


Slide 1



Nate Silver, Sabermetrician and Presidential Prognosticator says that “The signal is the truth.” and “The noise is what distracts us from the truth.” In the study of the event-related potentials the signal and the noise are less distinct and far more deeply involved with one another. Like the knight and his lady in this painting made by Kandinsky before he moved to abstraction. This talk will consider the amorous relations between signal and noise.


Slide 2

Signals and Noise	
I Ways of Looking	
II Time Analyses	Time and Frequency Domains
III Frequency Analyses	Spectra and Spectrograms
IV Component Structure	Electroencephalogram
V Induced Activity	Transient, Sustained and Steady-State Responses
VI Nature of the Response	
VII Overview	

There are many ways of looking. Parts of my talk will overlap with what Davis Steve has told you yesterday and what San Diego Steve will tell you later today. However, seeing things from different points of view increases understanding. So I shall first consider some different ways of looking at what we record. EEG signals may be considered in both the time and frequency domains. In particular, we can watch as the frequencies change over time in a spectrogram. The spectrogram of the EEG then

shows synchronization and desynchronization. The spectrogram of the event-related potential shows transient sustained and steady-state responses.

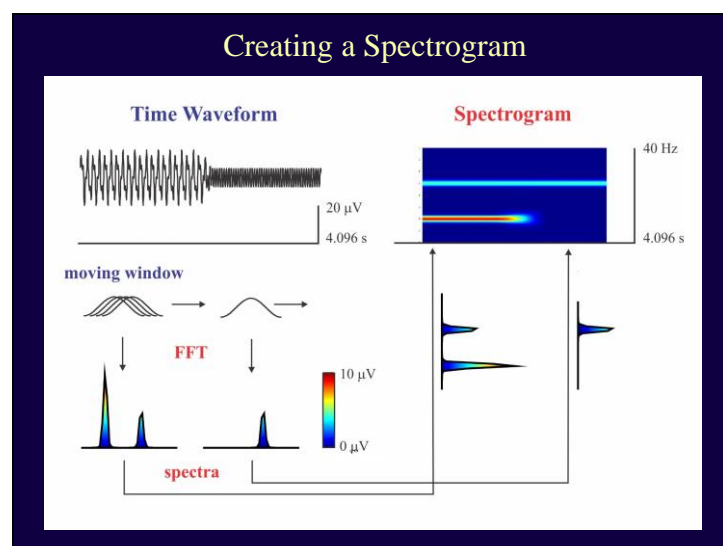
Slide 3



The EEG signal consists of a voltage that varies over time. Anything that occurs over time can also be considered in terms of frequency by submitting it to a Fourier Transform. Rapid temporal changes become fast frequencies and gradual changes become slow frequencies. In the EEG these are described using Greek letters: alpha beta gamma delta. Theta does not fit into the alphabet – it was thought to have come from the thalamus. Other rhythms are similar to the posterior alpha rhythm – mu in the somatomotor regions and tau in the temporal lobes.


This spectrum plots the amplitude at different frequencies. The full Fourier transform involves both amplitude and phase. Phase is important – it explains when the different frequencies occur in the temporal signal.

Slide 4



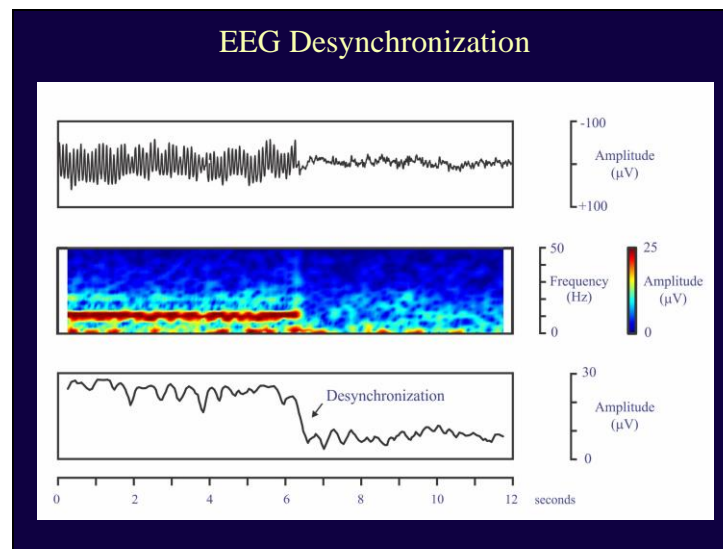
The frequency spectrum of a signal can change over time. In order to consider these changes we look at a spectrogram. This involves estimating spectra at multiple times during the signal. This is a simple model of the EEG that begins with a mixture of alpha and beta activity and then changes to just beta activity. We record multiple spectra, convert their amplitudes to a color scale, and then plot them in a three dimensional array with time on the x axis, frequency on the y-axis and amplitude on the z-axis.

Slide 5



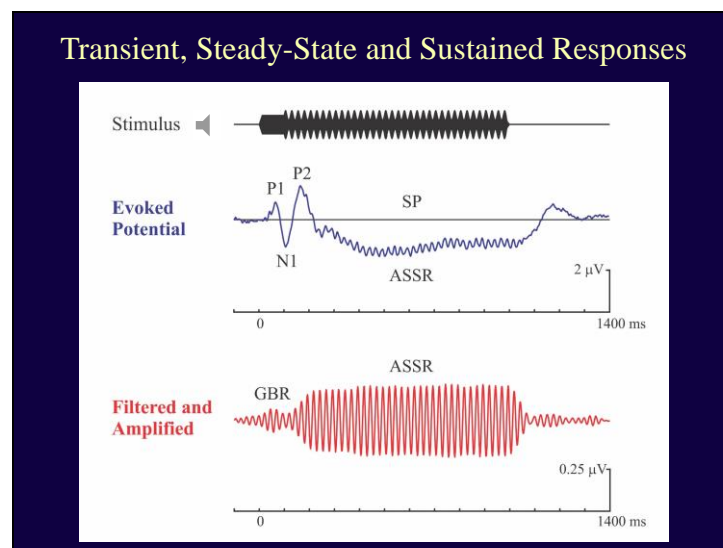
The spectrum of the human EEG signal contains many different frequencies. In this plot the frequency scale is logarithmic. The decrease in amplitude with increasing frequency is related to the amount of charge separated across membranes during the activity, the synchrony of the potentials, the distance from the generators, and the geometry of the head. Faster signals are better recorded from the cortical surface than from the scalp where multiple signals are averaged together. The scalp-recorded EEG is mainly generated in the brain – in the chatter of the axons, the rumble of the synapses and the hum of the glia. However, other parts of the head can also generate potentials and we must be careful to separate their effects from our signal if we wish to study thinking rather than blinking.

Slide 6



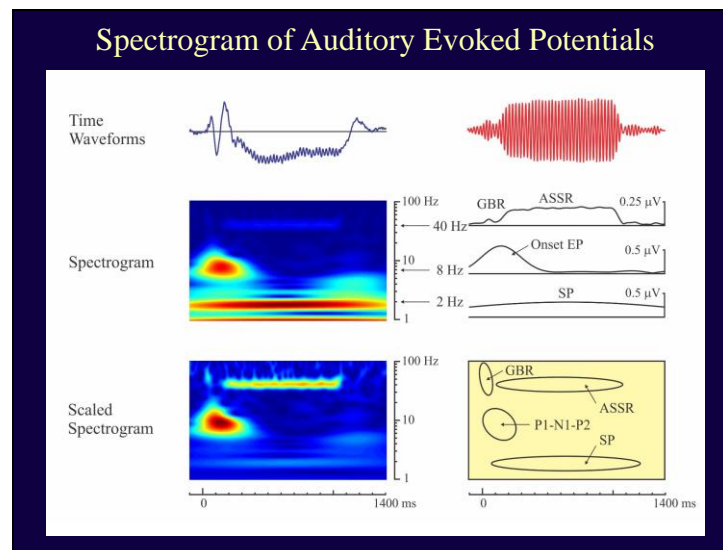
This is a real EEG signal recorded from the occiput. Midway through the recording the eyes open, visual input floods the cortex and the alpha rhythm goes away. The alpha rhythm is caused by synchronous oscillations between idle neurons. Ready to work but without anything yet to work on. Think of the rhythmic movements of a tennis player waiting for the serve. The decrease in the alpha rhythm is therefore called desynchronization.

Slide 7



We can also use spectrographic techniques to evaluate the evoked potentials. This is the response to a tone that after 100 ms becomes amplitude modulated at 40 Hz. The EP contains an onset-response, a sustained potential, and an auditory steady-state potential. If we filter and amplify the waveform, we can see the steady state response more clearly and also detect a small gamma-band response to the onset of the tone.

Slide 8



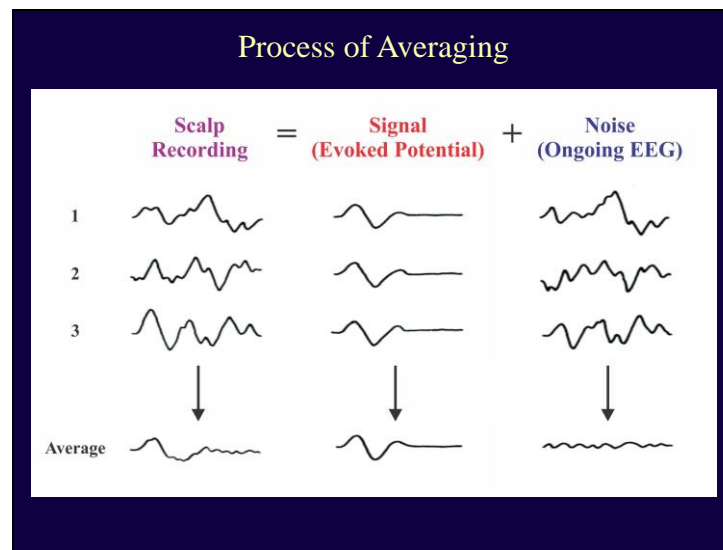
This shows the spectrogram of the evoked potential. As we noted before, the amplitude of the signal decreases with increasing frequency and the responses at higher frequencies are difficult to see. The spectrogram can be scaled to the amplitude in the period before the response to accentuate changes in these high frequencies.

Slide 9

Signals and Noise	
I Ways of Looking	
II Time Analyses	Process of Averaging
III Frequency Analyses	Latency Jitter
IV Component Structure	Artifact Rejection
V Induced Activity	Selective Averaging
VI Nature of the Response	Efficiency
VII Overview	

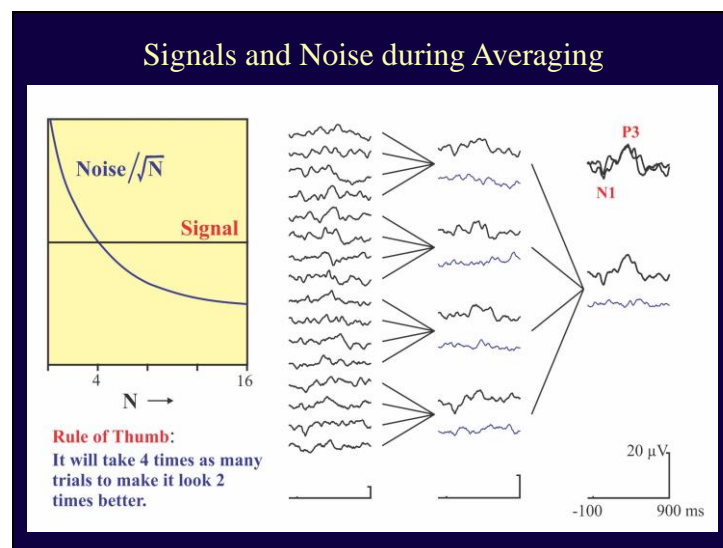
Now we have learned the ways of looking, become the masters of the spectrum and the lords of time. We shall now review the idea of averaging a time signal. We shall consider averaging a frequency signal later. Averaging is powerful but not perfect. We need some cautions and we need some tricks to make it better.

Slide 10



Averaging depends on the idea that what we record from the scalp in response to a stimulus is a combination of the evoked potential specifically related to processing the stimulus and the ongoing EEG related to all the other things that the brain is doing at the time that the stimulus occurs. Signal + Noise. With averaging the signal remains the same, the noise slowly cancels itself out, and finally the signal becomes visible in the average recording.

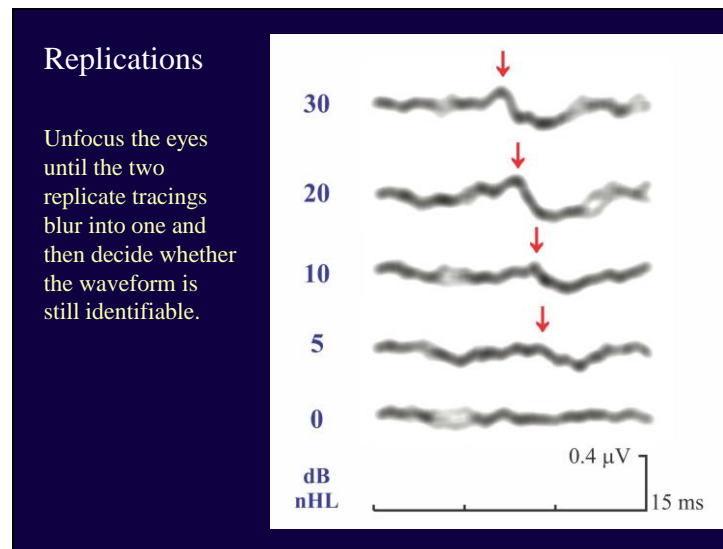
Slide 11



This shows the process of averaging. The noise decreases according to a square root rule. This is the same as going from standard deviation of the sample to standard error of the mean. Since the noise never goes away, we always have to consider if what we have recorded is a real signal or just unaveraged noise. One way of doing this is to calculate the +/- average. If we alternately add and subtract the waveforms prior to averaging we can cancel out the signal. The noise is pretty much the

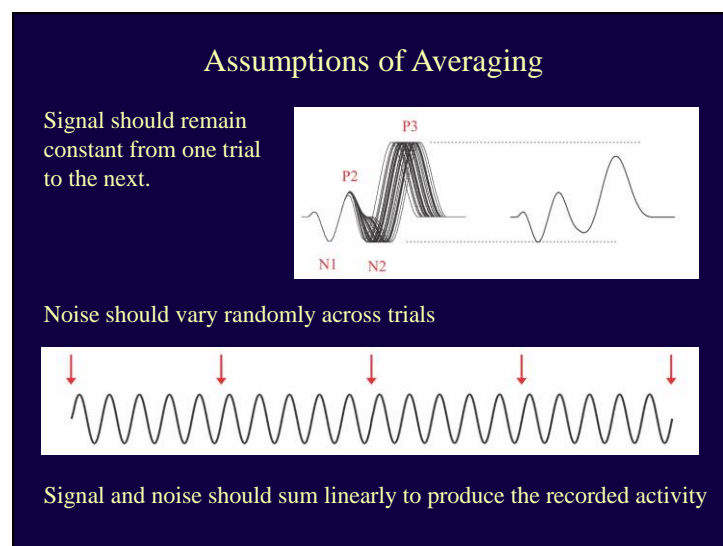
same right way up as upside down and so the blue \pm average gives us an estimate of the level of noise if there were no signal present. After averaging four responses the response is slightly bigger than the noise. After averaging 16, the average is clearly bigger than the \pm average. The figure changes the scale by a square root factor – so the noise stays constant and the signal gets bigger. Another way to evaluate the signal-to-noise is to look at replications. The distance between the replicate tracings measures the noise in the same way as the \pm average.

Slide 12



This shows some auditory brainstem responses near threshold. If we wish to determine whether a response is present or not we can consider the replications. The old man's way of doing this is to take of his glasses. A response is real if it remains recognizable in the blur.

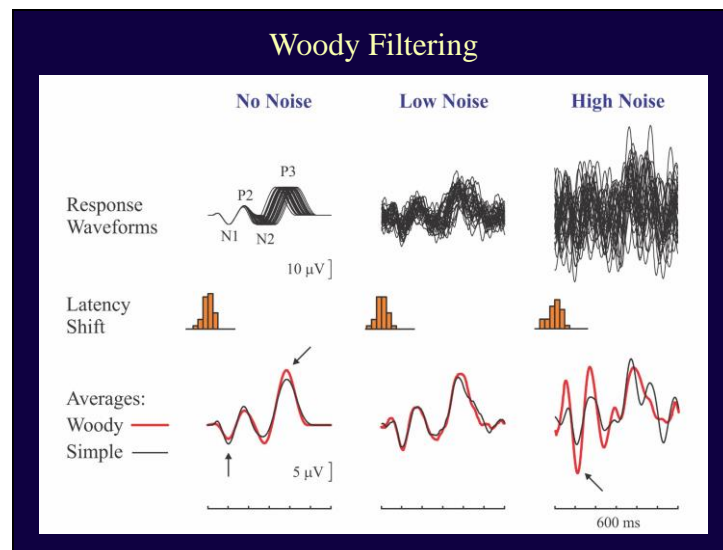
Slide 13



Averaging is based on three assumptions. The first is that the signal must remain constant from one trial to the next. Amplitude variability is not bad – the average response simply averages the amplitude. Latency variability is more bothersome – the average can become quite distorted. Amplitudes are decreased and the waveform is smeared out over time.

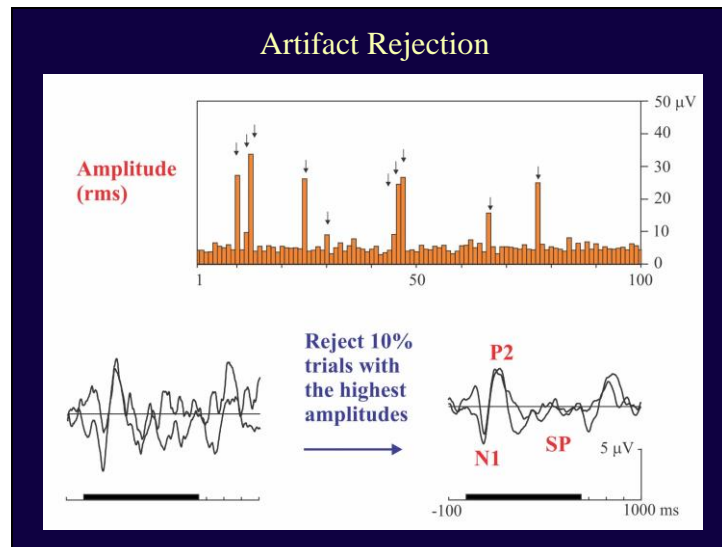
The second assumption can be a problem if the background activity (or artifactual noise) becomes somehow locked to the stimulus. One way to prevent this is to randomize the timing of the stimulus. The third assumption seems tautological – clearly the signal differs from the noise. But this may not be true. Sometimes the signal and the noise interact – we shall return to this later in the talk.

Slide 14



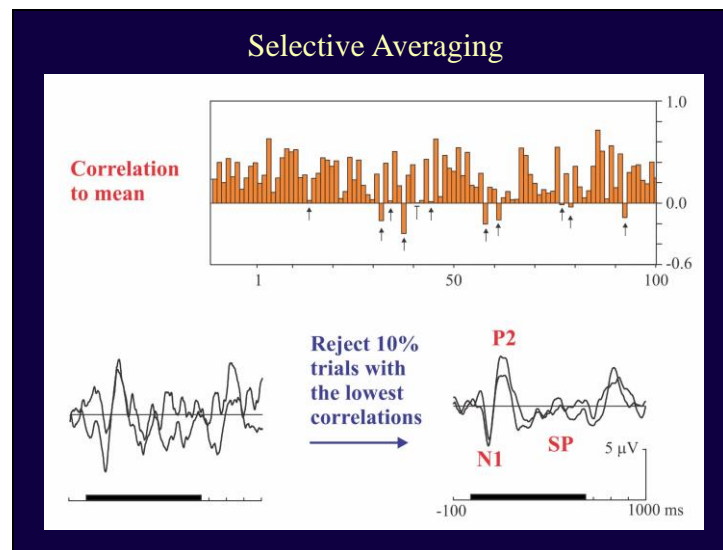
To return to the latency problem. One technique that can counter the effects of latency jitter is to correlate each individual waveform with the expected average waveform and then adjust the latency to maximize the correlation. It works well when the jitter affects the whole response. If only part of the waveform is jittered, however, correcting the latency for this part may distort the rest of the waveform. The averages are doubled in size to show the changes. While the P3 is corrected, the N1 is decreased. The process is also susceptible to noise. At high noise levels the N1 is actually accentuated as the Woody filter adjusts the latency of noise rather than signal.

Slide 15



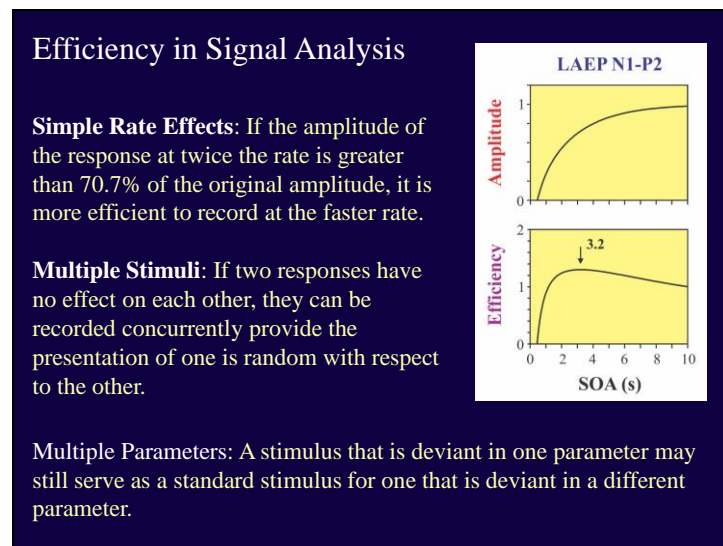
It is customary to reject from the averaging process those trials that have high amplitude. Since the signal is relatively constant, this is related to increased noise – as the subject blinks, moves, swallows, and tries to be a normal human being rather than an experimental subject. Rejection can be based on absolute amplitude or on the highest percentile.

Slide 16



We can also reject trials on the basis of how poorly they correlate with the average waveform. This seems like cheating but works very well.

Slide 17

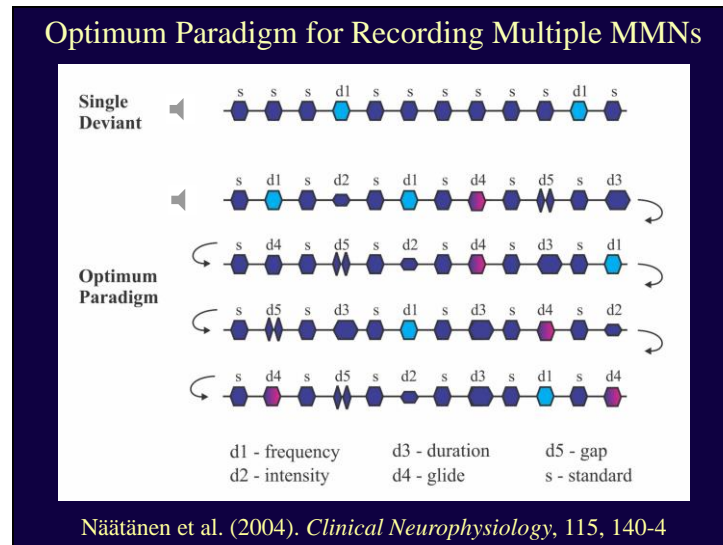


Efficiency is a worthwhile concept to bear in mind when averaging. When recording a single response we should use a rate that gives the largest signal-to-noise ratio in a given amount of time. This can be figured out by dividing the amplitude by the square root of the stimulus onset asynchrony. The late auditory EP is most quickly recorded when the stimuli come every 3 seconds or so.

We can also record responses to multiple stimuli simultaneously if the responses do not interact. If the stimuli are random each response becomes simply part of the noise for averaging the other responses.

The final example of efficiency occurs when recording the mismatch negativity. The conventional recording presents the deviant stimulus at low probability. However each deviant stimulus may serve as the standard stimulus for other types of deviant. By intercalating the deviants, we can record the response to five deviants in the same time as it takes to record one.

Slide 18



The conventional paradigm presents one deviant in a train of standards with a probability of 1/10. The sounds do not follow the illustration other than in principle. To record another type of MMN we need to repeat the whole recording.

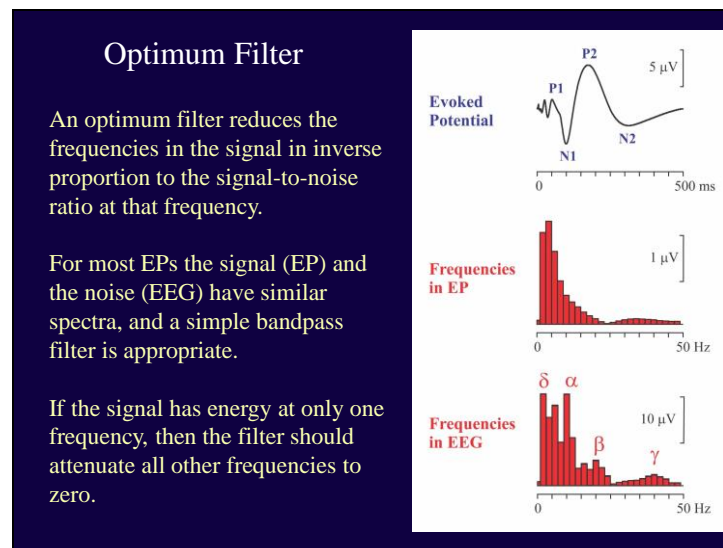
The multiple deviant paradigm presents each deviant at the same temporal probability but presents five deviants concurrently. Five mismatch negativities for the price of one.

Slide 19

Signals and Noise	
I Ways of Looking	
II Time Analyses	Optimum Filter
III Frequency Analyses	Steady-State Responses
IV Component Structure	Following Responses
V Induced Activity	Multiple Simultaneous Stimulation
VI Nature of the Response	Frequency-Tagging
VII Overview	

So we have looked at averaging in the time domain – and considered such things as the estimation of noise and the efficiency of the recording. How do we analyze signals in the frequency domain?


Slide 20



Filters remove unwanted frequencies from the recording. An optimum filter ... For most EPs ... Thus for the slow auditory evoked potential we can use a bandpass of 1 to 30 Hz.

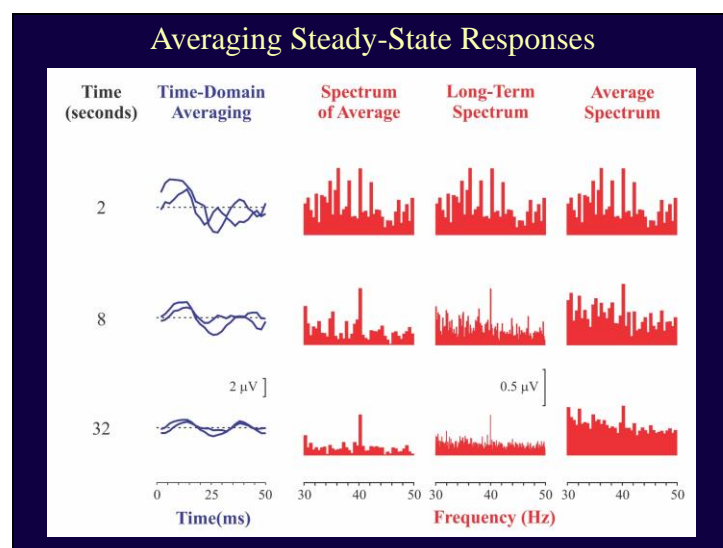
Steady-state responses occur at only one frequency and can be efficiently recorded using filtering. The most effective filter is to look only at the exact frequency of stimulation and reject all others.

Slide 21



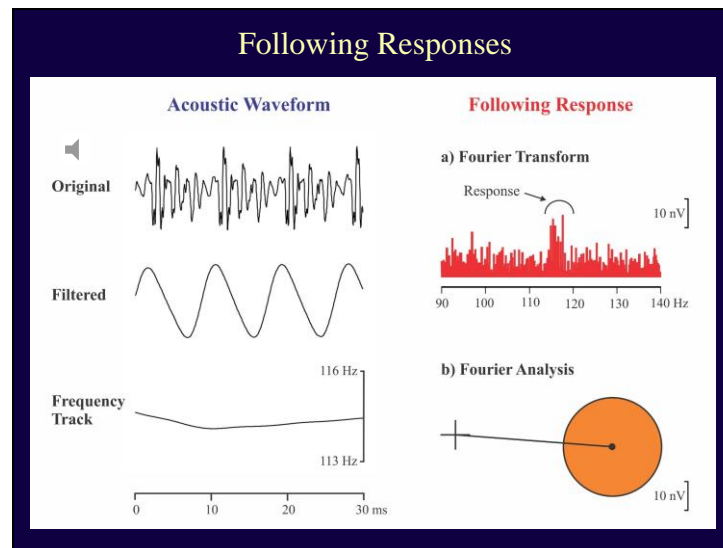
This is the steady-state response of the brain to a low-frequency tones amplitude modulated at 40 per second. The response averaged in the time domain shows a sinusoidal pattern that follows the envelope of the stimulus. It maintains its amplitude and phase throughout the presentation of the sound – making this particular envelope-following response also a steady-state response. The response viewed in the frequency domain shows a prominent single line at the frequency of the stimulus. The other parts of the spectrum show the background EEG activity – alpha, beta. The amplitude spectrum does not contain the phase information – this can be seen in a polar plot where the amplitude is the length of the line and the phase its orientation relative to the axes. Phase is important – it determines the delay of the response relative to the signal.

Slide 22



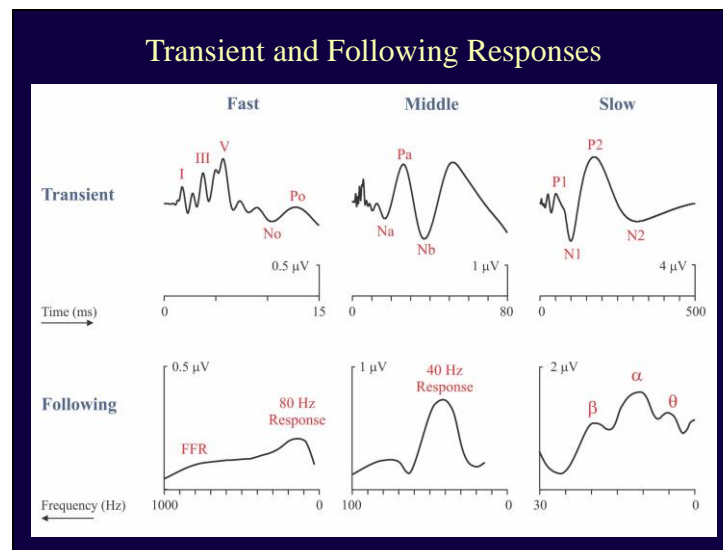
Steady state responses can be averaged in either the time or frequency domain. This shows a 40-Hz response averaged over 2 to 32 seconds. We can take spectra of the time-domain average, or we can just take a spectrum of the whole recorded signal. Only a portion of the spectrum is plotted – from 30 to 50 Hz. In either case the response is clearly visible at 40 Hz. If we take multiple amplitude spectra from every 2-seconds worth of the recording we can then average these spectra. This averaging is performed without regard to phase. The resultant spectrum represents the average of what you can see in 2 seconds, not what you can see over the whole recording period. It represents the background activity more than the signal, which is only just visible. We shall return to this later and you will hear again the mantra that the spectrum of the average is not the same as the average spectrum.

Slide 23



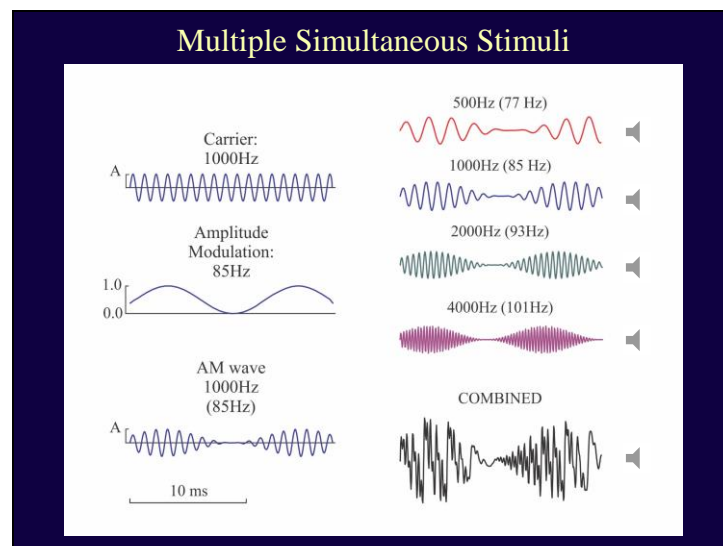
The frequency of a recurring stimulus in real life may not remain steady. This is prolonged vowel 'a' spoken by Steve Aiken. Its frequency varies over time. If we use the Fourier transform we can see a response but it is smeared across multiple frequency bins. Other techniques must be used to measure the response that follows a natural vocalization. This shows the results of a Fourier Analysis (as distinct from a Transform) – the orange area represents the standard error of the mean in two dimensions.

Slide 24



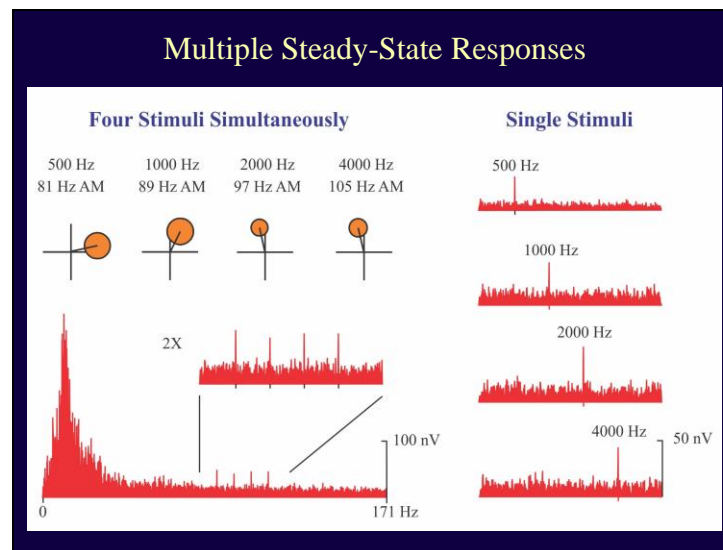
Using such techniques, following responses can be recorded at all frequencies of stimulation. This plot shows the transient evoked responses to a simple click. The response begins in the ear continues through the brainstem to the cortex and then spreads across the association areas of cortex. The following responses are plotted backward – from high to low frequency. The slow following is related to the slow transient responses and fast following is likely mediated by the same neurons as generate the fast transient response

Slide 25



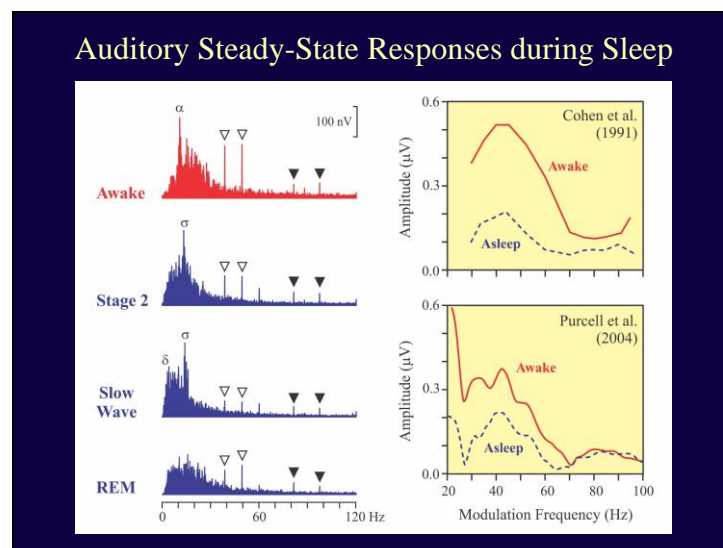
Steady state responses allow us to increase the efficiency of our recordings. This shows how we can construct multiple different stimuli with each carrier frequency having a specific modulation frequency. Then we combine the stimuli and record the brain's response.

Slide 26



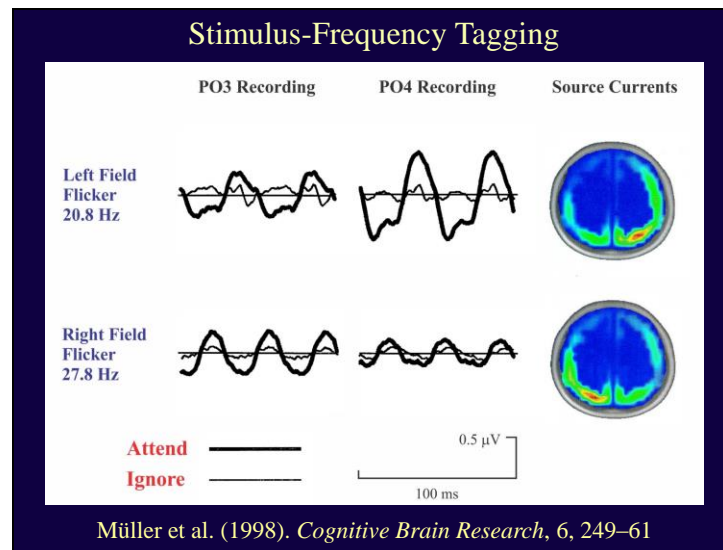
This is the complete spectrum. At the modulation-frequencies of each stimulus we can see a response. The region around these frequencies is amplified. Each response can also be viewed on a polar plot. If we present each stimulus by itself we find is no difference in amplitude from when the stimuli were presented concurrently. It is therefore much more efficient to record the responses to multiple stimuli presented simultaneously than to record multiple separate responses.

Slide 27



This slide shows the response to stimuli modulated at frequencies near 40 Hz and at frequencies near 80 Hz. When a subject falls asleep, the responses near 40 Hz decrease, but the responses at higher frequencies are unaffected. This combination of stimuli might be used to monitor the state of arousal of a subject – the 40-Hz responses show the state and the faster responses show that the ear and brainstem are processing sounds normally.

Slide 28



Tagging stimuli with a signature frequency can be used to track attention. This shows some responses in a paper that Old Steve and Old Terry were both involved in. Stimuli were presented concurrently to the left and right visual fields. Responses that are specific to the frequency-tagged stimuli vary with whether the stimuli are attended or ignored. Responses to stimuli in the left visual field are specifically enhanced by attention to the left and have their sources in the right parietal area and vice versa. Old Steve will consider such studies in greater detail.

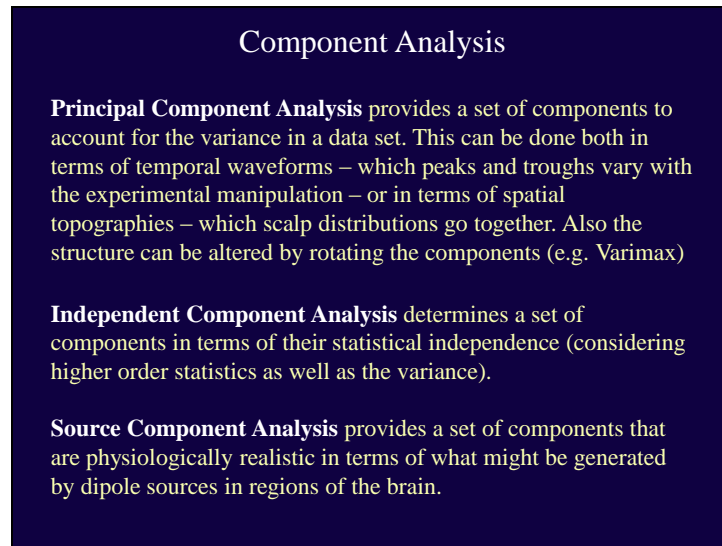
Slide 29

Signals and Noise	
I Ways of Looking	
II Time Analyses	Component Analyses
III Frequency Analyses	Scalp Topography
IV Component Structure	Dipole Source Analysis
V Induced Activity	Ocular Artifacts
VI Nature of the Response	Three Wizards
VII Overview	

We have reviewed the signal and the noise in both the time and frequency domains. We shall now consider how we decide what the signal means. Understanding speech involves parsing out the words

that make up what we hear. Understanding the evoked potentials involves separating out the different components that make up what we record. We shall look at some dipoles and talk to some wizards.

Slide 30



Component Analysis

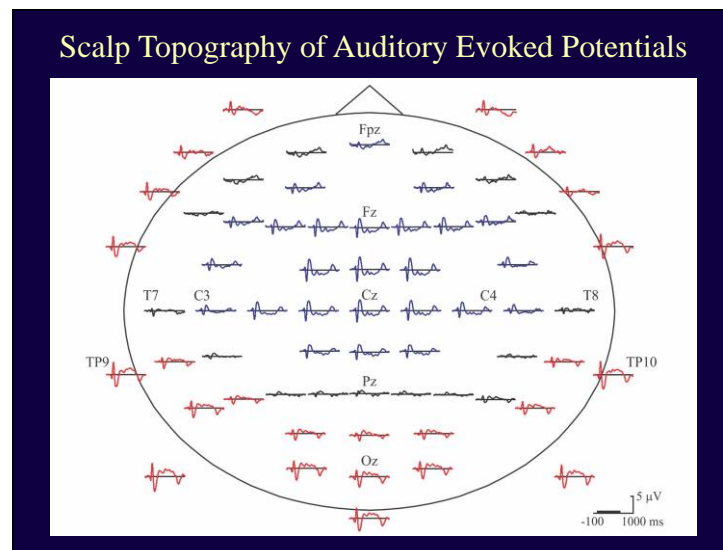
Principal Component Analysis provides a set of components to account for the variance in a data set. This can be done both in terms of temporal waveforms – which peaks and troughs vary with the experimental manipulation – or in terms of spatial topographies – which scalp distributions go together. Also the structure can be altered by rotating the components (e.g. Varimax)

Independent Component Analysis determines a set of components in terms of their statistical independence (considering higher order statistics as well as the variance).

Source Component Analysis provides a set of components that are physiologically realistic in terms of what might be generated by dipole sources in regions of the brain.

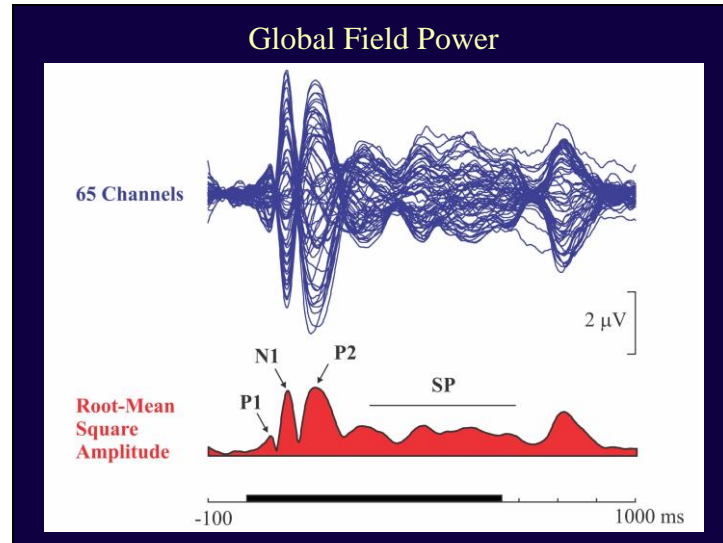
The components of the event related potentials can be analyzed in several ways. Davis Steve will talk more about this later, but this will present my way of looking and my biases. Principal component This statistical approach first came to prominence in the study of personality factors. Independent component ... PCA is a lower order version of ICA. ICA originated in the discrimination of different auditory sources – how the CIA could listen to one conversation and not others. My bias is toward source components. I would like to see the when and where in the brain of what is going on. However, source component analysis is not easy – it is best done following a preliminary analysis of how the waveforms react to experimental manipulation. Then physiology and psychology can work together.

Slide 31



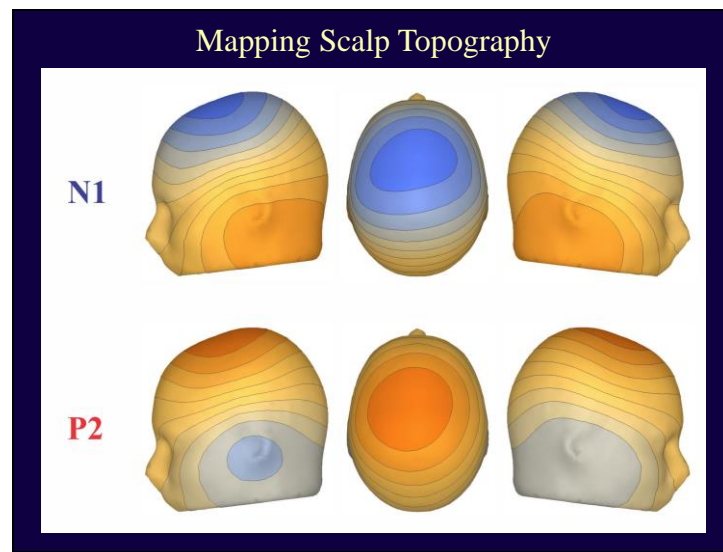
The next few slides will quickly go through a source analysis of the slow auditory evoked potentials. We begin with a 65 channel scalp recording. We note that the waveforms are generally of opposite polarity in the frontal regions compared to more inferior and posterior regions of the scalp.

Slide 32




We can get some idea of the different waves present by superimposing all the recordings in a butterfly plot or by calculating a global field power or RMS measurement. The pattern of the butterfly plot suggests that there may be overlapping components in the region of P2 - the lines cross in the middle of the wings.

Slide 33



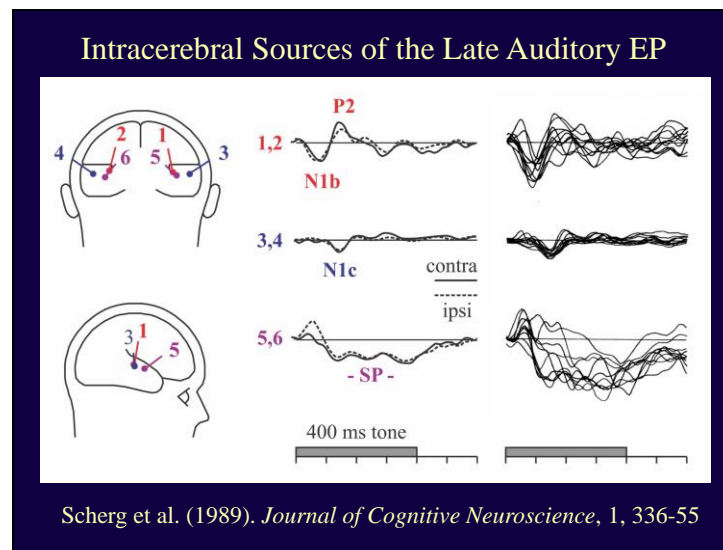
We can look at maps of the scalp topography. The N1 wave is negative frontally and positive inferiorly whereas the opposite is true for the P2 wave.

Slide 34




We can look at individual waveforms and see a small negative wave in the mid temporal regions that occurs between the vertex N1 and the P2.

Slide 35



Source analysis involves placing dipoles in the brain, determining what part of the recorded topography they might explain and calculating how well they fit the data. When the fit is optimized the analysis shows that the N1b and P2 waves result from dipole fields on the surface of the both temporal lobes. The N1c wave reflects activity on the lateral surfaces of the temporal lobes. The sustained potential has a sources on the superior surfaces of the temporal lobes anterior to the N1-P2 sources.

Slide 36



The concepts we have been considering may be summarized by observing the three wizards of evoked potential research. The physiological wizard operates with the anatomy of the dendrites, the unit recordings of the microelectrode, and the neurochemistry of the sliced brain. He is a quiet and hard-working fellow whose scientific endeavor is well-founded upon the other sciences — his cauldron is

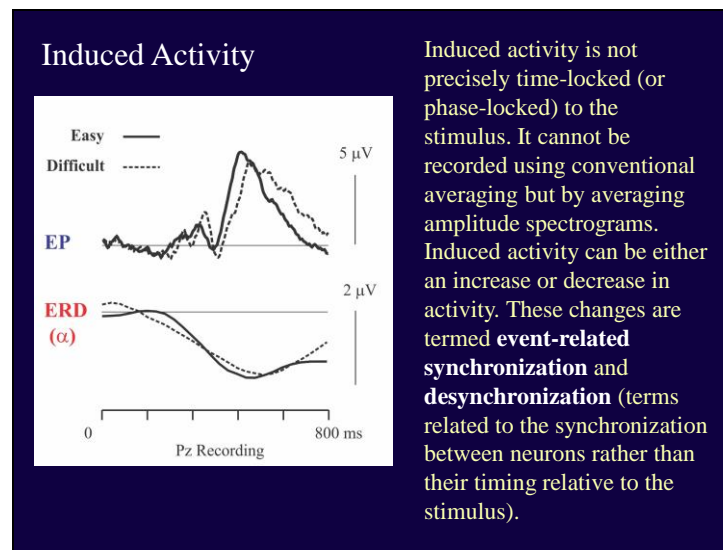
set upon a four-legged stool and buttressed by the stake of mathematics. His activity is "so elegant, so intelligent" but with respect to event-related potentials, the vat is particularly empty. The psychological wizard is more wild-eyed and flamboyant in nature. Reaction times, questionnaires, verbal reports and various stimuli are all entered into "the clatter and the chatter" of his mixing vat. There are only three legs to the stool. Perhaps his scientific endeavor is less soundly based than that of physiology, or perhaps, by being more independent of the other sciences, it wobbles less. It is difficult to tell whether the darkness of the cauldron is related to the inaccessibility of black boxes or to the mystery of black magic. The statistical wizard appears somewhat aloof. He is perhaps intoxicated by the steam of insignificant variance that rises from the boiler where the bones of his data are being picked clean. His eyes are focused upon some far-away matrix, and he is performing one of his most amazing rotations — a new oblique version that tests the power of his machine to its utmost. From out of the statistical still slowly drip the refined components of human psychophysiology. For a moment we can perhaps discern a complex that seems similar to N1-P2-P3, but then, before we can be sure which way is up, it dies away and all is quiet.

Slide 37

Signals and Noise	
I Ways of Looking	
II Time Analyses	
III Frequency Analyses	
IV Component Structure	
V Induced Activity	Induced and Evoked Activity
VI Nature of the Response	Event-Related Spectrogram
VII Overview	Synchronization and Desynchronization

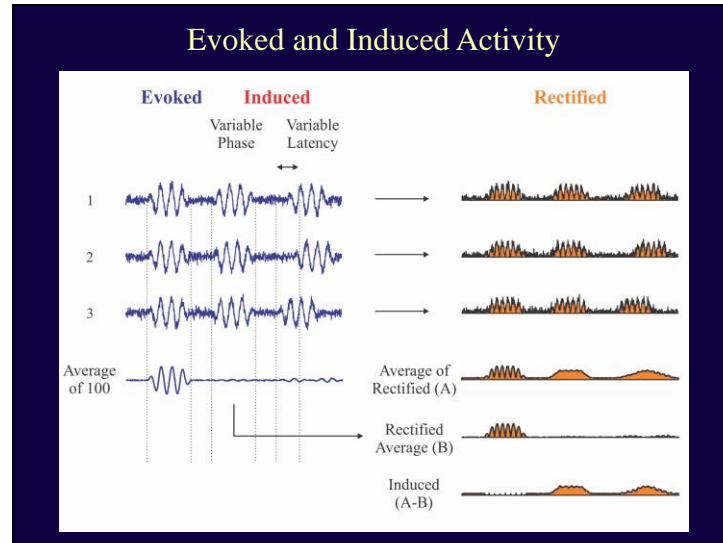
The three wizards give us many different components. The most magical of these are the induced activities. Now you see them but then you don't.

Slide 38



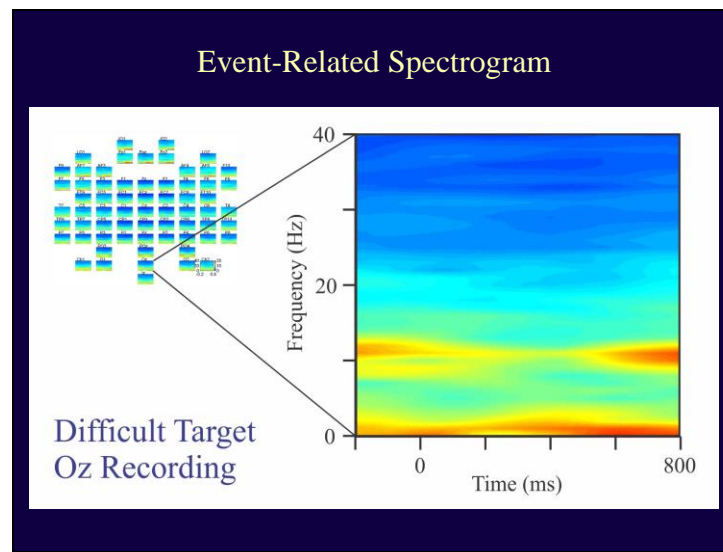
Induced activity ... The figure shows that the alpha activity desynchronizes during the processing a visual stimulus. Indeed we might suggest that the beginning of the even-related desynchronization represents the onset of perceptual processing whereas the P3 represents its end.

Slide 39



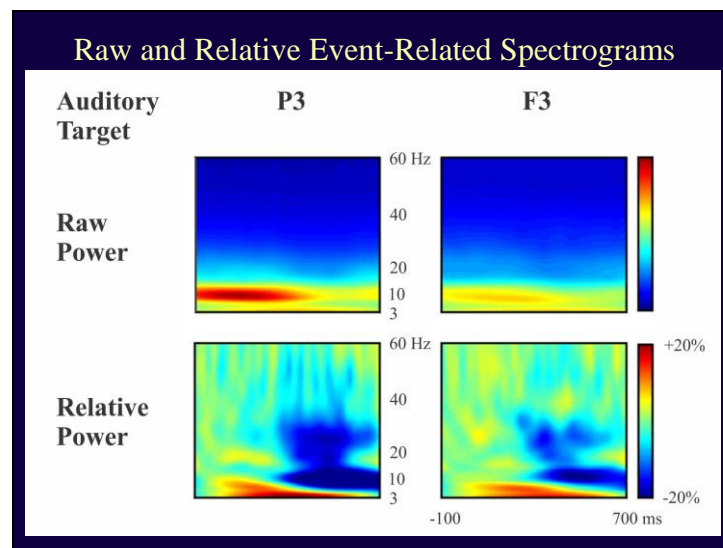
Induced activity varies in its phase or in its timing. With averaging, induced activity cancels itself out and becomes invisible. It may be seen in the single trials in a spectrogram which rectifies the activity so that it does not cancel itself out. The average spectrogram shows both evoked and induced activity. The spectrum of the average shows only the evoked activity. The difference shows the induced activity.

Slide 40




Now we can look at the spectrogram during visual processing. The view is like at looking at a beach with a lagoon, a reef with an opening where the alpha activity desynchronizes, the ocean and the sky. The general tendency to blue is related to the decrease in amplitude of the EEG with increasing frequency.

Slide 41



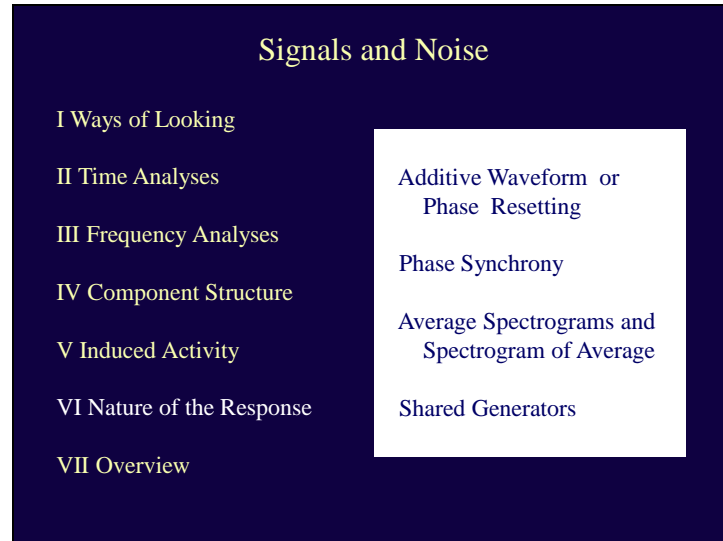
We can compensate for this by normalizing the spectrogram to the activity level at each frequency in the period before the stimulus. Islands of activity then arise from the deep blue sea. Bursts of gamma activity occur and the beta and alpha activity show striking decreases.

Slide 42



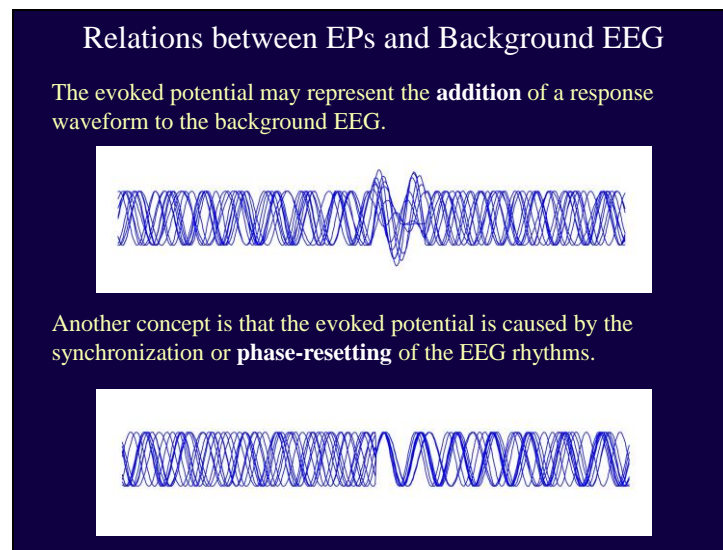
An average spectrogram of the activity associated with processing a visual target shows a burst of gamma activity, a desynchronization in the alpha and beta frequency range and a burst of theta activity as the stimulus is processed.

Slide 43




Induced activity is intriguing. It can show us as much about the processing of stimuli as the evoked activity. The background EEG is not just noise – it is a signal in its own right. Even more intriguing is the relation between the EEG and the evoked potential.

Slide 44




In recent years controversy has arisen about how the evoked potentials are generated. One idea - the one that we have been following, the one that explains the process of averaging - is that the evoked potential is added onto the background EEG noise. Averaging removes the noise and leaves the EP signal. Another idea - one originating with Sayers in the 1970s and promoted by Scott Makeig in the 21st century is that the EP represents the phase-resetting of the background EEG.

Slide 45



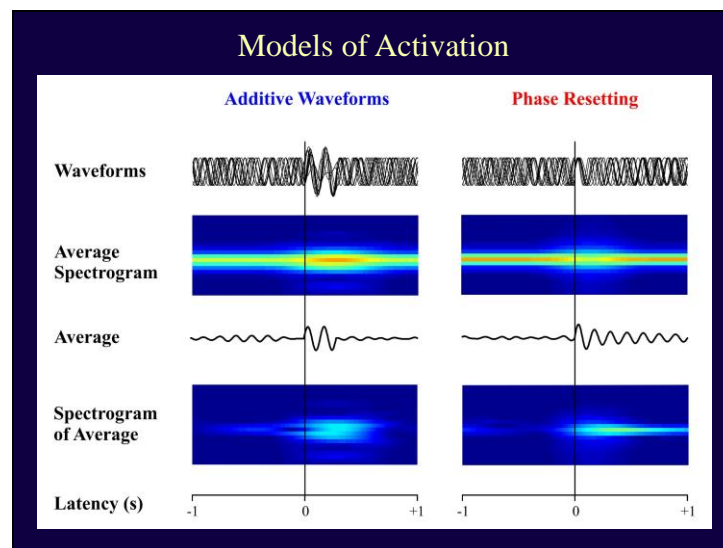
Ali Mazaheri and I looked at these two ideas. First we considered some visual EPs that we had recorded and looked at the phases from one trial to the next. Before the stimulus the phases were random. After the stimulus, the phases became similar from one trial to the next - inter-trial phase coherence. Was our idea that the EP signal was independent of the background EEG noise wrong?

Slide 46




We made a Matlab model. We arbitrarily added a signal to some random EEG signals. The phase showed coherence from trial to trial even though there was no phase-resetting of the modeled EEG. Theoretically we should have been able to discriminate the added activity in the spectrum but this difference did not show up as significant.

Slide 47



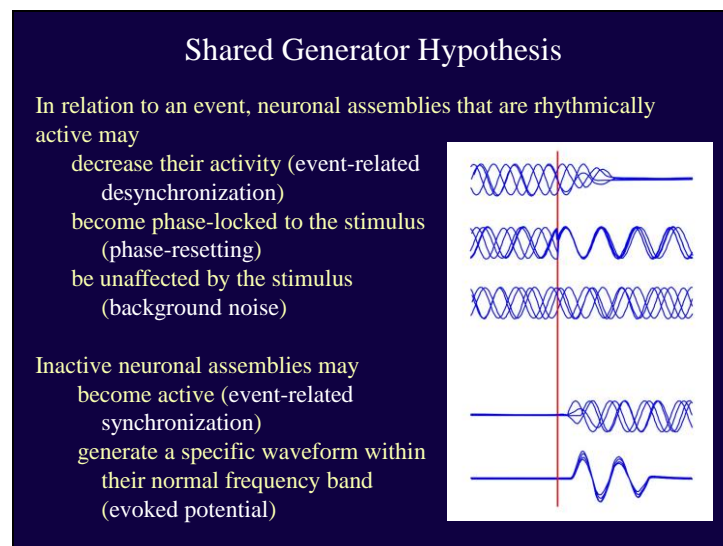
If we add a signal to the ongoing EEG there is an increase in activity in the average spectrogram. If we reset the phases there is no change or a slight decrease in the average spectrogram. Both processes show an average EP.

Slide 48



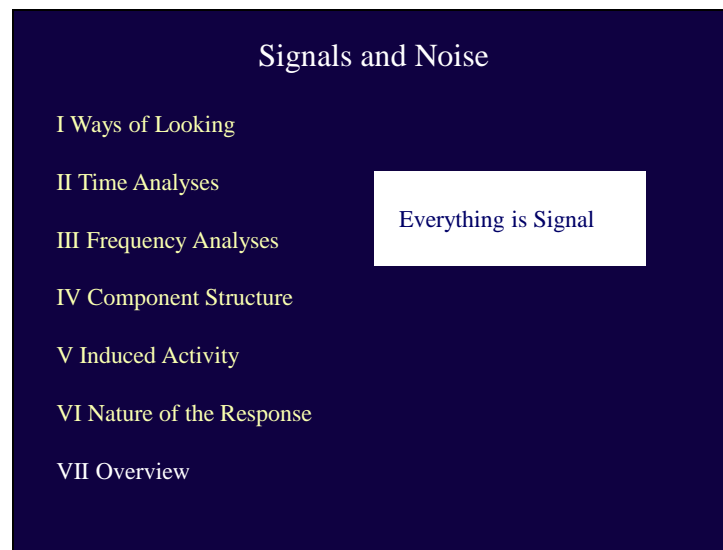
We looked back at the burst of gamma activity that occurs early in the processing of a visual stimulus. The average spectrogram shows little change. The spectrogram of the average EP shows a clear burst of gamma activity. The difference indicates the induced activity – here a desynchronization of gamma activity. The evoked gamma activity ...

Slide 49




The generators are shared in the sense that they are located in the same region of cortex and have similar rhythmic characteristics. The neuronal assemblies actually affected by an event will vary with the event and with the subject experiencing it.

Slide 50



We have considered different ways of looking at the signal and the noise – in particular how they are measured in time and frequency domains. We have looked at the components of the response and met the three wizards in the land of ERP. We have wondered at the magic of induced activity that is visible on single trials and invisible in the average. Finally we looked at the relations between the EP signal and the EEG noise – and found them intertwined. Everything is signal.

Slide 51



It just depends on what you wish to look at and how you choose to look. Kandinsky's painting shows that each floor of a building tells a different story.