Published in final edited form as:

Int J Psychophysiol. 2007 April ; 64(1): 62–74.

Decomposing delta, theta, and alpha time–frequency ERP activity from a visual oddball task using PCA

Edward M. Bernat^{a,*}, Stephen M. Malone^a, William J. Williams^b, Christopher J. Patrick^a, and William G. Iacono^a

^a Department of Psychology, University of Minnesota, Elliott Hall, 75 East River Road, Minneapolis, MN 55455, USA

^b Department of EECS, University of Michigan, Ann Arbor, Michigan, USA

Abstract

Objective—Time–frequency (TF) analysis has become an important tool for assessing electrical and magnetic brain activity from event-related paradigms. In electrical potential data, theta and delta activities have been shown to underlie P300 activity, and alpha has been shown to be inhibited during P300 activity. Measures of delta, theta, and alpha activity are commonly taken from TF surfaces. However, methods for extracting relevant activity do not commonly go beyond taking means of windows on the surface, analogous to measuring activity within a defined P300 window in time-only signal representations. The current objective was to use a data driven method to derive relevant TF components from event-related potential data from a large number of participants in an oddball paradigm.

Methods—A recently developed PCA approach was employed to extract TF components [Bernat, E. M., Williams, W. J., and Gehring, W. J. (2005). Decomposing ERP time-frequency energy using PCA. *Clin Neurophysiol*, 116(6), 1314–1334] from an ERP dataset of 2068 17 year olds (979 males). TF activity was taken from both individual trials and condition averages. Activity including frequencies ranging from 0 to 14 Hz and time ranging from stimulus onset to 1312.5 ms were decomposed.

Results—A coordinated set of time–frequency events was apparent across the decompositions. Similar TF components representing earlier theta followed by delta were extracted from both individual trials and averaged data. Alpha activity, as predicted, was apparent only when time–frequency surfaces were generated from trial level data, and was characterized by a reduction during the P300.

Conclusions—Theta, delta, and alpha activities were extracted with predictable time-courses. Notably, this approach was effective at characterizing data from a single-electrode. Finally, decomposition of TF data generated from individual trials and condition averages produced similar results, but with predictable differences. Specifically, trial level data evidenced more and more varied theta measures, and accounted for less overall variance.

Keywords

Time-frequency; P300; ERP; PCA

* Corresponding author. Tel.: +1 612 624 5063; fax: +1 612 626 2079. E-mail address: ebernat@umn.edu (E.M. Bernat).

1. Introduction

Time–frequency (TF) analysis has become an important tool for assessing electrical and magnetic brain activity from event-related paradigms. The P300 event-related potential (ERP) response (Pritchard, 1981) is perhaps the most widely studied response in work with electrical potential brain data. In TF analyses, theta and delta activities have been shown to underlie P300 ERP activity, whereas alpha is inhibited during the P300 response. Measures of delta, theta, and alpha activities are often evaluated using TF methods. However, methods for extracting relevant activity do not commonly go beyond taking means of windows on the TF surface, analogous to measuring activity within a defined P300 window in time-only signal representations. The current objective was to derive meaningful TF components from event-related potential (ERP) data in an oddball paradigm using a previously developed PCA method as a data driven approach. The large number of participants utilized for the current purposes (2068) allows the development of particularly stable solutions, and thus a detailed characterization of TF activity in a representative, population-based sample.

1.1. Time–frequency and oddball ERP responses

P300 responses contain variation in TF amplitude and topography, forming a processing sequence that varies across time. Topography of P300 responses has been well-investigated, and evidence suggests both earlier anterior and later posterior contributions to P300 responses. It has long been known that P300 responses to target stimuli involve parietal activations that vary with the degree of cognitive processing (e.g. Comerchero and Polich, 1999). Anterior contributions were found initially in response to unexpected and novel stimuli, and later to task irrelevant infrequent stimuli more generally, where stronger orienting responses produced an earlier and more anterior P300 response (Comerchero and Polich, 1999; Friedman et al., 2001). However, current models suggest that deviant stimulus processing generally (i.e. both target and novel stimuli) involves both earlier anterior and later posterior activations, but that the relative activation of each depends on the extent of the orienting response and degree of evaluative cognitive processing (Spencer et al., 2001). A number of investigators have also evaluated ERP activity to novel and target stimuli using time–frequency methods, and findings suggest that the anterior to posterior processing sequence also varies in frequency. Specifically, the anterior activation is generally higher in frequency (e.g. theta) relative to the posterior activation (e.g. delta). Both theta and delta activities have been related to oddball target responses (Basar-Eroglu et al., 1992; Basar-Eroglu and Demiralp, 2001; Basar-Eroglu et al., 2001). Delta, generally measured from zero to around 3 Hz, has been most directly related to target P300 amplitude. Theta activity has also been strongly implicated in oddball processing, and is generally measured from around 3 Hz up to the lower bound of alpha, around 7 Hz. Theta generally precedes delta in the P300 response, and is more anterior in topography, while delta is later and more posterior (Demiralp et al., 2001a,b; Porjesz et al., 2005). Interestingly, theta is selectively enhanced during novelty stimulus presentations, linking it to the orienting processes associated with novelty processing (Demiralp et al., 2001a,b). Thus, the overall processing sequence involves an anterior theta response first, more closely tied to orienting, and then a posterior delta response more closely tied to cognitive processing. It is worth noting that recent work suggests that the anterior response has sources in the anterior cingulate cortex and the posterior response in the temporal–parietal junction (Dien et al., 2004). In the current study, because only target responses in a single parietal electrode will be assessed, no topography differences or selective influence of novelty on these processes can be evaluated. However, contributions of parietally recorded delta and theta activities to target stimulus processing can be evaluated.

Alpha reductions measure increased cognitive processing during the oddball task (Yordanova and Kolev, 1998), and fast and slow alpha have been related to specific aspects of cognitive

processing indexed by P300 (Yordanova et al., 2001). This decrease is often measured as an event-related desynchronization (ERD) of alpha around the P300 (Yordanova et al., 2001). Taken together, delta, theta, and alpha represent a coordinated set of processes that are activated in sequence during oddball processes, and are strongly related to the cognitive events present during target stimulus processing. Higher frequency activity is also modulated during responses in oddball tasks, and is related to perceptual and sensory processes as well as cognitive processes (Karakas et al., 2001). While higher frequency activity is of great interest in understanding the complete set of processes involved in target processing, such frequencies were not available in the current data (a lowpass hardware filter precluding this activity was applied during data collection, see Methods section for details).

1.2. Time–frequency transform methods

Before TF measures can be obtained, a TF representation of the signal must be created (a TF transform). TF transform methods have advanced substantially in the recent past, and two classes are now well-developed: wavelets and Cohen's class. A wavelet is a simple oscillating amplitude function that is localized in time, and in this general form it is referred to as the mother wavelet (Samar et al., 1999). To create TF signal representations, the mother wavelet is systematically stretched and reduced (scaled) in time to be sensitive to lower or higher frequency activity, respectively. There are an infinite number of possible mother wavelets, many of which have been closely studied. The wavelet is passed along the length of the signal multiple times (once for each scale) to filter it for the frequencies consistent with that scale. The resulting collection of filtered time series can then be arranged into rows of a matrix, creating a TF surface where the columns represent different time points, and the rows represent frequency (in scale). Advanced Cohen's class approaches to creating TF signal representations produce similar TF surfaces, but are based on a different approach. In the reduced interference distribution (RID; Williams, 2001), for example, mathematical combinations of spectrograms, which incorporate time activity around the point of interest, are used to estimate 'instantaneous frequency' (the energy in each frequency bin for one time point). This contrasts with wavelets where each iteration in time of a wavelet of a given scale estimates a small 'tile' of energy from a specific range of times and frequencies (defined by the scale of the wavelet). The algorithm for estimating instantaneous frequency is referred to as the kernel, and an infinite number of kernel functions can be implemented within Cohen's class of TF transforms (Williams, 2001). A brief review of wavelet and Cohen's class methods is detailed by Bernat et al. (2005), which also contains references to more detailed reviews and comparisons. It is worth noting that both methods are capable of generating high resolution TF representations of ERP activity, and the PCA decomposition method has been demonstrated to work well with both kinds of transforms (Bernat et al., 2005). It is also worth noting that some advantages offered by the RID relative to wavelets, as all TF surfaces in the current report were created using the RID. Particularly relevant for the current purposes is that the RID creates TF surfaces with uniform resolution across the surface, where the resolutions in wavelet transforms are non-uniform (the concept of scale translates into lower frequency resolution at higher frequencies and lower time resolution at lower frequencies). While wavelets can be tuned to have similar resolution within a particular range of interest for a given decomposition, they must be retuned for different analyses. Additionally, TF energy in the RID, and not wavelets, "satisfies the marginals" (Williams, 2001). This means that summed energy across the columns (time) or rows (frequency) of the TF surface will equal the energy in the signal, offering a more accurate accounting of the energy in the signal than wavelets.

1.3. Time–frequency and PCA

The PCA approach employed here is a general data reduction technique for TF signal representations. Because analysis of TF representations is rapidly developing, there is currently a strong demand for new TF data reduction methods. Current TF data reduction and

measurement methods are often simplistic. For example, because TF representations can be treated as a collection of time series, each filtered to include only certain frequencies, one approach has been to assess each frequency time series separately, using peak and mean analyses common to ERP investigators. Another common approach is to take single tiles (from wavelet TF transforms) and evaluate their activity. However there is also an interest in assessing complete TF surfaces. With a surface, the entirety of the energy in the signal is available at once, which makes for a rich data representation, but complex. The most common method for analyzing high-resolution TF surfaces is to assess a TF region of interest, or to 'cut out' an area of the surface. Like peak and mean measures in time-only signal representations (e.g. P300), this method is reasonably effective when the frequency-and time-course of the activity is understood. However, because TF analytic methods are still being developed, these parameters are often not known, or it can be of interest to evaluate data with completely unknown characteristics. The PCA method employed here was recently developed as an approach to this problem (Bernat et al., 2005), offering a data driven method for decomposing a dataset of TF surfaces. Because this approach is an extension of PCA as commonly applied in the time and frequency domains separately, its operation, as well as many of its advantages and drawbacks, is well-understood.

1.4. Current study

Because TF activity has been shown to be relevant to cognitive processes, there is a strong interest in methods that can effectively extract meaningful TF activity. In the current study a PCA approach to extracting TF components (Bernat et al., 2005) was employed with data from a large dataset of ERP responses in an oddball task. Current approaches to characterizing TF activity from oddball tasks have relied on a priori or post-hoc methods for extracting measures from TF surfaces, and found significant modulations of delta, theta, and alpha during target processing. One goal of the current study was to investigate whether a data driven method will produce measures consistent with previous effects. A second goal was to investigate whether PCA can be an effective method for succinctly characterizing TF activity. Based on previous work, as discussed above, a progression of early higher frequency activity (theta) to later lower frequency activity (delta) is predicted to characterize the time around and during the P300 response, as well as an inhibition of alpha (ERD) during the P300.

2. Methods

2.1. Participants

The sample consisted of twins in the Minnesota Twin Family Study (MTFS), a longitudinal and epidemiological investigation of the origins and development of substance use disorders and related psychopathology. The MTFS employed a population-based ascertainment method in which all same-sex twins born in the state of Minnesota within specified birth years were identified through public birth records. Twins were drawn from the 1972 through 1984 birth cohorts, inclusive. Families were eligible if they lived within a day's drive of the University of Minnesota and if the twins did not have a physical or mental handicap that would preclude completing a day-long assessment. At least 90% of eligible families were located for each cohort, with 17% of eligible families refusing participation. Comparisons of participating and nonparticipating families with respect to demographic, socioeconomic, and mental health measures indicated that participating families were slightly, but significantly, better educated than nonparticipating families, averaging 0.3 more years of education (Iacono et al., 1999). No other differences were significant, indicating that the sample is indeed representative of the population of Minnesota from which it was drawn.

The sample for this investigation consisted of all male and female twin participants for whom ERP data were available from the study's psychophysiological assessment. The sample totaled

2068 adolescents, assessed as close as possible to a target age of 17 ($M = 17.7$; $SD = 0.5$; range = 16.7 to 20.0). This sample combined subjects from the two age cohorts of the MTFs: subjects in one cohort were 17 years old at intake (501 males, 595 females) whereas subjects in the other were approximately 11 years old at intake (478 males, 494 females). Data for the latter came from their second three-year follow-up assessment. Consistent with the demographics of the state of Minnesota at the time study participants were born, nearly all were Caucasian. Parents in the sample spanned all socioeconomic strata. Sixty-two percent had a high school degree or its equivalent, approximately 22% had a college degree, and 7% of fathers and 4% of mothers had advanced degrees.

2.2. Experimental procedure

A visual oddball task was used (Begleiter et al., 1984). Each of the 240 stimuli comprising this task was presented on a computer screen for 98 ms, with the inter-trial interval (ITI) varying randomly between 1 and 2 s. A small dot, upon which subjects were instructed to fixate, appeared in the center of the screen during the ITI. On two-thirds of the trials, participants saw a plain oval to which they were instructed not to respond. On the remaining third of the trials, participants saw a superior view of a stylized head, depicting the nose and one ear. These stylized heads served as “target” stimuli. Participants were instructed to press one of the two response buttons attached to each arm of their chair to indicate whether the ear was on the left side of the head or the right. Half of these target trials consisted of heads with the nose pointed up, such that the left ear would be on the left side of the head as it appeared to the subject (easy discrimination). Half consisted of heads rotated 180° so that the nose pointed down, such that the left ear would appear on the right side of the screen and the right ear would appear on the left side of the screen (hard discrimination).

2.3. ERP data collection and preparation

All participants completed the assessment at approximately the same time in the morning. They sat in a comfortable high-backed chair while electroencephalographic (EEG) data were recorded from three parietal scalp locations, one on the midline (Pz) and one over each hemisphere (P3 and P4), although due to computational constraints, only data from the Pz electrode is reported here. Linked earlobes served as reference and an electrode on the right shin as ground. Blinks and eye movements were recorded with a pair of biopotential electrodes arranged in a transverse montage, one electrode superior to the right eye and the other over the outer canthus. A Grass Model 12A Neurodata acquisition system was used to collect EEG data, with each signal passed through an amplifier with a bandpass of 0.01 to 30 Hz (half-amplitude) and a roll-off of 6 dB per octave. For each trial, 2 s of EEG, including a 500-ms pre-stimulus baseline, were digitized to 12 bits resolution at a rate of 256 Hz. If participants failed to respond on a given trial, or if any EEG signal exceeded the range of the A–D converter, the trial was repeated. Trials repeated more than twice were excluded from further analysis. The procedure of Gratton et al. (1983) was used to correct for blinks and other ocular artifacts in the EEG. Trials with activity exceeding 100 μV between stimulus onset and 1000 ms, relative to a 500 ms baseline, were excluded from further processing. For decompositions based on averaged waveforms, averages were constructed separately for easy and hard discrimination conditions. Comparisons between the easy and hard targets are not the focus of the current analyses and are thus not reported. Also, although data from standard (frequent) stimuli were collected, they were not analyzed for the current report, and thus overall target condition responses served as the basis for all decompositions and analyses were presented. All TF transforms were computed using Cohen’s class RID transform (for further detail see Bernat et al., 2005). TF transforms were created using the maximum available time epochs, from 500 ms before stimulus onset to 1500 ms after stimulus onset, so that edge effects from the transform could be discarded.

2.4. PCA decomposition method

Details of the method employed to apply PCA to time–frequency surfaces have been described previously (Bernat et al., 2005). Here, only the primary aspects of this approach will be reiterated. The application of PCA to time–frequency energy is much the same as its application to signals in the time or frequency domain. The primary difference is that each time–frequency surface is rearranged into a vector, recasting the time–frequency energy into the same terms as amplitude decomposition of time signals, with a matrix of trials in rows and different points of activity in columns (each representing a different time–frequency point). First the covariance matrix is decomposed, then varimax rotation is applied to maximize simple structure, and finally the trials are rearranged into TF surfaces. The number of components to extract was based on breaks in the plot of singular values (the ‘elbow’). Such breaks suggest a qualitative shift such that additional components account for variance at a much lower rate than those before the break. In the current application, decomposition for two breaks were conducted. The first break extracted a fewer number of components to evaluate the outcome at a gross levels of detail, while the second break allowed evaluation of a finer level of detail by extracting a larger number of components. The data handling and decomposition steps were carried out in Matlab (version 6.5, Mathworks, Inc) using a generalized set of scripts developed for this purpose.¹ A graphical depiction of the steps involved in the decomposition process is given in Fig. 1.

2.5. Application of time–frequency PCA decomposition

Three decompositions were undertaken. First, a decomposition of trial level data was conducted using frequencies ranging from 0 to 14 Hz, and time ranging from stimulus onset to 1312.5 ms post-stimulus. This decomposition was intended to cover a broad range of time and a full range of frequencies through alpha. A decomposition of averaged data with these time and frequency ranges was also conducted, but yielded no alpha activity, and is thus not presented. Next, decompositions for both trial and averaged level data were conducted on more narrow time and frequency ranges, to focus on lower frequency delta and theta activities. These included frequencies ranging from 0 to 5.75 Hz (excluding alpha) and time ranging from stimulus onset to 1000 ms post-stimulus (excluding the 1000–1312.5 ms time range as it primarily measured the return of alpha post-P300). Below is an accounting of the rationale for this set of decompositions.

Generating TF surfaces using individual trials or using averages has an important relationship to the nature of the activity that will be represented. Using condition averages will enhance activity that is consistently phase-locked to the stimulus — i.e. has a similar phase from trial to trial, while other activity will be strongly attenuated. For example, this can manifest in ‘component jitter’ where the latency of a peak component measure is inconsistent from trial to trial, creating underestimates of the activity when averaged. At the same time, analysis of averaged data has the most direct parallel with the extant body of research on target P300 responses in oddball tasks, and thus offers an important connection to previous work. Generating TF surfaces from trial level data will contain all activity in the signal, phase-locked and non-phase-locked. While one can specifically assess phase information and differentiate these kinds of activity, the analysis here will focus on analysis of the energy and will not seek to directly differentiate phase activity.

In order to assess the impact of the level of averaging on the PCA decomposition of the ERP data, separate analyses were conducted using TF transforms of trial level and averaged level data. In the trial level approach, PCA decomposition was done on TF surfaces generated from the trials. This approach consumes a great deal more computational resources, and for this

¹Scripts available from the first author.

approach the number of target trials was reduced to 20 per participant (every other trial taken from the middle 40 trials of the trials left after artifact rejection) in order to perform the decomposition across the large number of participants. Data for all decompositions were constrained to the Pz electrode for comparability, due to the computational resources required for the trial level analyses.

Data were decomposed for two different ranges of frequency on the TF surfaces. For trial level data, a decomposition from 0 to 14 Hz was conducted to include alpha range activity. As described earlier, decompositions including alpha activity for averaged level data were attempted, but predictably yielded no alpha components and were not pursued further. A second range of activity was decomposed to exclude alpha, and to focus on delta, theta, and low frequency activity below 1 Hz. These decompositions included frequencies ranging from 0 to 5.75 Hz and time ranging from stimulus onset to 1000 ms post-stimulus, and were conducted with both trial and averaged level data. These decompositions stopped at 5.75 Hz, instead of including activity closer to alpha (e.g. 7 Hz) because power in the alpha range tended to be large and dominate the solution (next to delta) and extended down between 6 and 7 Hz. The range was adjusted down to avoid this problem, and allow better decomposition of theta, delta, and low frequency activity. As described above, two levels of extraction were conducted for each TF range, one with fewer components, and one with a larger number of components, making 6 total decompositions.

2.6. Cross-validation

To evaluate the stability of the components from the decompositions, a cross-validation evaluation was conducted. For each of the 6 solutions evaluated, the original data were split into even and odd subject numbers and decomposed separately, with the same number of components as the originals. To evaluate the cross-validation decompositions, each extracted component was matched between odd and even decompositions. Stability of any given component was inferred if it was extracted from both cross-validation sets independently.

3. Results

3.1. Decomposition of full frequency range (0–14 Hz)

Decomposition of trial level data including the full range of assessed frequencies are presented in Fig. 2. Two decompositions were undertaken based on the singular values (presented as a scree plot in Fig. 2), one with fewer components (6, accounting for 34% of the variance) and one with a greater number of components (14, accounting for 50% of the variance). For the six component solution, three lower frequency and three alpha frequency components were apparent. Alpha activity was characterized by a fast and slow components before P300 and one alpha component after P300. These are representative of the alpha ERD. The three lower frequency components followed a sequence from earlier higher frequency activity (spanning 1–3 Hz, corresponding to the leading edge of P300), to a component close to the peak of P300 (spanning 1–2 Hz), and finally to the later longer duration low frequency component (primarily contained in 1 Hz, corresponding to a slow wave).

The 14 component solution offered much greater separation of TF activity spanning the frequency range, and suggested a progressive sequence of activity including alpha, theta, delta, and low frequency (1 Hz and below). As apparent in Fig. 2, the earliest components are fast and slow alpha, indexing alpha continuing from the baseline. As these alpha components subside, starting around 150 ms and before P300, three types of TF activity become apparent. First, a single alpha component takes the place of fast and slow alpha components, indexing an alpha modulation antecedent to the alpha ERD. Second, two similar theta components are apparent, whose primary difference is one latency bin. The highly similar nature of these two

theta components, and that they appear to be simply shifted in time, suggests that together they characterize latency shifts of activity in this timeframe by weighting one or the other more strongly. The third type of activity in this pre-P300 window is a low frequency activation (centered from 1 to 2 Hz). Next starts the P300 sequence apparent in the 6 component solution, but in greater detail. This starts with a higher frequency component on the leading edge of P300 centered at 3 Hz (the lower edge of what is generally considered theta). Next is a component corresponding in time to the peak of P300 (centered at 2 Hz) which transitions into a component (centered across 1 and 2 Hz) on the trailing edge of P300. Post-P300 slow wave activity is apparent in two components centered at 1 Hz. Finally, alpha power returns as the P300 complex subsides in a mirror image of its pre-P300 activity: a single component appears first, which transitions to separate fast and slow alpha components.

3.2. Decomposition of frequency activity below alpha

To investigate the theta–delta processing sequence in more detail, a second set of decompositions was conducted excluding the alpha frequencies. These were conducted for trial level data (Fig. 3), producing a five component (42% of the variance) and a 16 component solution (70% of the variance), as well as with condition averages, producing a five component (79% of the variance) and an 11 component solution (91% of the variance). Five component solutions for both trial and averaged level data produced the same result, and is apparent in Figs. 3 and 4. In order of time, first is a low frequency component, centered between 150 and 200 ms and 1 Hz. The next four components represent the P300 sequence detailed earlier, from a component on the leading edge of P300 on the border between theta and delta, through P300 peak and trailing edge delta components, and finally a slow wave component centered between .5 and 1 Hz.

The higher component solutions for the averaged data (11 component) and trial data (16 component) replicated the components from the 5 component solution, but with notable additions. First, a very low frequency component is apparent in both the averaged and trial decompositions, centered at .5 Hz and covering a broad time range from around 200 to 700 ms. The P300 sequence already described is again apparent, but the lowest frequencies of each component in the sequence now belong to this low frequency component. Also notable in both averaged and trial decompositions is theta activity occurring during the P300 peak. While the earlier theta (150–200 ms) is similar for the full range decomposition and these excluding the alpha frequencies, only these decompositions extract theta during this time period. Some differences between the averaged and trial level decompositions are notable. First is the greater variance accounted for by the decomposition of averages relative to the decomposition of trials. While predictable, this difference is striking. Also, for the trial level data, the two theta components apparent during the P300 in the averaged decomposition are split into five components, and a late theta component is present which is not for the averaged data. Delta activity is also characterized by one additional component at the end of the time window in the trial level decomposition, accounting for the 5th additional component. It seems likely that these components are possible due to variation at the trial level that is attenuated for during averaging. Overall, similarity in the structure of the extracted components is more striking than the differences, although the additional components for the trial level decomposition likely represent meaningful inter-trial variation not available in the averages.

3.3. Cross-validation

To evaluate the stability of the decompositions, a visual cross-validation evaluation was conducted, as described earlier. For each of the six presented solutions, the data were split into even and odd subject numbers and recomputed separately, with the same number of components as the originals. The results of these decompositions are presented in Fig. 5. Notably, all components, across all decompositions, matched between the cross-validation sets.

This matching is visually apparent in the plots. This suggests that the components extracted represent stable activations.

4. Discussion

Using the principal components analysis approach, delta, theta, alpha, and low frequency activity (1 Hz and below) were extracted from the ERP data, which yielded a detailed picture of activity to target stimuli from this oddball paradigm. A predictable progression of activity, over time and across frequencies, was readily apparent. The following discussion details the progression of activity, suggests areas of further analysis, and evaluates some of the benefits and limitations of the current analytic approach.

Current results are consistent with the previous research in showing that ERP responses to target stimuli in oddball tasks involve a progression from higher frequency to lower frequency activity across time (Basar-Eroglu et al., 1992; Basar-Eroglu and Demiralp, 2001; Basar-Eroglu et al., 2001; Demiralp et al., 2001a,b). However, the current results offer a more detailed characterization of the nature of the activity than previous work. An interesting question about the extracted components is their topographical distribution. As described earlier, generally, higher frequency theta activity is stronger frontally while the lower frequency delta is stronger parietally. Because data in the current study was only reported from the parietal Pz electrode, it is likely that the ratio between theta and delta is skewed towards delta in the decompositions presented here. Thus, the extent of theta activity present during the task may not be fully represented in the current decompositions. At the same time, it will be of interest to apply the PCA decomposition to multi-channel data from the same paradigm to assess whether the observed progression from earlier higher frequency to later lower frequency activity is still apparent, and additionally whether it shows the anterior to posterior progression. Insofar as it does, the PCA method represents an excellent approach to characterizing the nature of this processing sequence in and time and frequency.

The theta to delta progression has already been shown to exhibit reduced activity in substance abuse psychopathology relative to controls (Porjesz et al., 2005). Similarly, both alcoholics (Kamarajan et al., 2004), and children of alcoholics (Kamarajan et al., 2005), have shown reduced theta and delta activities for nogo stimuli during go-nogo tasks. This relationship is particularly relevant to the primary aims of the MTFs (cf. Iacono et al., 1999), from which the current data was drawn. Specifically, because theta has been related to more anterior processes, and delta to more posterior processes, the current decomposition method may offer possibilities for evaluating whether certain psychopathologies disproportionately affect these processes in much greater detail than previous approaches. This would be particularly true with data including anterior leads, but would be suggestive even with the current data from only the Pz electrode.

With regard to the comparison of average and trial level decompositions for activity below alpha, there was striking similarity in structure for theta, delta, and low frequency activity. Due to the nature of PCA, the decomposition method is most sensitive to activity that occurs in the same bin (or collection of bins) for each row of the decomposition matrix. Thus, activity that occurred at the same TF bins from trial to trial contributed most strongly to extracted components across the trial and averaged decompositions. Additionally, the trial level analyses accounted for less overall variance than the averaged level decomposition. Because solutions were similar, and more variance was accounted for by the averaged decomposition, it seems likely that the two levels of analysis extracted the same activity, but simply with a higher signal to noise ratio in the case of the average. This is in keeping with effects generally found for averaging signals in the time domain for data from oddball tasks: phase-locked activity that drives ERP components is enhanced while other activity is attenuated. Trial level

decompositions are still relevant, however. First, there can be situations where the trial level data is required to measure a given effect, such as alpha only being apparent at the trial level in the current report. Also, single trial variation in P300 has been an important goal for many researchers, and TF approaches have been helpful in creating such characterizations (e.g. Demiralp et al., 1999). Single-trial information is available from the current decompositions, but further analyses would be helpful in clarifying the utility of the PCA approach for characterizing single trial activity. For example, the amplitude and latency of the extracted components could be compared to trial level behavioral data such as reaction time and accuracy. This kind of analysis offers the possibility of linking physiological processes with their immediate behavior output, allowing a more detailed measure of the neural processes that generate behavioral responses from moment to moment.

Low frequency activity (e.g. 1 Hz and below) was apparent in the decompositions, and was separated from delta activity generally measured in relation to P300. This suggests that a very slow wave process may be a separable contributor to ERPs from target detection tasks. Beyond appearing in the decompositions as separate components, additional analyses would be relevant to discern whether this activity is separate and meaningful. Stronger inferences could be made if, for example, low-frequency activity carried information such as condition or individual differences, or differentiable neural sources could be identified for the activity. The current findings offer a method for extracting this kind of activity, and suggests that further research on the nature of this activity would be of interest.

With regard to measured alpha activity, it was effectively extracted from the trial level data. Reduction of parietal alpha activity that is not phase-locked to the stimulus has been demonstrated to represent increased cognitive activation, and has been differentiated from anterior phase-locked alpha activity (Yordanova and Kolev, 1998). Because the observed activity was evident primarily in the trial level data in the recorded parietal Pz electrode, the components are most likely representative of non-phase-locked activity. The ERD of alpha during P300 was appropriately modeled in the decomposition, where alpha activity was greatly attenuated during the P300. This is consistent with the notion of alpha suppression during the P300. Interestingly, fast and slow alpha emerged as components at the earliest and latest points in the signal, but only a single alpha component was apparent just before and after P300. Fast and slow alphas have been shown to have significant relationships to the cognitive processes related to P300 (Yordanova et al., 2001), and the current decomposition offers an appropriate method for extracting measures of this activity.

4.1. Summary and future directions

A stable decomposition of the ERP activity at the Pz electrode in this visual oddball task was derived in the present analyses. These decompositions produced predictable measures of the TF activity, but in greater detail than the other TF data reduction methods. The current results could be extended in an interesting manner by focusing on task conditions to further elucidate the nature of the TF activity in relation to specific cognitive processes. For example, comparing decompositions between target and non-target conditions would offer the possibility of inferring TF activity that is common across tasks, activity that is simply attenuated or enhanced by target processing, and activity that is qualitatively different for target processing. Qualitative differences would suggest additional neural generators or cognitive processes recruited only for target processing. Additionally, evaluation of the easy and hard target conditions would offer similar information about the neural and cognitive processes involved in task difficulty. Finally, given the large epidemiological nature of the current set of data, a thorough investigation of gender differences could be of general interest.

Some benefits of the current approach are worth noting specifically. First, the PCA approach is data driven. Thus, using this approach, a priori definitions of the nature of the time–frequency

activity are not necessary. Next, the logic of PCA for TF is fundamentally the same as PCA for the time or frequency domains alone. For this reason, results derived using this approach are easily compared with work applying PCA to ERPs. Also, many core benefits and limitations of applying PCA to physiological signals are well understood (for a more detailed review see e.g. Chapman and McCrary, 1995). A benefit demonstrated in the current analyses is that this approach can offer complex representations of data from single channel recordings. Future work may focus on multi-channel data to inform the topographical distribution of the extracted TF activity, or to make inferences about neural generators using source localization. Finally, it will be of great interest to apply the decomposition to other frequency ranges such as beta and gamma, to aid in making inferences about the structure of high-frequency activity.

Acknowledgements

This study was funded by grant DA05147 from the National Institute on Drug Abuse and AA09367 from the National Institute on Alcohol Abuse and Alcoholism.

References

- Basar-Eroglu C, Demiralp T. Event-related theta oscillations: an integrative and comparative approach in the human and animal brain. *Int J Psychophysiol* 2001;39 (2–3):167–195. [PubMed: 11163896]
- Basar-Eroglu C, Basar E, Demiralp T, Schurmann M. P300-response: possible psychophysiological correlates in delta and theta frequency channels. A review *Int J Psychophysiol* 1992;13 (2):161–179.
- Basar-Eroglu C, Demiralp T, Schurmann M, Basar E. Topological distribution of oddball ‘P300’ responses. *Int J Psychophysiol* 2001;39 (2–3):213–220. [PubMed: 11163898]
- Begleiter H, Porjesz B, Bihari B, Kissin B. Event-related brain potentials in boys at risk for alcoholism. *Science* 1984;225 (4669):1493–1496. [PubMed: 6474187]
- Bernat EM, Williams WJ, Gehring WJ. Decomposing ERP time–frequency energy using PCA. *Clin Neurophysiol* 2005;116 (6):1314–1334. [PubMed: 15978494]
- Chapman RM, McCrary JW. EP component identification and measurement by principal components analysis. *Brain Cogn* 1995;27 (3):288–310. [PubMed: 7626278]
- Comerchero MD, Polich J. P3a and P3b from typical auditory and visual stimuli. *Clin Neurophysiol* 1999;110 (1):24–30. [PubMed: 10348317]
- Demiralp T, Ademoglu A, Schurmann M, Basar-Eroglu C, Basar E. Detection of P300 waves in single trials by the wavelet transform (WT). *Brain Lang* 1999;66 (1):108–128. [PubMed: 10080867]
- Demiralp T, Ademoglu A, Comerchero M, Polich J. Wavelet analysis of P3a and P3b. *Brain Topogr* 2001a;13 (4):251–267. [PubMed: 11545154]
- Demiralp T, Ademoglu A, Istefanopulos Y, Basar-Eroglu C, Basar E. Wavelet analysis of oddball P300. *Int J Psychophysiol* 2001b;39 (2–3):221–227. [PubMed: 11163899]
- Dien J, Spencer KM, Donchin E. Parsing the late positive complex: mental chronometry and the ERP components that inhabit the neighborhood of the P300. *Psychophysiology* 2004;41 (5):665–678. [PubMed: 15318873]
- Friedman D, Cycowicz YM, Gaeta H. The novelty P3: an event-related brain potential (ERP) sign of the brain’s evaluation of novelty. *Neurosci Biobehav Rev* 2001;25 (4):355–373. [PubMed: 11445140]
- Gratton G, Coles MG, Donchin E. A new method for off-line removal of ocular artifact. *Electroencephalogr Clin Neurophysiol* 1983;55 (4):468–484. [PubMed: 6187540]
- Iacono WG, Carlson SR, Taylor J, Elkins IJ, McGue M. Behavioral disinhibition and the development of substance-use disorders: findings from the Minnesota Twin Family Study. *Dev Psychopathol* 1999;11 (4):869–900. [PubMed: 10624730]
- Kamarajan C, Porjesz B, Jones KA, Choi K, Chorlian DB, Padmanabhapillai A, et al. The role of brain oscillations as functional correlates of cognitive systems: a study of frontal inhibitory control in alcoholism. *Int J Psychophysiol* 2004;51 (2):155–180. [PubMed: 14693365]
- Kamarajan C, Porjesz B, Jones K, Chorlian D, Padmanabhapillai A, Rangaswamy M, et al. Event-Related Oscillations in Offspring of Alcoholics: Neurocognitive Disinhibition as a Risk for Alcoholism. *Biol Psychiatry*. 2005

- Karakas S, Basar-Eroglu C, Ozesmi C, Kafadar H, Erzen OU. Gamma response of the brain: a multifunctional oscillation that represents bottom-up with top-down processing. *Int J Psychophysiol* 2001;39 (2–3):137–150. [PubMed: 11163893]
- Porjesz B, Rangaswamy M, Kamarajan C, Jones KA, Padmanabhapillai A, Begleiter H. The utility of neurophysiological markers in the study of alcoholism. *Clin Neurophysiol* 2005;116 (5):993–1018. [PubMed: 15826840]
- Pritchard WS. Psychophysiology of P300. *Psychol Bull* 1981;89 (3):506–540. [PubMed: 7255627]
- Samar VJ, Bopardikar A, Rao R, Swartz K. Wavelet analysis of neuroelectric waveforms: a conceptual tutorial. *Brain Lang* 1999;66 (1):7–60. [PubMed: 10080864]
- Spencer KM, Dien J, Donchin E. Spatiotemporal analysis of the late ERP responses to deviant stimuli. *Psychophysiology* 2001;38 (2):343–358. [PubMed: 11347879]
- Williams, WJ. Reduced interference time–frequency distributions: scaled decompositions and interpretations. In: Debnath, L., editor. *Wavelet Transforms and Time–Frequency Signal Analysis*. Birkhauser; 2001. p. 381–417.
- Yordanova J, Kolev V. Event-related alpha oscillations are functionally associated with P300 during information processing. *Neuroreport* 1998;9 (14):3159–3164. [PubMed: 9831444]
- Yordanova J, Kolev V, Polich J. P300 and alpha event-related desynchronization (ERD). *Psychophysiology* 2001;38 (1):143–152. [PubMed: 11321615]

Illustration of decomposition process.

Step 1 - Transform time signals to time-freq. surfaces

Step 2 - Decompose time-freq. surfaces

Fig. 1.

Illustration of the decomposition process. In step 1, a dataset of waveforms is transformed into a dataset of time–freq surfaces, one for each waveform. In step 2, a principal components analysis is conducted. First the number of components is chosen from the singular values plot. The components are then extracted (and varimax rotated) and displayed graphically. Finally, descriptive labels are assigned.

Fig. 2. Decomposition of trial level activity across the full range of assessed frequencies (0–14 Hz). Grand average time and time–freq. plots are presented at the top (see Fig. 1 for surface and color keys). Two principal components analysis decompositions are presented below the grand averages. Scree plot contains singular values (units not relevant) for the largest 30 components, depicting the relative variance accounted for by each component. Within the scree plot, blue indicates the components extracted for decomposition of a lower number of components (6), and red the additional components extracted for decomposition of a higher number of components (14). Numbers in blue or red next to each component detail the order of variance explained for the unrotated solution, and thus correspond to the order in the scree plot. Numbers

in black correspond to the order of variance explained for each component in the varimax rotated solution — i.e. the displayed component surfaces. Components are sorted by time.

Fig. 3.

Decomposition of averaged activity below the alpha range (0–5.75 Hz). Grand average time and time–freq. plots are presented at the top (see Fig. 1 for surface and color keys). Two principal components analysis decompositions are presented below the grand averages. Scree plot contains singular values (units not relevant) for the largest 30 components, depicting the relative variance accounted for by each component. Within the scree plot, blue indicates the components extracted for decomposition of a lower number of components (5), and red the additional components extracted for decomposition of a higher number of components (11). Numbers in blue or red next to each component detail the order of variance explained for the unrotated solution, and thus correspond to the order in the scree plot. Numbers in black

correspond to the order of variance explained for each component in the varimax rotated solution — i.e. the displayed component surfaces. Components are sorted by time.

Fig. 4.

Decomposition of averaged activity below the alpha range (0–5.75 Hz). Grand average time and time–freq. plots are presented at the top (see Fig. 1 for surface and color keys). Two principal components analysis decompositions are presented below the grand averages. Scree plot contains singular values (units not relevant) for the largest 30 components, depicting the relative variance accounted for by each component. Within the scree plot, blue indicates the components extracted for decomposition of a lower number of components (5), and red the additional components extracted for decomposition of a higher number of components (16). Numbers in blue or red next to each component detail the order of variance explained for the unrotated solution, and thus correspond to the order in the scree plot. Numbers in black

correspond to the order of variance explained for each component in the varimax rotated solution — i.e. the displayed component surfaces. Components are sorted by time.