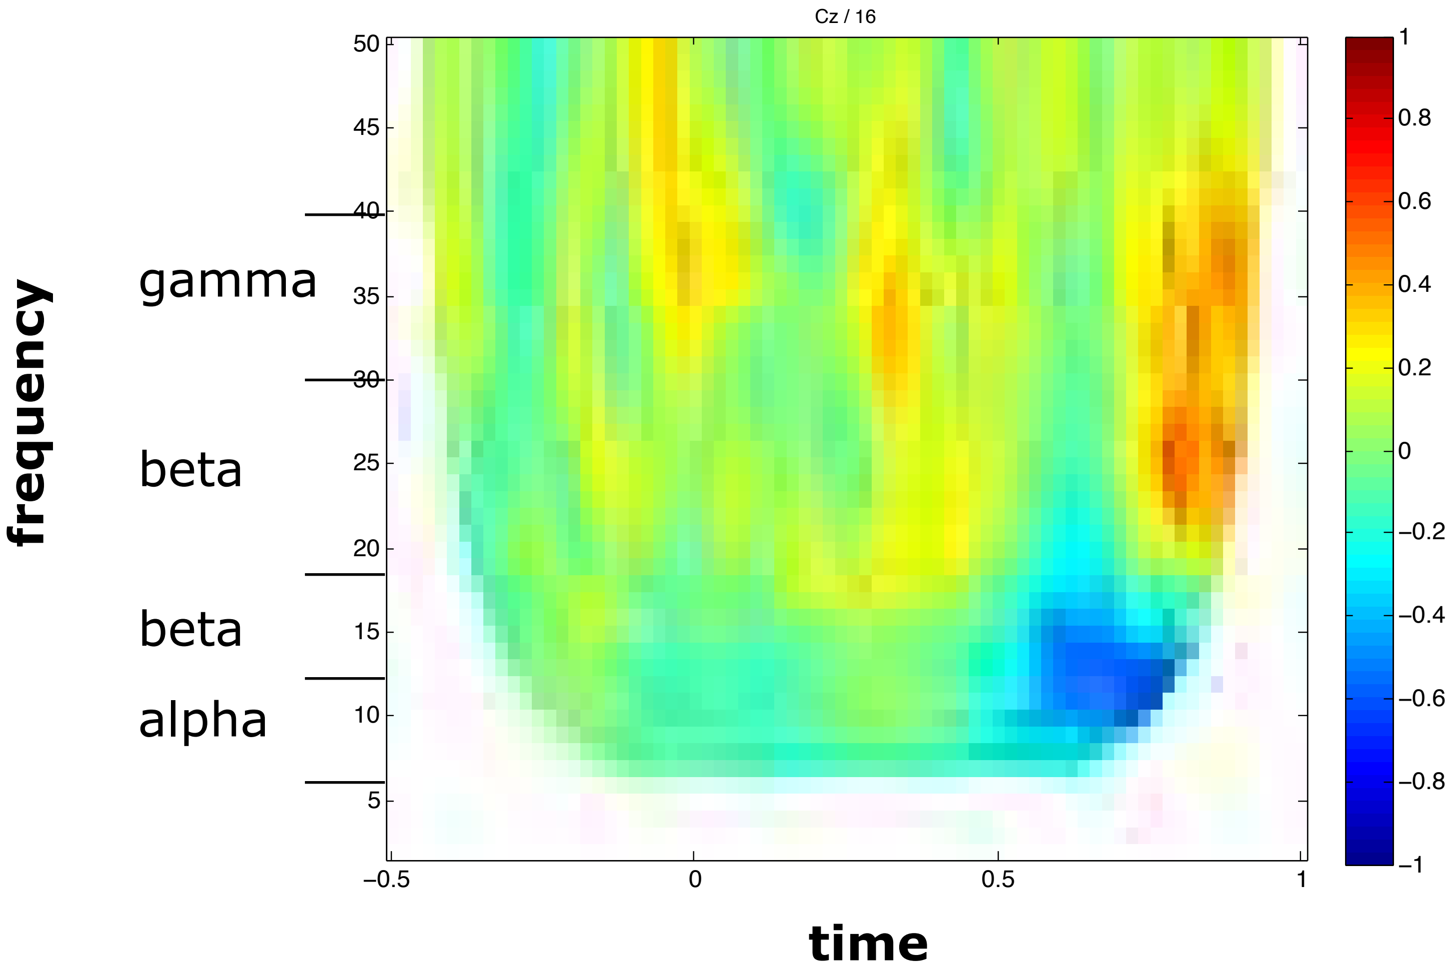

Case Violation (all electrodes)

The cameraman knew that the mayor would honor **them**...

The cameraman knew that the mayor would honor **they**...

Case

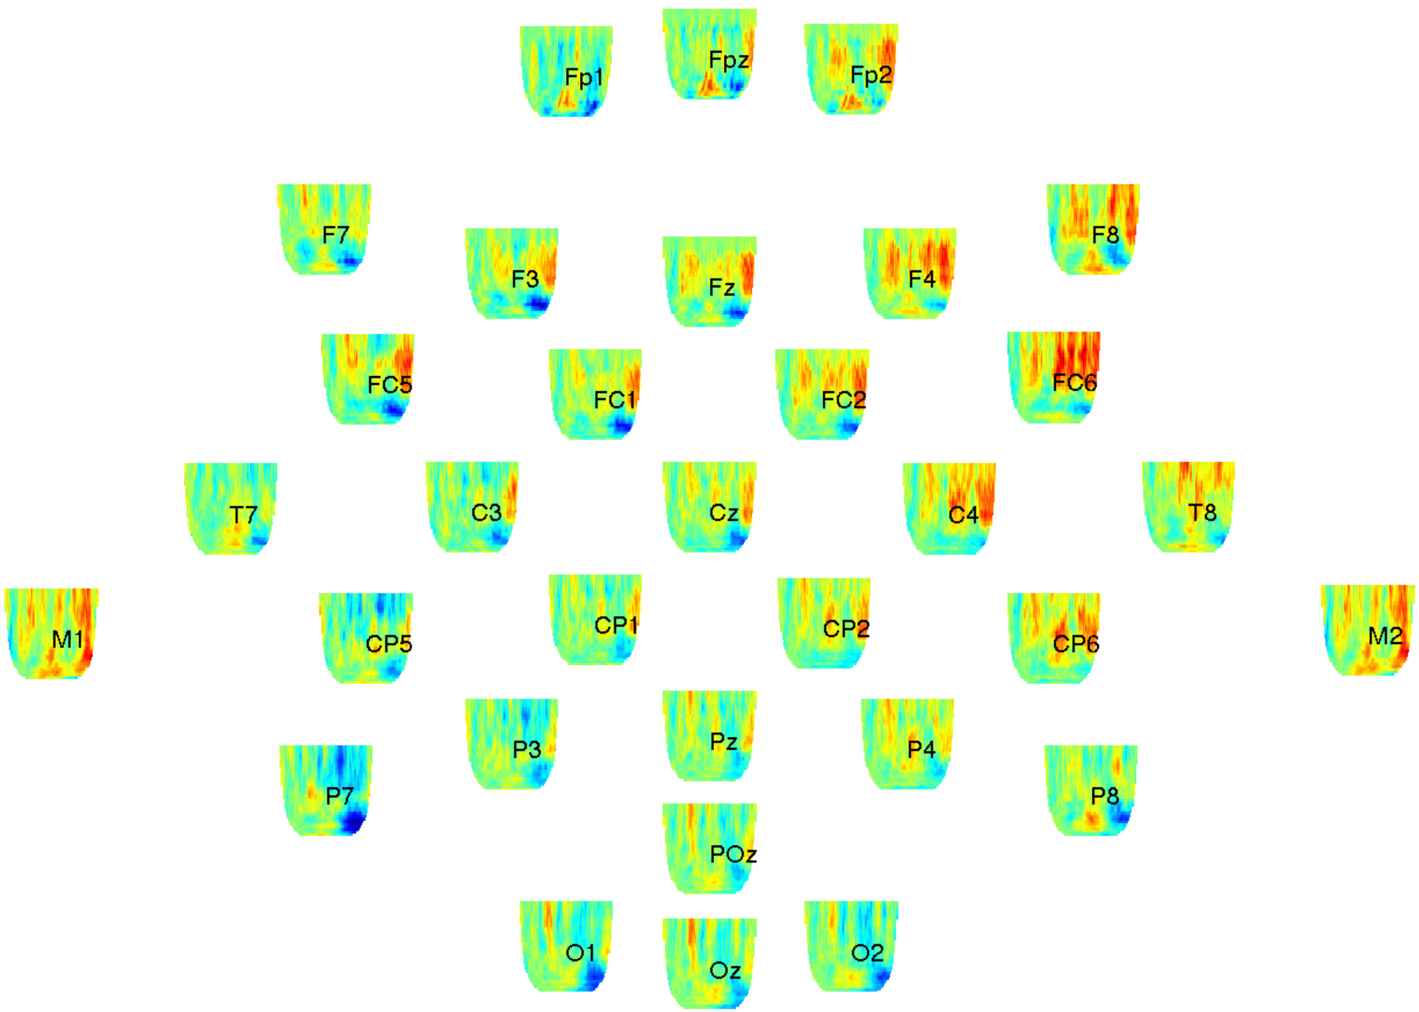

Agreement Violation (electrode Cz)

The cameraman knew that the mayor would **honor** the soldiers...

The cameraman knew that the mayor would **honors** the soldiers...

Agreement

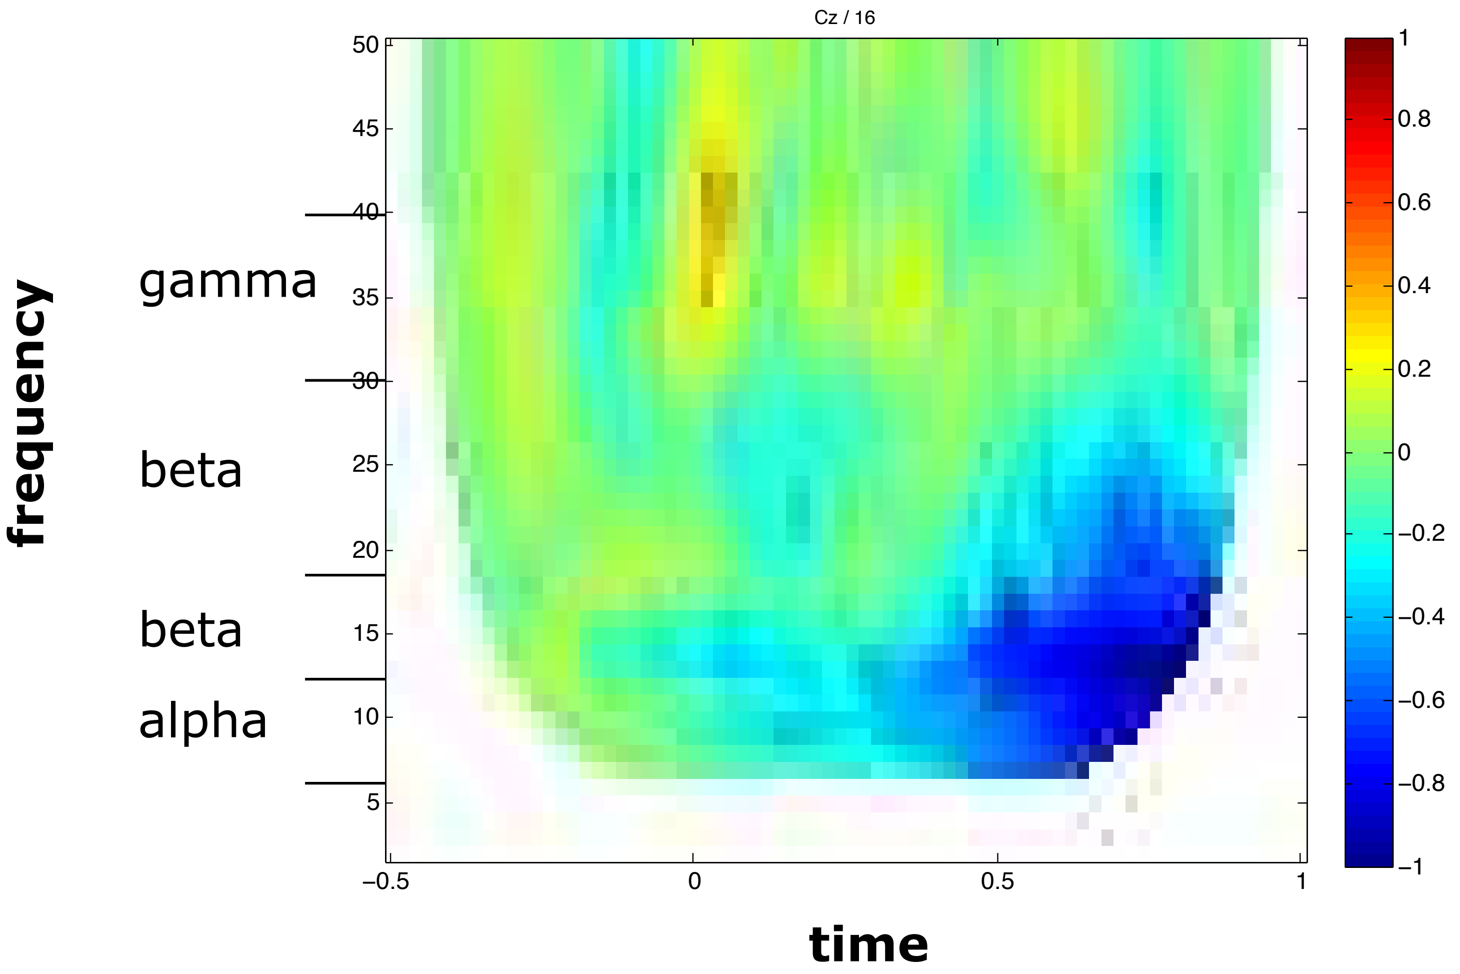

Agreement Violation (all electrodes)

The cameraman knew that the mayor would **honor** the soldiers...
The cameraman knew that the mayor would **honors** the soldiers...

Agreement

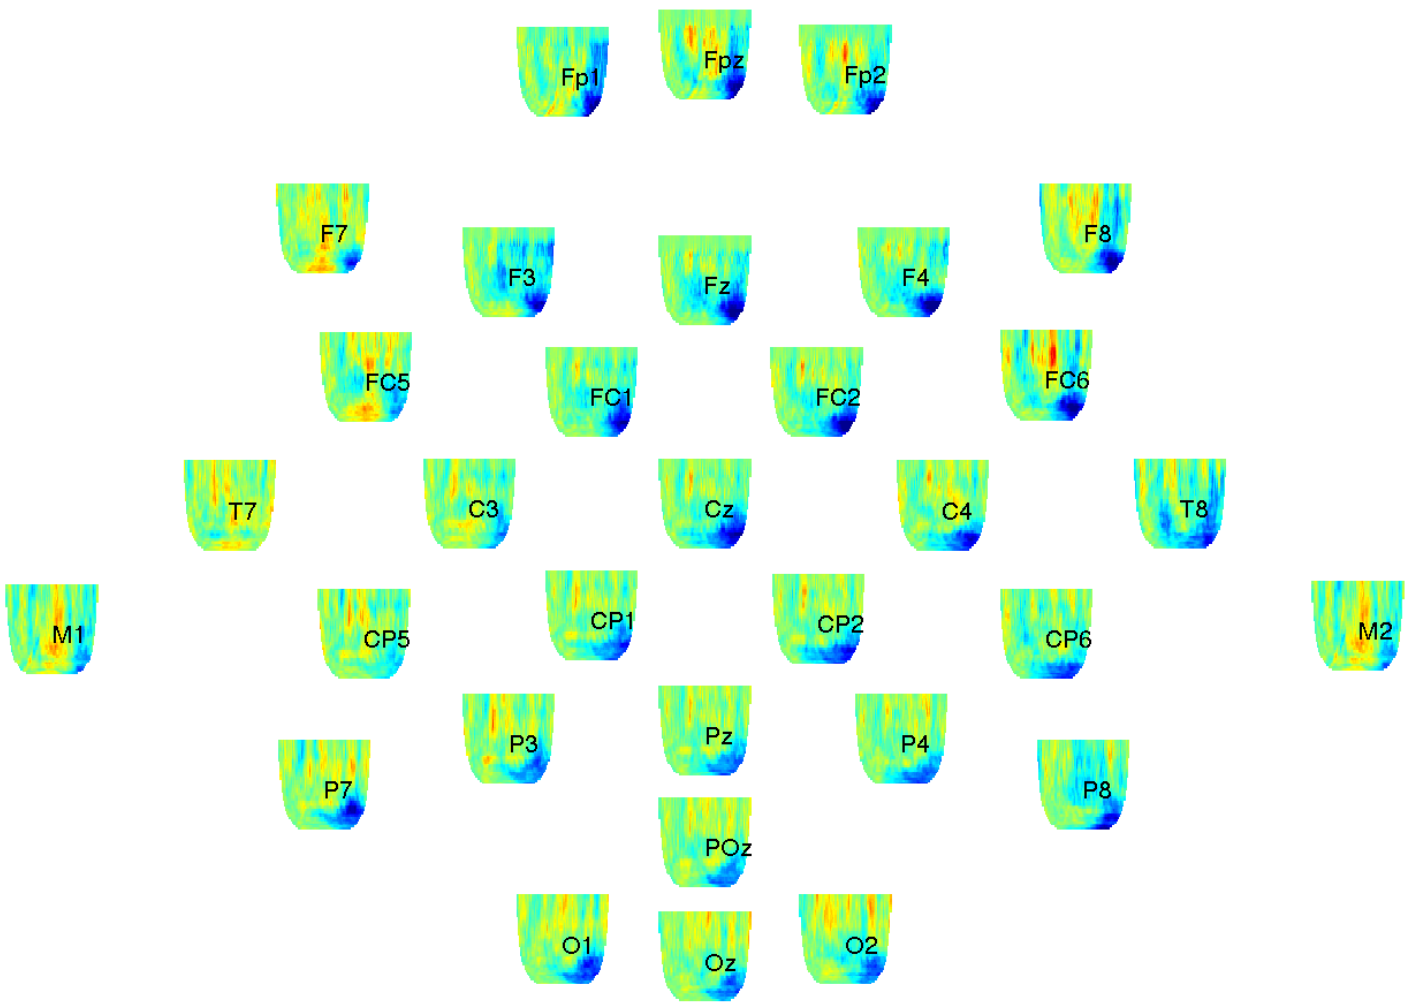

Theta Violation (electrode Cz)

The cameraman knew that the mayor would honor **the soldiers**...

The cameraman knew that the mayor would honor __ **before** ...

Theta

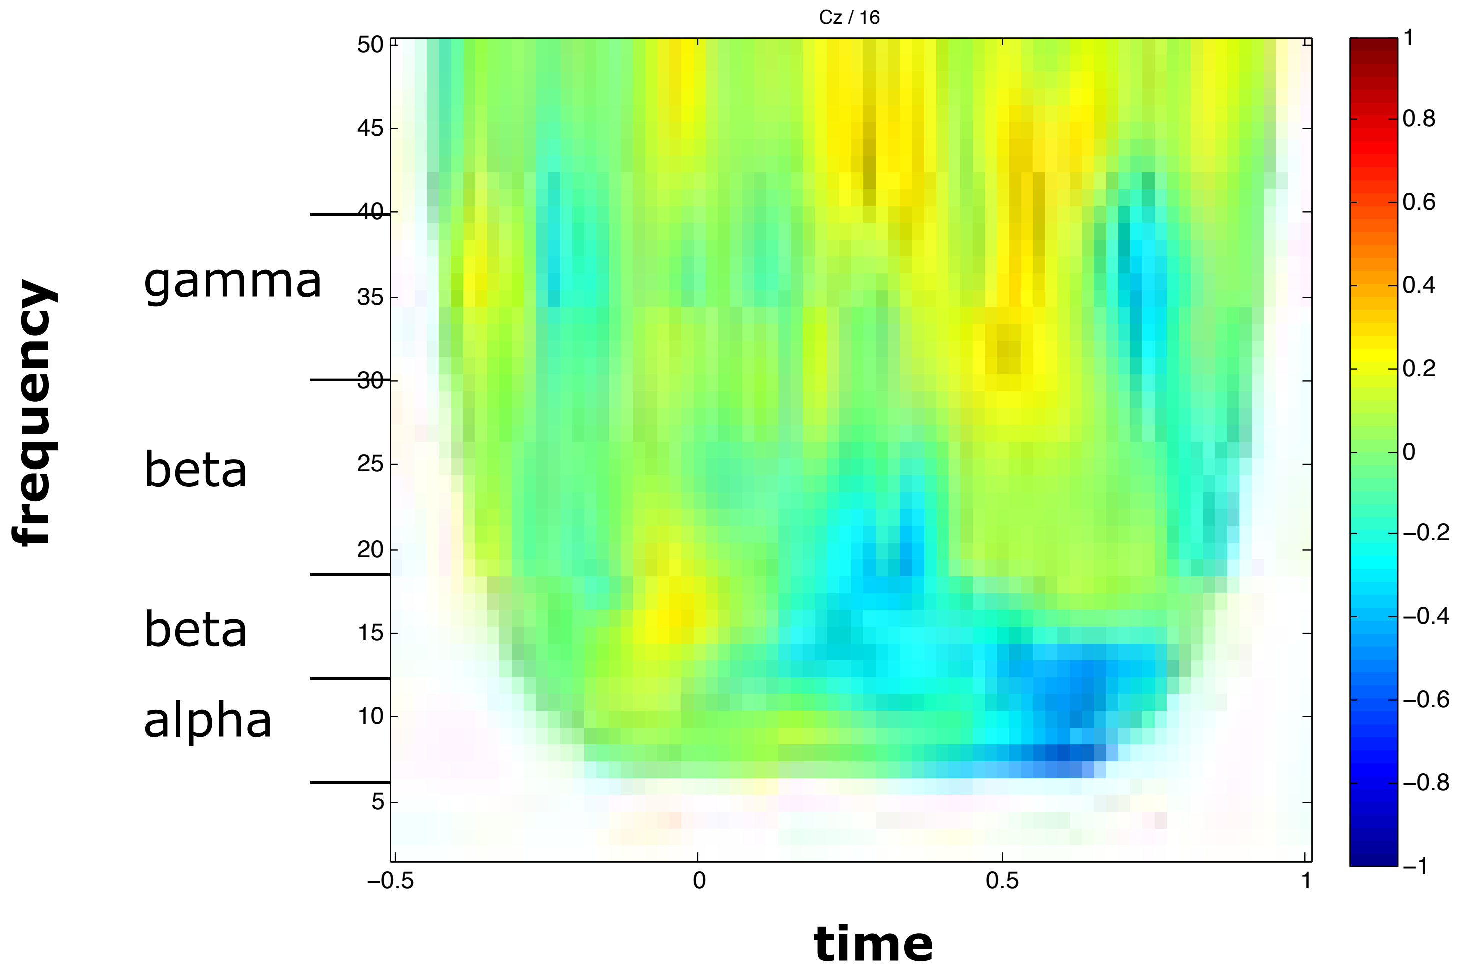

Theta Violation (all electrodes)

The cameraman knew that the mayor would honor **the soldiers**...

The cameraman knew that the mayor would honor __ **before** ...

Theta

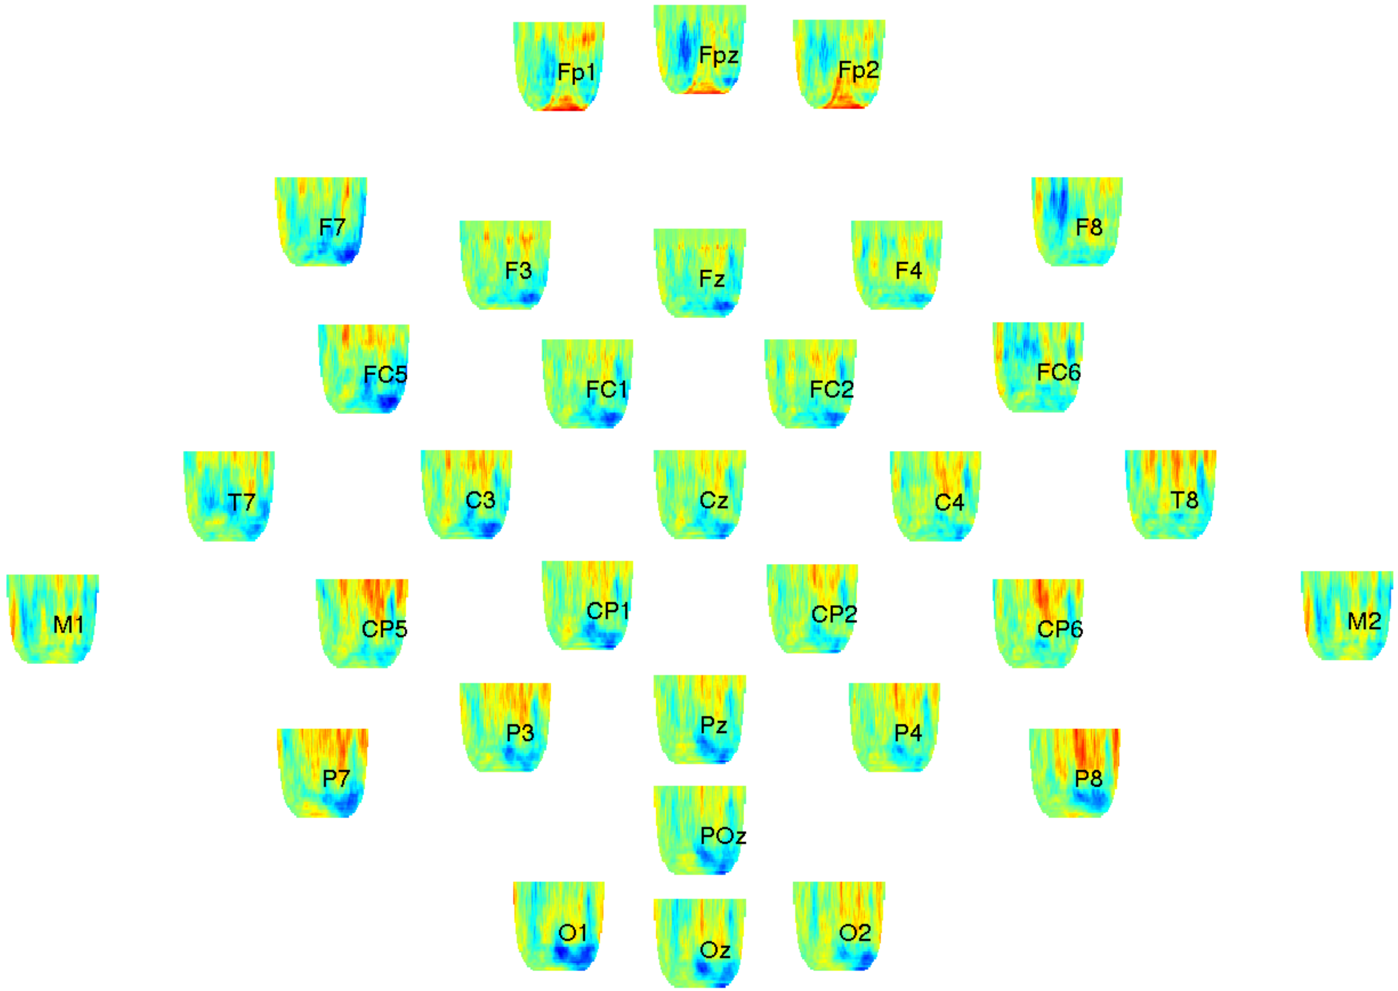

Study 2:
Look at normal syntactic structure-building

Violations are fun to study because they are generally big, robust effects. But ultimately we want to find the neuronal oscillations for syntactic structure-building, not just violation-detection.

One idea I have been playing with is trying to look for neuronal oscillations that are the same between two (or more) constructions that share a single structure and/or structure-building operation.

Unergatives: The child danced.

Unaccusatives: The towel dried <the towel>.

A-movement

Actives: The child was opening the present.

Passives: The present was opened <the present> by the child.

Looking across constructions has a lot of advantages (if it works), as it helps control for any processing differences that might be construction-specific.

Unaccusative - Unergative (electrode Fz)

Unergatives: The child danced.

Unaccusatives: The towel dried <the towel>.

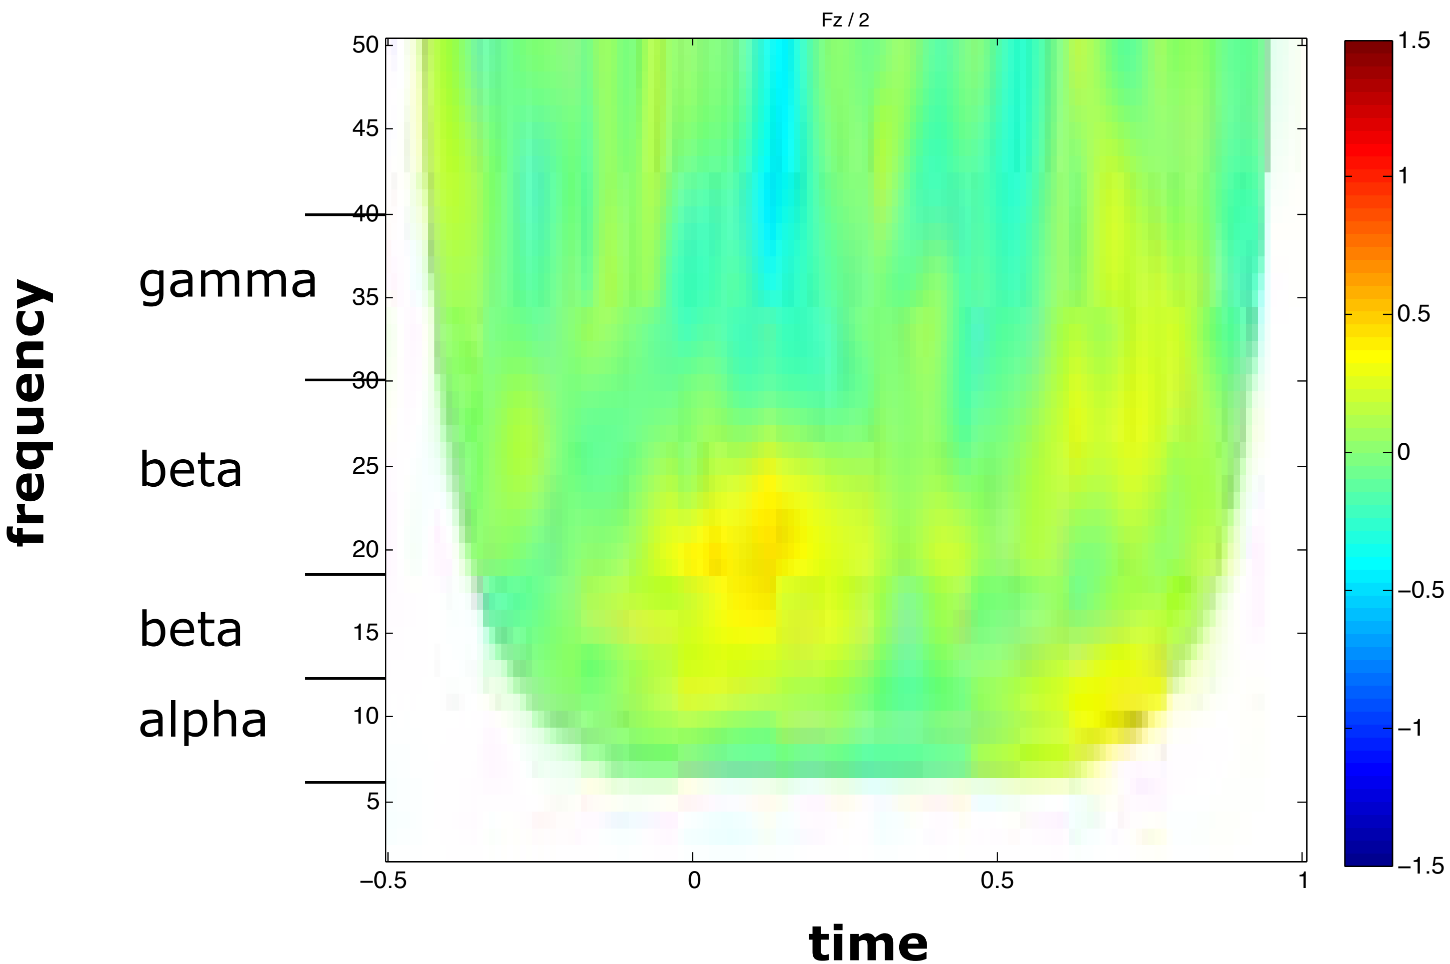

Unaccusative - Unergative (all electrodes)

Unergatives: The child danced.

Unaccusatives: The towel dried <the towel>.

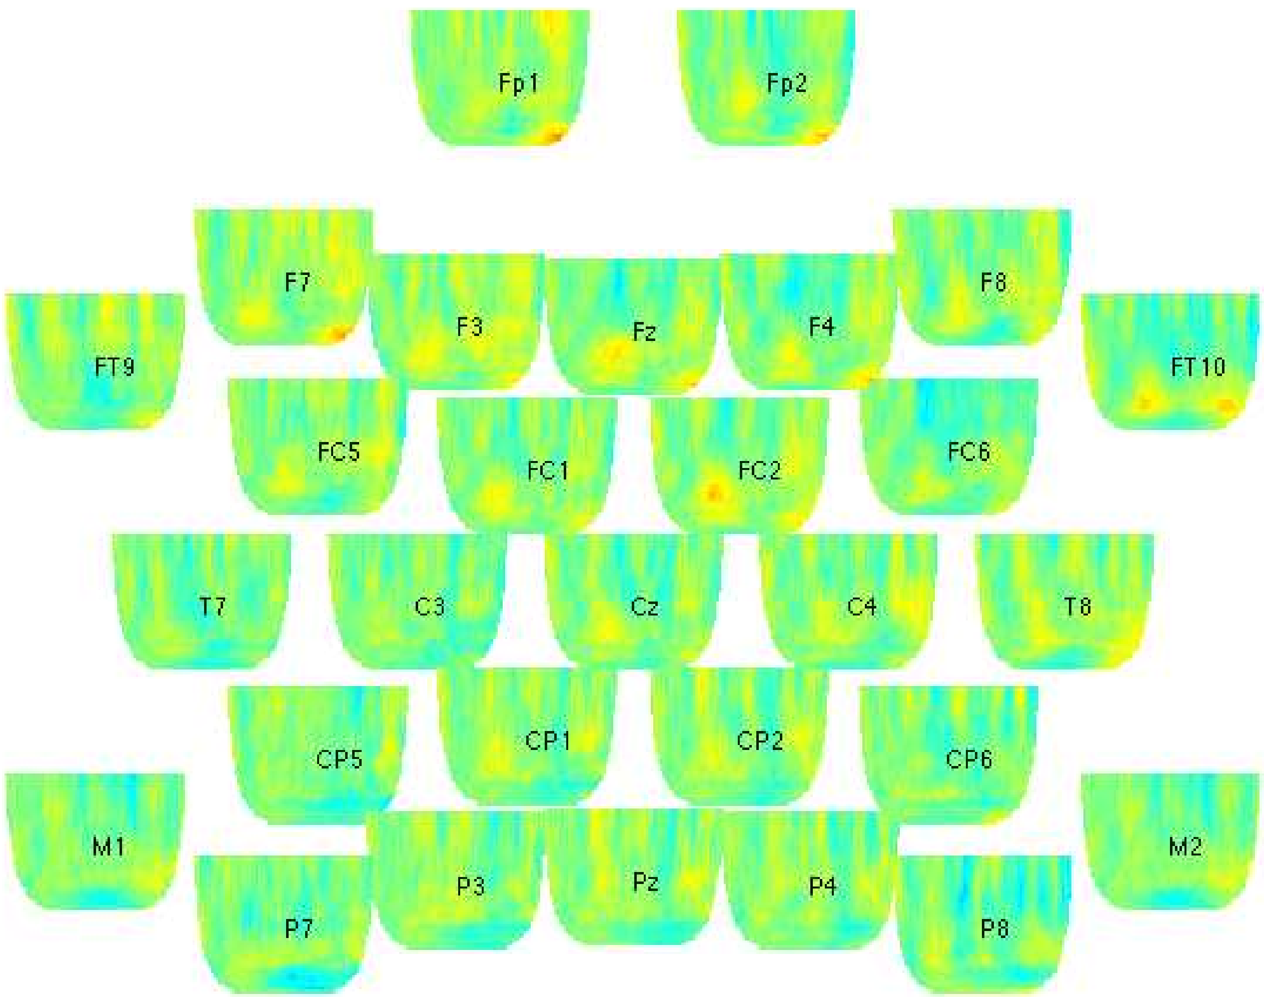

Passive - Active (electrode Fz)

Actives: The child was opening the present.

Passives: The present was opened <the present> by the child.

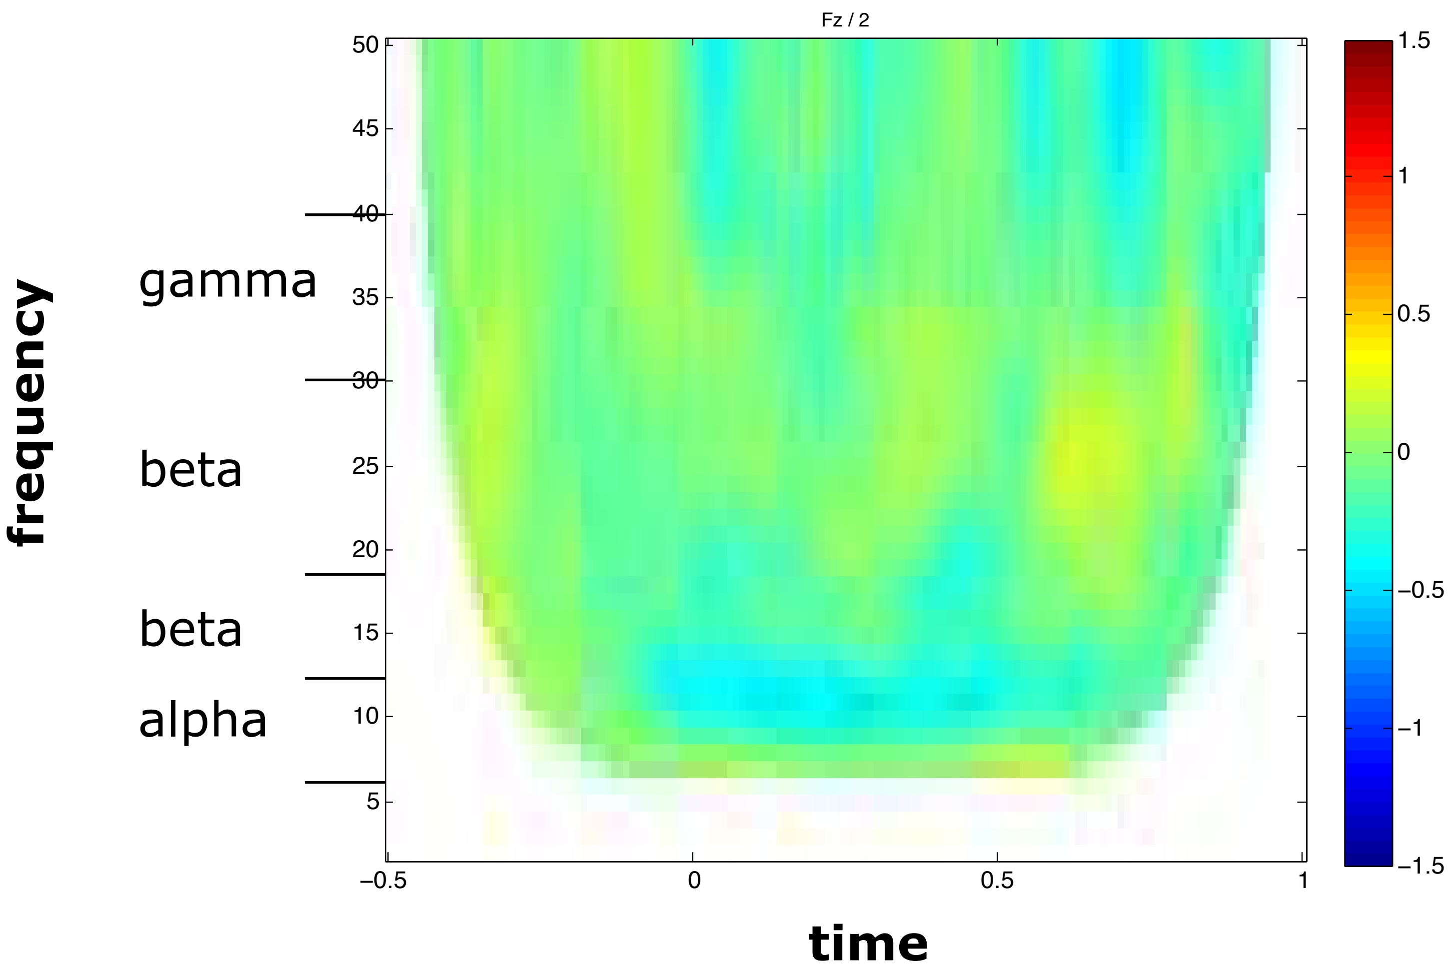

Passive - Active (all electrodes)

Actives: The child was opening the present.

Passives: The present was opened <the present> by the child.

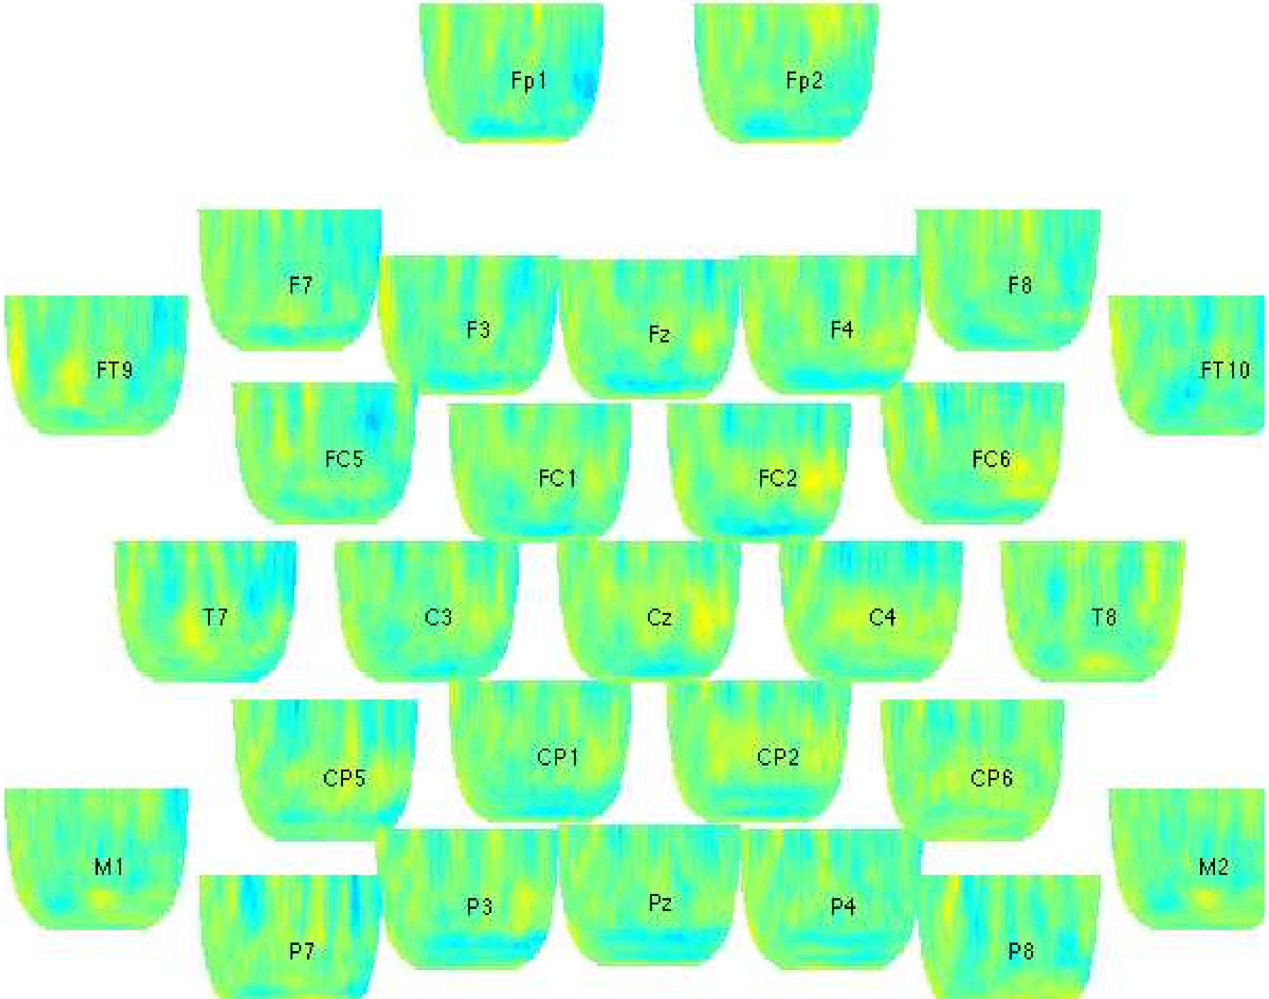

This is work in progress, but my currents thoughts are something like this:

Syntactic violations:

We see a decrease in alpha, lower beta, and sometimes upper beta. This accords well with some previous research on oscillations to syntax, and the decreases may indicate disruptions to the oscillations of syntactic structure-building.

Unaccusatives vs Unergatives:

We see a trend of an increase in alpha, lower beta, and upper beta. This tends to overlap with the decrease in violations, suggesting that unaccusatives might require a bit more oscillatory power to be constructed. We need to run more subjects to see if the trend continues.

Passives vs Actives:

We see a decrease in alpha. This is different from both the violations and the unaccusatives. There are two possibilities: one is that unaccusatives and passives are different after all. The other is that some other property of passives is overriding the (weak) beta increase that should come from A-movement. One possibility is that the passives are more predictable than the actives ('was' may more strongly predict passive than progressive).

The (longterm) plan

Islands: Continue to systematically test languages (and island phenomena) to characterize the patterns and effect sizes.

Acquisition: Continue to systematically build (bias-based) models of acquisition. Right now we are working on argument structure. The goal is to build a bunch of specific models to identify potential UG biases, and then try to incorporate them into a single model.

EEG and time-frequency: Continue to systematically test both violations and syntactically-important constructions to get a sense of the oscillatory responses for a range of violations, structures, and processes (raising, control, passives, binding, phases, Agree, etc).

As we have seen, it is quite possible that none of these projects will really pan out over the longterm. But I think that if we are truly going to take formal experimentation seriously, we need to see what the world looks like with a systematic data set in each case (not just one off “proof-of-concept” studies). Then we can evaluate the consequences for syntactic theory.

From Cognitive Theories to Neural Systems

	phonology	**morphology**	**syntax**	**semantics**
representations:	[+/- feature]	V + ed	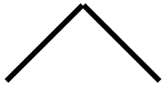	λx.λy
computations:	conversion	fusion/fission	merge/move	$f(x)/\wedge$

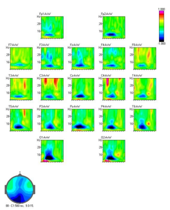

Use electrophysiology to identify large-scale neural correlates of representations and/or computations

Use neuroimaging to identify networks underlying electrophysiological correlates

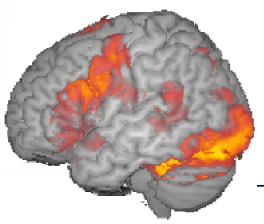

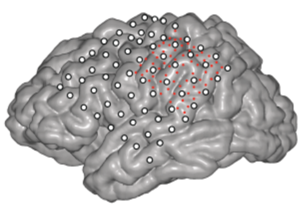

Use ECoG to identify single-unit activation in the relevant networks

Use theoretical/computational neuroscience to postulate plausible encodings of representations and computations.

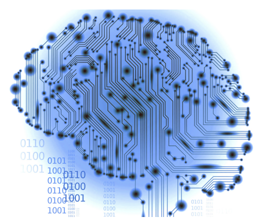

THANK YOU!
and thank you to my generous collaborators!

Jessamy Almquist
UC Irvine

Susi Wurmbrand
UConn

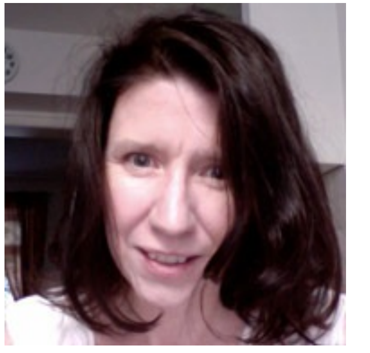

Greg Hickok
UC Irvine

William Matchin
UC Irvine

Extra slides about EEG

Electrodes are placed at several locations on the scalp

Scalp potentials differ at different locations on the scalp, therefore we place electrodes all over the scalp to record the different patterns of activity:

Each of these little cups contains an electrode:

Electrodes are placed at several locations on the scalp

There is a convention when it comes to naming electrodes. Each name corresponds to a location on the scalp.

F indicates frontal electrodes.

T indicates temporal electrodes.

C indicates central electrodes.

P indicates parietal electrodes.

O indicates occipital electrodes.

Odd numbers indicate left side.

Even numbers indicate right side.

Z indicates midline.

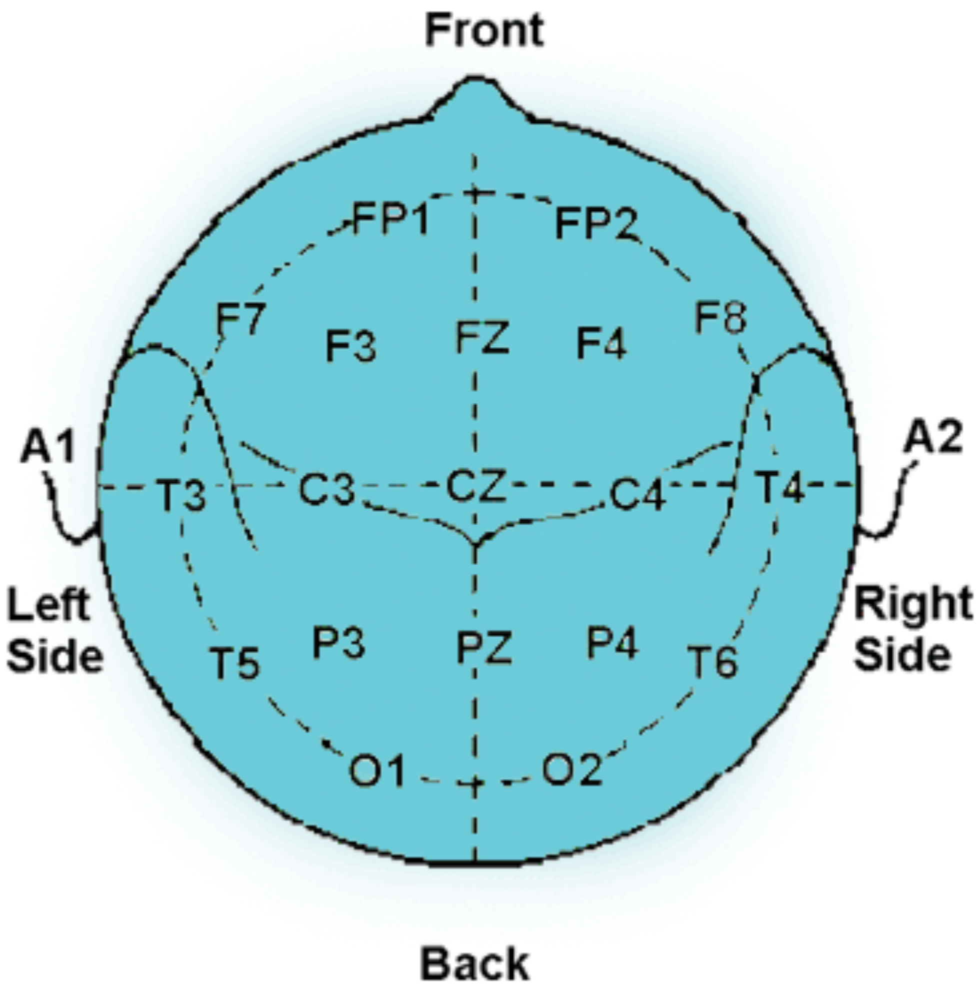

Electrodes are placed at several locations on the scalp

The electrode naming scheme is based on the names of the areas of cortex that are under the scalp:

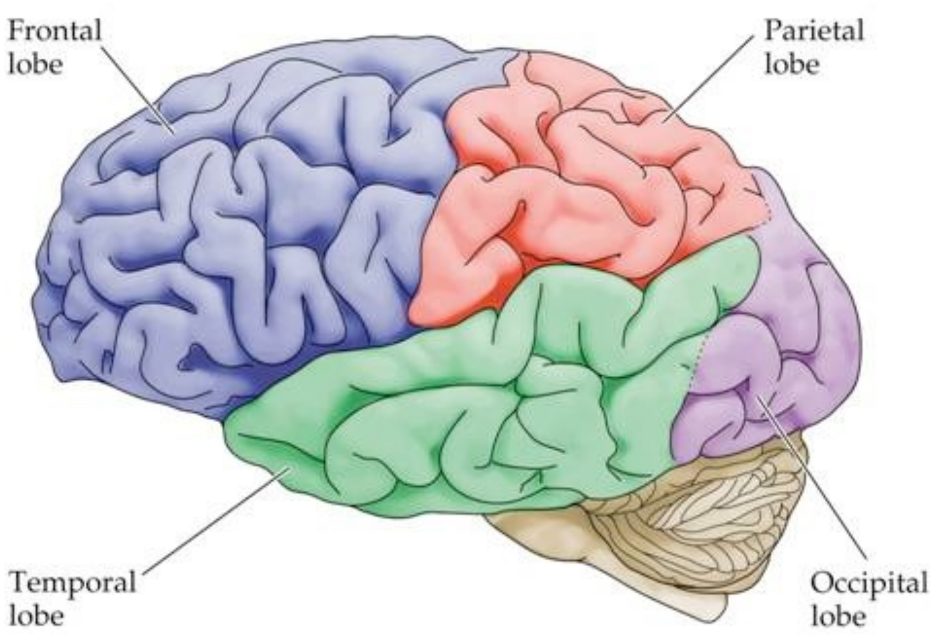

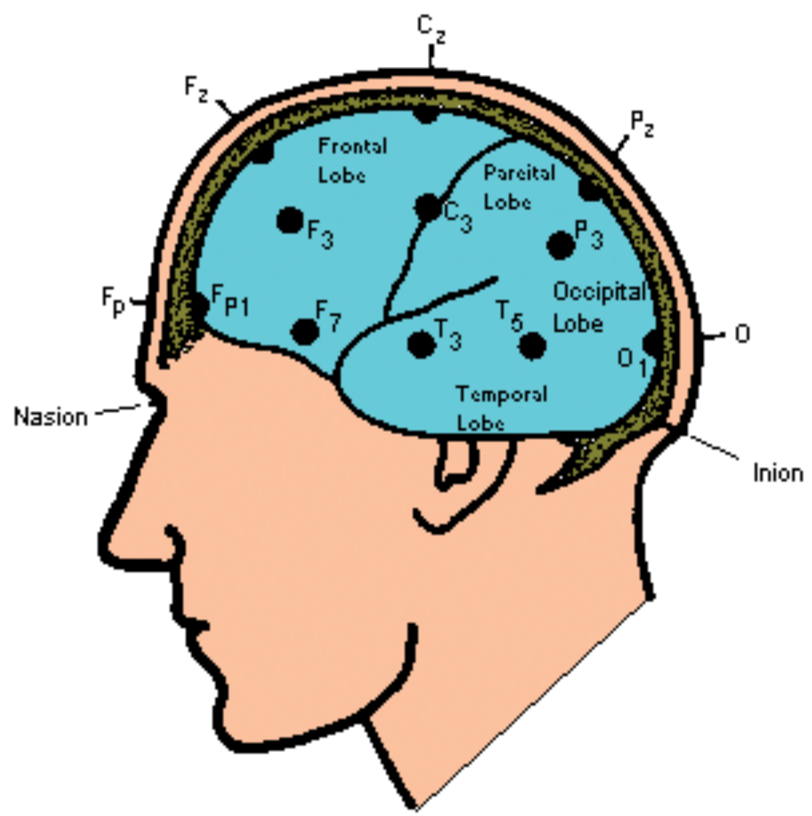

Electrodes are placed at several locations on the scalp

The (international) standard for electrode placement is known as the **10-20 System**, so named because the placement of electrode is based on 10% or 20% of the distance between the nasion and inion.

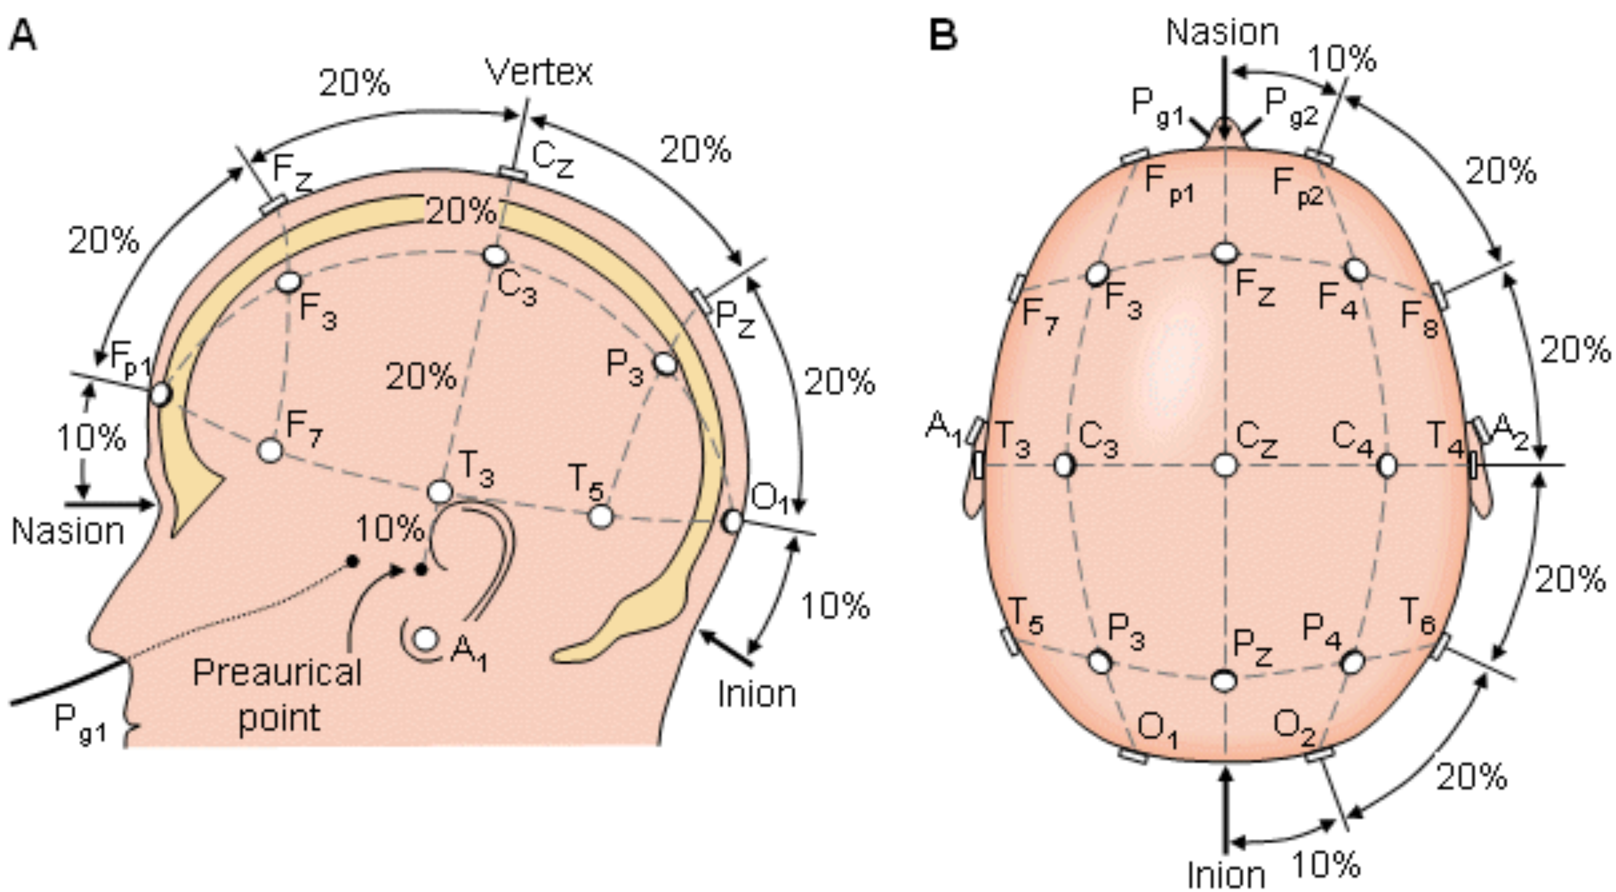

Very poor spatial resolution

What this means in practice is that we can talk about the scalp location of a potential, such as “at electrode site F7”.

But this does not mean that the potential was generated by the left frontal lobe.

Instead, EEG recordings remain agnostic about the cortical location of the underlying neural generators.

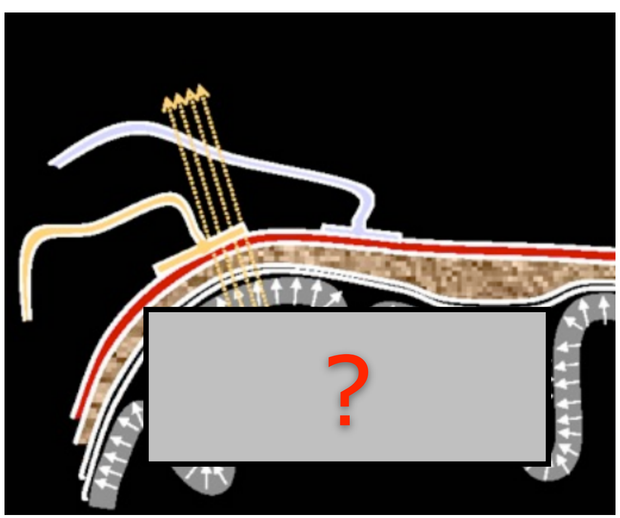

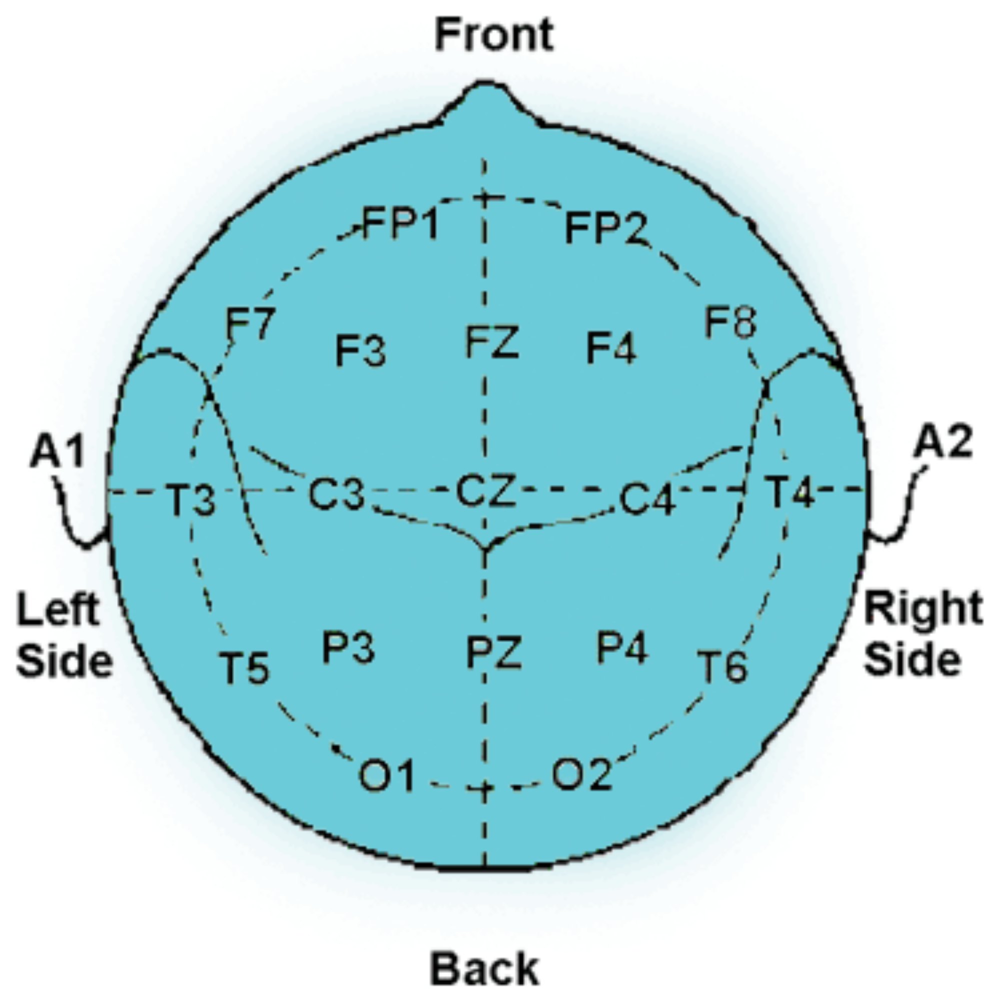

Waves also have a frequency

Frequency is a measure of the number of cycles that a wave completes in a given unit of time.

A complete cycle consists of peak, trough, and return to baseline.

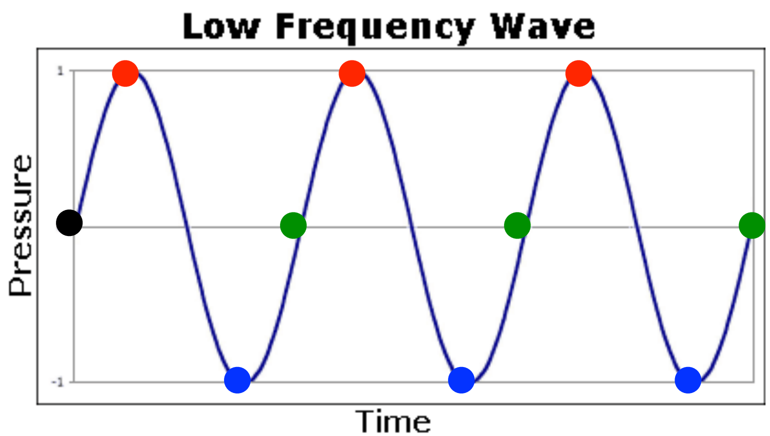

3 cycles per second

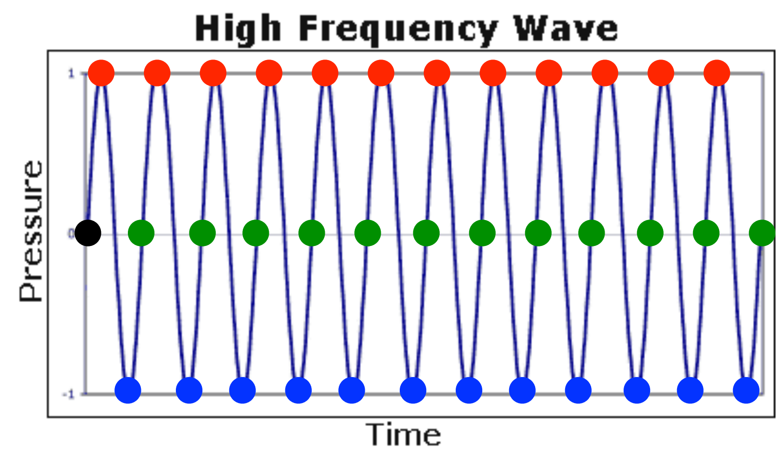

12 cycles per second

Frequency is measured in **Hertz (Hz):** cycles per second.