

The VESPA: a Method for the Rapid Estimation of a Visual Evoked Potential

Edmund C. Lalor^{*†} Barak A. Pearlmutter[‡] Richard B. Reilly^{*†}
Gary McDarby[§] John J. Foxe^{†¶}

March 17, 2006

(CVS: p.tex 1.143)

Abstract

Faster and less obtrusive means for measuring a Visual Evoked Potential would be valuable in clinical testing and basic neuroscience research. This study presents a method for accomplishing this by smoothly modulating the luminance of a visual stimulus with a stochastic process. Despite its visually unobtrusive nature, the rich statistical structure of the stimulus enables rapid estimation of the visual system's impulse response. The profile of these responses, which we call VESPAs, correlates with standard VEPs, with $r = 0.91$, $p < 10^{-28}$ for the group average. The time taken to obtain a VESPA with a given signal-to-noise ratio compares favorably to that required to obtain a VEP with a similar level of certainty. Additionally, we show that VESPA responses to two independent stimuli can be obtained simultaneously, which could drastically reduce the time required to collect responses to multiple stimuli. The new method appears to provide a useful alternative to standard VEP methods, and to have potential application both in clinical practice and to the study of sensory and perceptual functions.

Keywords: EEG, Visual Evoked Potential, Experimental Design, System Identification, Magnocellular, Striate Cortex

1 Introduction

Since its earliest descriptions (Cobb and Dawson, 1960; Vaughan Jr. and Hull, 1965), the visual evoked potential (VEP) has become a routinely used and extremely valuable tool in both research and clinical settings for the evaluation of visual sensory and perceptual processing. In clinical testing, the so-called transient VEP is typically evoked by the repeated presentation of a visual stimulus at a rate of less than or equal to two presentations per second and extracted from the EEG using signal averaging techniques. When recorded to such repetitive stimulation, it shows several distinct components (*e.g.* C1, P1, N1) with distinctive scalp topographies over the occipital scalp. It is acknowledged that the greatest advantage of the VEP technique is its exquisite temporal resolution which is limited only by the sampling rate of the measurement device.

As well as being used to evaluate optic neuritis and tumors (Kupersmith et al., 1981), retinal disorders (Alexander et al., 2005) and demyelinating diseases such as multiple sclerosis (Halliday et al., 1972; Matthews et al., 1977), more recent work has also shown that certain components of the transient VEP are affected in disorders such as schizophrenia (Foxe et al., 2005), autism (Kemner et al., 1994) and depression (Fotiou et al., 2003).

^{*}School of Electrical, Electronic and Mechanical Engineering, University College Dublin, Belfield, Dublin 4, Ireland

[†]Cognitive Neurophysiology Laboratory, St Vincent's Hospital, Fairview, Dublin, Ireland

[‡]Hamilton Institute, National University of Ireland Maynooth, Co. Kildare, Ireland, barak@cs.nuim.ie, phone: +353 1 708 6100. **Corresponding author**

[§]University College Dublin, Belfield, Dublin 4, Ireland

[¶]Program in Cognitive Neuroscience, Department of Psychology, The City College of the City University of New York, New York, USA

If the rate of repetitive presentation of the visual stimulus exceeds 4–8 Hz, the separate components of the transient VEP are no longer seen due to refractoriness, and a periodic frequency-following response known as the steady-state visual evoked potential (SSVEP) is elicited (Regan, 1989). The periodicity of this response matches that of the stimulus, and provided stimulus presentation is precise, SSVEP power extends over an extremely narrow bandwidth. Spectral analysis with high frequency resolution allows rapid and continuous quantification of the SSVEP magnitude with a high signal-to-noise ratio (SNR). The advantage of the ability to rapidly obtain the SSVEP comes at the cost of the intrinsic timing information that comes with the distinct peaks of the transient VEP.

Both methods have also been used to study attention mechanisms in the brain (Yamaguchi et al., 1995; Shibata et al., 1999; Müller et al., 2000, 2003; Gruber et al., 1999). It has been reported that SSVEP amplitude modulations correlate strongly with certain peaks of the transient VEP but not with others (Müller and Hilliard, 2000), suggesting that SSVEP studies which measure only response amplitude changes as a function of attention will exhibit degraded performance as compared to techniques which also monitor response latencies or the complete response profile. While studies of attentional modulation of transient VEPs abound, this technique is hampered by the need to average over many trials to derive a stable response profile with a typical minimum number in the region of 60 trials and as many as 200–300 being preferable. As mentioned above, each of these trials needs to be separated by at least 500 ms in order to obtain a sufficiently high SNR. The length of time required to acquire this number of trials and the discrete nature of the trials present a serious challenge to the continuous measurement of short term focusing and shifting of attention. That is, many of the experimental paradigms necessitated by this arrangement become decidedly monotonous and can be extremely taxing for subjects, making them noticeably non-environmental in nature. It would be of great use to have a method for rapidly and continuously measuring the visual evoked response

where a complete temporal profile could be obtained without the necessity of such cumbersome paradigms.

White noise signals are commonly used in both linear and nonlinear system identification of physiological systems (Marmarelis and Marmarelis, 1978; Coppola, 1979). By considering the brain in simplified form as a linear system, with isolated events as input and EEG as output, the average event related potentials (ERPs) can be said to approximate the system's time-domain impulse response functions, which is the linear part of the event related dynamics. However, in reality, events are not isolated. Rather, inputs occur in a rapid and continuous stream and their associated electrophysiological responses often overlap in time. Consequently, an average ERP may obscure the brain's response dynamics.

In this paper we describe a method which facilitates the rapid acquisition of a visual evoked potential with a complete temporal profile and high SNR. This is accomplished by smoothly modulating the luminance of a visual stimulus using underlying waveforms to rapidly estimate the time-domain impulse response, which we have termed the VESPA (for Visually-Evoked Spread Spectrum Response Potential). These underlying waveforms have the property that their power is spread over a range of frequencies and as a result are termed spread spectrum waveforms. Spread spectrum communications is a technique in which a signal is transmitted on a bandwidth considerably larger than the frequency content of the original information (Markey and Antheil, 1942). We compare the profile of the VESPA with that of the VEP elicited using standard methods. We also compare, across a range of SNRs, the impulse response acquisition time using spread spectrum methods versus standard methods. Several results are provided which demonstrate the potential of this method, and several applications and extensions are proposed.

2 Methods

2.1 Subjects

Twelve subjects (three female) aged between 21 and 41 participated in the study. All had normal or corrected-to-normal vision. All subjects provided written-informed consent once the goals of the experiment were explained to them. All procedures were approved by the Ethics Committee of St. Vincent's Hospital, Fairview, Dublin. Subjects were paid a modest fee for their participation.

2.2 Hardware

Subjects were seated 60 cm from a 19 inch computer monitor driven by an NVIDIA GeForce FX5200 video card, at a refresh rate of 60 Hz. EEG data were recorded from 64 electrode positions, filtered over the range 0–134 Hz and digitized at a rate of 512 Hz using the BioSemi Active Two system (<http://www.biosemi.com/faq/cms&drl.htm>). Synchronization between the video display and the EEG signals was ensured by including the signal on the parallel port of the presentation computer, controlled by the presentation software, among the signals acquired by the analog-to-digital converter bank.

The response properties of the video monitor used for stimulus presentation were measured using a Nuclear Associates photometer, model 07-621, with an ambient light shield. The monitor was found to have a gamma of 2.3.

2.3 Images

Two basic images were used in this study. The first was a snowflake image (Fig. 1(a).) This contains a large uniform area while also having numerous sharp edges that, it was hoped, would increase the activation of V1. Striate cortex holds columns of neurons that become active when lines or edges are presented, with each column responding to a specific orientation (Hubel and Wiesel, 1959). The snowflake image subtended visual angles of 5.25° vertically and horizontally.

The second was a standard checkerboard pattern, as seen in Fig. 1(c). This consists of equal numbers of black and white checks. Each check

subtended a visual angle of 0.65° both horizontally and vertically, while the checkerboard as a whole subtended visual angles of 5.25° vertically and horizontally. In the case of both the snowflake images and the checkerboard patterns, the luminances of the black and white areas were measured as 0.1 cd/m^2 and 164 cd/m^2 respectively, giving a Michelson contrast of 99.9%.

2.4 Standard Pattern Reversal

To allow direct comparison between standard methods and the spread spectrum method employed in this study, the standard method of pattern reversal was used. In the case of the snowflake image, the pattern reversal method was implemented using the images shown in Fig. 1(a). These images consisted of equal numbers of black and white pixels ensuring there was no change in the mean luminance level during the course of the tests. The presentation of these images was alternated every 1000 ms.

The pattern reversal method was also implemented using the original checkerboard and a phase reversed checkerboard image again ensuring that there was no change in mean luminance level with a change in checkerboard phase. Again, the presentation of these images was alternated every 1000 ms.

2.5 Spread Spectrum Stimuli

For all of the spread spectrum stimuli the image being displayed is controlled directly by an underlying spread spectrum modulation waveform. Using the Nyquist sampling theorem and given that EEG power above 30 Hz is very low, the monitor refresh rate was set to 60 Hz.

Waveforms with any desired statistical properties can be pre-computed and stored. This is accomplished by first choosing a target power spectrum, then shaping Gaussian noise appropriately. As waveforms are pre-computed the use of a causal filter is not required. For this reason, the shaping filter is a simple linear zero-phase filter, performed by simply scaling the coefficients of the noise in the Fourier domain and then converting back to the time domain.

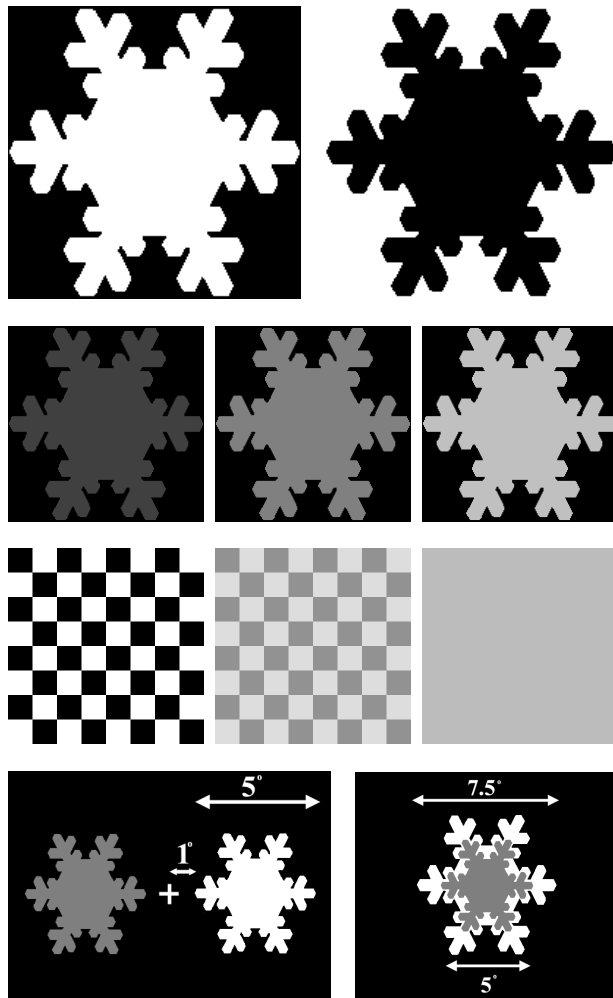


Figure 1: Stimuli used. **(a)** The snowflake and inverted snowflake images used for the pattern reversal method. **(b)** Snowflake images with gray-scale levels of 64, 128 and 192, respectively. **(c)** Constant Mean luminance checkerboards 0, 34 and 67. **(d)** Examples of the two-snowflake stimuli, shown with one snowflake at gray-scale level 255 and the other at 127.

In this study, unless otherwise stated, normally distributed noise waveforms with uniform power over the range 0–30 Hz were used.

2.5.1 Snowflake

Using the white snowflake image of Fig. 1(a) as a template, 256 snowflake images were generated where the white area of each image was assigned a gray-scale value between 0 and 255. Fig. 1(b) shows three such images. The underlying spread spectrum waveform was mapped to the luminance level according to a linear relation, with the zero-point of the waveform corresponding to a luminance of 50%, and scaled to allow \pm three standard deviations within the displayable dynamic range. On every refresh of the computer monitor, the snowflake image corresponding to the current sample of the input waveform was displayed and the EEG data was tagged with the corresponding value of the luminance.

2.5.2 Constant Mean Luminance Checkerboards

68 checkerboards were generated where the mean of the luminance of the lighter checks and the darker checks were approximately equal for each checkerboard. For example the checkerboard consisting of dark checks of gray-scale level 0 and light checks of gray-scale level 255 has a mean luminance of approximately 82 cd/m^2 . Similarly the checkerboard consisting of dark checks of gray-scale level 129 and light checks of gray-scale level 230 has a mean luminance of approximately 82 cd/m^2 . Finally the uniform image consisting of pixels at gray-scale level 188 also has a mean luminance of 82 cd/m^2 . The underlying spread spectrum waveform was mapped to these images according to a linear relation, with the zero-point of the waveform corresponding to checkerboard 34, and scaled to allow \pm three standard deviations within the range of the images. Again, on every refresh of the computer monitor, the checkerboard image corresponding to the current sample of the input waveform was displayed. In this case, because the mean luminance of all the checkerboards was the same, the EEG data was tagged with the value of the luminance of the light checks minus the luminance of the dark checks. Fig. 1(c) shows three of these constant mean luminance checkerboards.

2.5.3 Multiple Simultaneous Stimuli

To verify the hypothesis that it is possible to detect responses to more than one spread spectrum stimulus simultaneously, two experimental set-ups were employed. These arrangements can be seen in Fig. 1(d). The first consisted of two snowflakes situated 1° to the right and left of a central fixation point marked by a cross hair. The second consisted of a small snowflake occluding a larger snowflake.

For both of these arrangements subjects undertook trials where the modulating waveforms were different instantiations of the same random process, and therefore had identical statistics. For the purposes of illustrating that the input waveform can be shaped as desired and still elicit the desired response, subjects also undertook trials where one of the waveforms was filtered by scaling coefficients corresponding to frequencies below 1 Hz by a factor of 0.1 and those corresponding to frequencies between 1 Hz and 10 Hz by a factor of 0.3. In the case of the bilateral stimuli this shaping was carried out on the waveform controlling the right stimulus and in the case of the concentric stimuli it was the waveform controlling the inner snowflake.

2.6 Experimental Procedure

Subjects were instructed to maintain visual fixation on the centre of the screen for the duration of each testing session. While abstaining from eye-blinks was not possible given the trial lengths, subjects were instructed to keep the number of eye-blinks to a minimum during both the standard and spread-spectrum trials. Subjects were also instructed to keep all other types of motor activity to a minimum during testing.

Each subject underwent three sessions of 120 pattern reversals using the snowflake images and two sessions, again of 120 pattern reversals using checkerboards. Each subject also undertook three sessions of 120 seconds each for the spread spectrum snowflake stimulus and two sessions of 120 seconds for the constant mean luminance checkerboard stimulus. In the case of both of the two-stimuli set-ups two sessions of 120 seconds were performed by each subject for both the case where

the modulating waveforms had identical statistics and the case where one waveform was filtered. This gave a total of eight two-stimuli sessions. The order of presentation of stimuli was counterbalanced between subjects, and no setup was ever undertaken twice in succession.

2.7 Signal Processing

In this case of the spread spectrum stimuli we assume that the EEG response consists of a convolution of the stimulus brightness waveform with an unknown impulse response waveform $w(\tau)$, plus noise. Given the known stimulus waveform and the measured EEG signals, we fit the free parameters of this model, *i.e.* the impulse response function, to the data. The details are shown in Appendix A.

Hereafter we refer to the impulse response $w(\tau)$ as the VESPA. As can be seen in Fig. 2, the VESPA can be thought of as the impulse response which, superimposed of many impulse responses (one per frame), each scaled by the associated input value.

2.7.1 Pre-Processing

Some pre-processing steps were taken. The visual input signal was calculated as the square-wave commands to the monitor, convolved with the video monitor's response function. The EEG was filtered with a high-pass filter with a passband above 2 Hz and a -60 dB response at 1 Hz and a low-pass filter with a 0–35 Hz passband and a -50 dB response at 45 Hz. Impulse responses were measured using a sliding window of 500 ms of data starting 100 ms pre-stimulus. It should be noted that the same filtering was applied to the EEG obtained during the pattern reversal sessions.

2.8 Quantification of Performance

In order to compare the visual evoked potentials obtained by the standard method with the VESPA obtained using the spread spectrum stimuli, three comparison methods were used.

First, correlation values were determined between VEPs and VESPAs for each subject and for

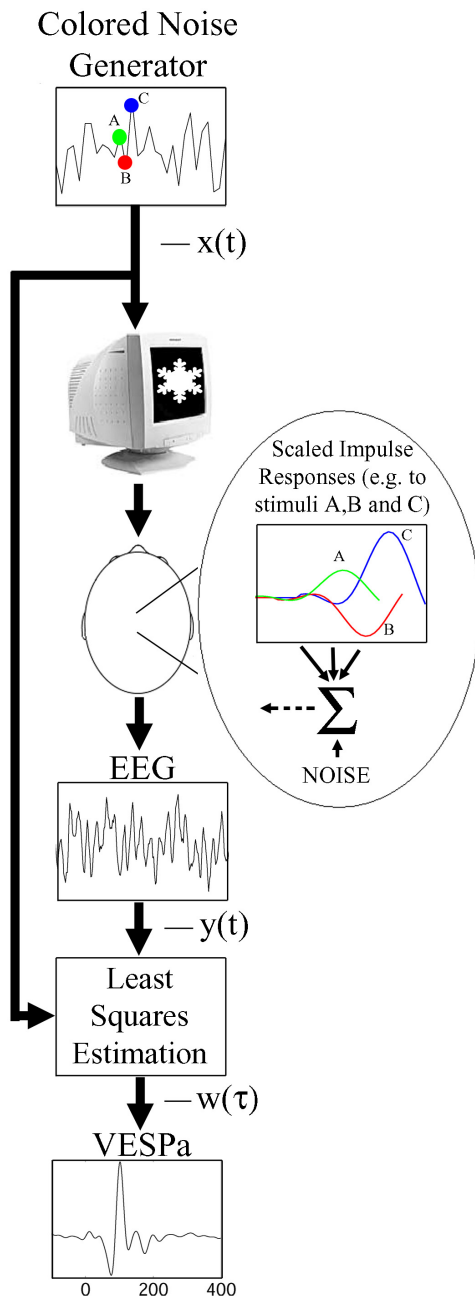


Figure 2: Flow diagram of VESPA acquisition. The EEG is modeled as a sum of overlapping impulse responses scaled by the corresponding stimuli, plus noise. Three such scaled impulse responses are shown, corresponding to stimulus values A, B and C.

each type of stimulus. For this calculation it was assumed that both the VEP and VESPA occurred in the interval 35–175 ms post-stimulus. The correlation C was calculated as

$$C = \frac{\text{COV}(w, v)}{\sqrt{\text{VAR}(w) \text{VAR}(v)}} \quad (1)$$

where $w(t)$ is the VESPA, $v(t)$ is the VEP, COV indicates covariance of a waveform, and VAR the variance. Correlation values were also determined for the group averaged VEP and VESPA for both stimulus types.

Second, the reproducibility of the VESPA was compared to that of the VEP on a subject by subject basis. This was accomplished by calculating correlation values between VEPs obtained from each session with VEPs obtained from every other session for each stimulus type and for each subject and averaging across sessions. In a similar way correlation values between VESPAs were calculated for each stimulus type and each subject and averaged across sessions.

Third, the SNR was calculated for the VEP and for the VESPA at five seconds intervals of every session for each subject and each stimulus type. This resulted in 24 SNR measures per session per subject per stimulus type. These were then averaged across sessions and subjects for each stimulus type to compare the SNR over time for the VESPA and the VEP. These SNRs were calculated by defining the noise as the mean of the squared values in the 100 ms interval immediately preceding the stimulus and the signal was defined as the mean of the squared values in the interval 35–175 ms post-stimulus.

2.9 Software

Calculations were done using Matlab. The scalp maps of Figures 8 were generated using the EEGLAB toolbox (Delorme and Makeig, 2004, <http://sccn.ucsd.edu/eeglab/>). Stimuli were presented using the PRESENTATION stimulus presentation program.




Figure 3: Group average VESPA obtained using constant mean luminance checkerboards at electrode location Oz.

3 Results

3.1 VESPA Responses

Fig. 3 shows the group average VESPA obtained using the constant mean luminance checkerboard stimulus at electrode site Oz. A large response to the stimulus is visible with clear negative peaks at around 75 ms, 125 ms and 175 ms and with clear positive peaks at around 100 ms and 150 ms.

3.2 Comparison of VEP and VESPA

The VEPs obtained by the pattern reversal method using the snowflake stimulus along with the VESPAs obtained using the spread spectrum snowflake stimulus at channel Oz for all subjects are plotted in Fig. 4. The group average is also shown. The VEPs obtained by the pattern reversal method using the checkerboard stimulus along with the VESPAs obtained using the spread spectrum constant mean luminance checkerboard stimulus at channel Oz for all subjects are plotted in Fig. 5. Again, the group average is shown. The VESPAs are plotted as solid lines while the dashed lines correspond to the VEPs.

Table 1(a) shows correlation values between VEPs and VESPAs for both stimulus types for each subject and for the group average. While the

group averages show strong correlations between the VEP and the VESPA ($r = 0.91$, $p < 10^{-28}$ for the checkerboard stimulus and $r = 0.41$, $p < 10^{-3}$ for the snowflake) there is a large degree of variation in the results for individual subjects. For subjects 3, 8 and 12 the VEP and the VESPA are significantly correlated ($p < 10^{-3}$) for both stimulus types. For subjects 1, 4, 5, 6 and 10 the VEP and VESPA are significantly correlated for one stimulus type (with the VEP and VESPA being significantly negatively correlated for subjects 5 and 6 for the other stimulus type). Subjects 2, 7, 9 and 11 show no significant correlation between VEP and VESPA for either stimulus type.

Table 1(b) shows correlation values indicating the reproducibility of the VEPs and VESPAs for both stimulus types for each subject as well as the mean across subjects.

The reproducibility (as determined by within-subject correlations) of the VESPA and VEP resulted in significant correlations for all subjects and for both stimulus types ($p < 10^{-8}$). The mean reproducibility is slightly lower for the VESPA for both stimulus types. Subjects 5, 7, 9 and 10 all had somewhat lower correlation values ($r < 0.85$) for the VESPA for at least one of the stimulus types. For the other subjects the reproducibility was comparable for both VEP and VESPA. In general the checkerboard stimuli appeared to give more reproducible VEPs and VESPAs.

3.3 Required Estimation Time

Fig. 6 shows a plot comparing the SNRs achieved by the various methods at 5 second intervals at electrode Oz averaged across subjects and sessions. The spread spectrum method using the constant mean luminance checkerboards achieves a SNR higher than that obtained using either of the pattern reversal methods at every time point and is almost 5 dB greater than both methods after 120 seconds. The spread spectrum method using snowflake stimuli does not achieve as high a SNR, in general, although it is still higher than both pattern reversal methods after 120 seconds.

As a further demonstration of how rapidly the VESPA can be obtained, Fig. 7 shows the evolution of the VESPA at Oz to the constant mean

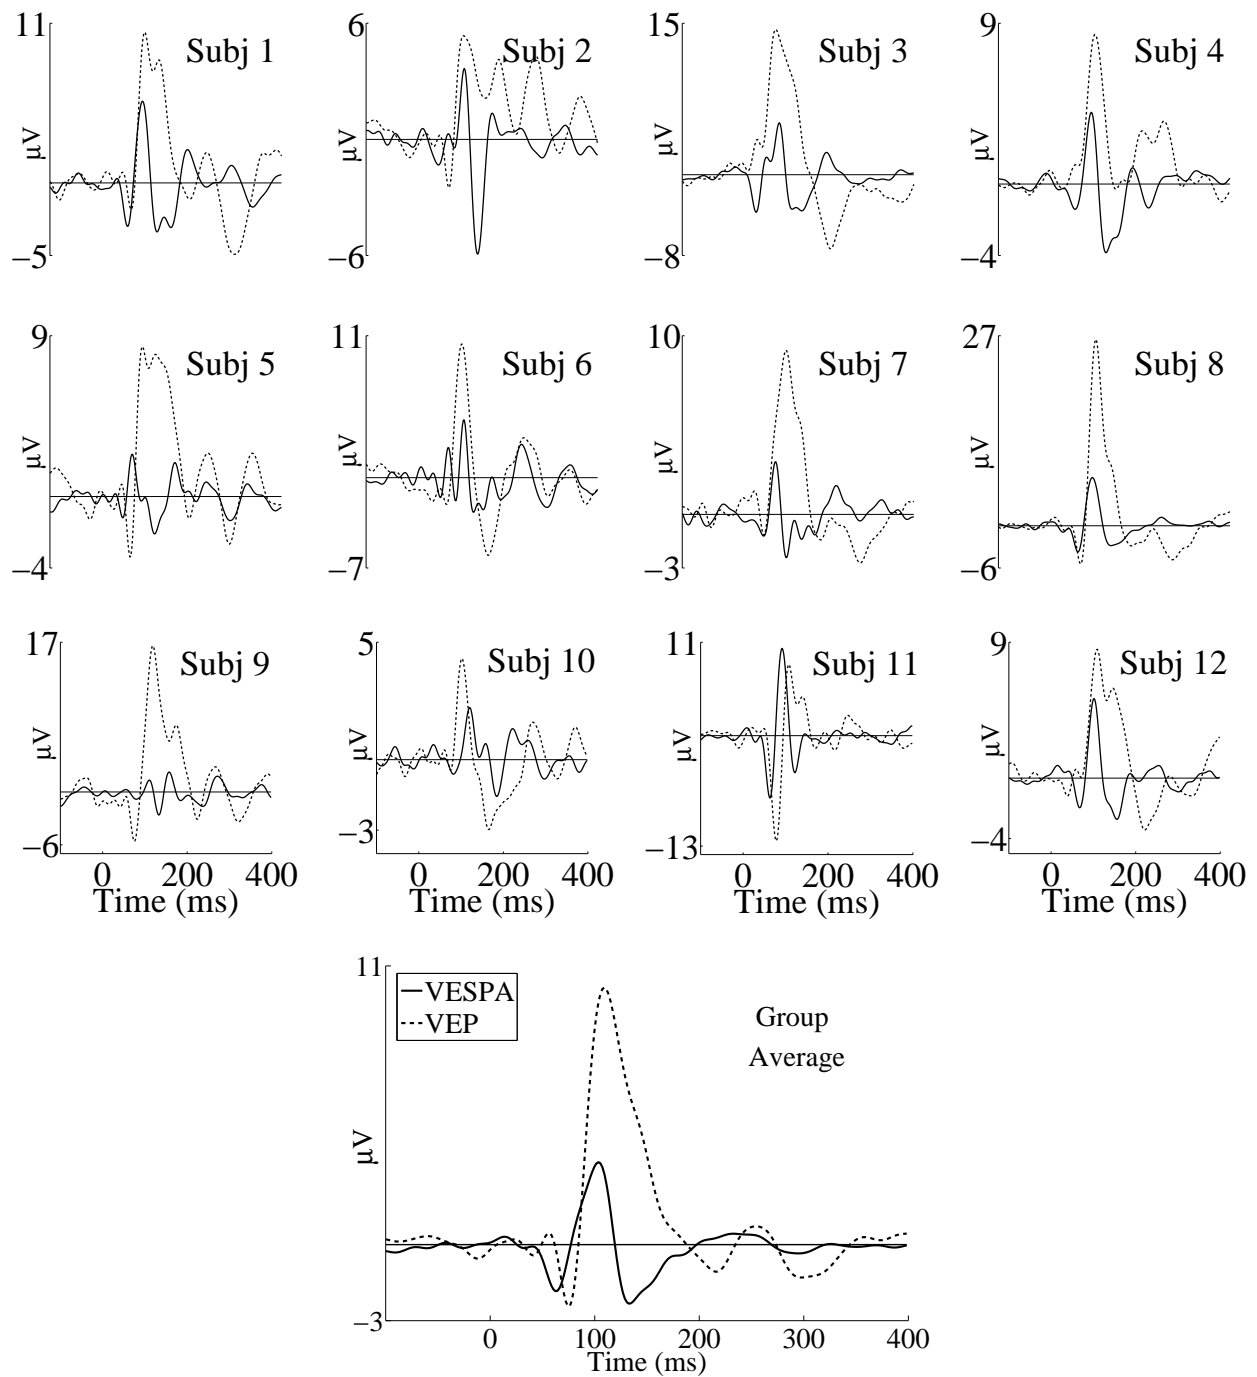


Figure 4: VEPs for all twelve subjects elicited by pattern reversal with snowflake stimulus, and VESPAs elicited by snowflake stimulus. Group average also shown. Potentials measured at Oz. VESPAs plotted as solid lines, and VEPs as dashed lines.

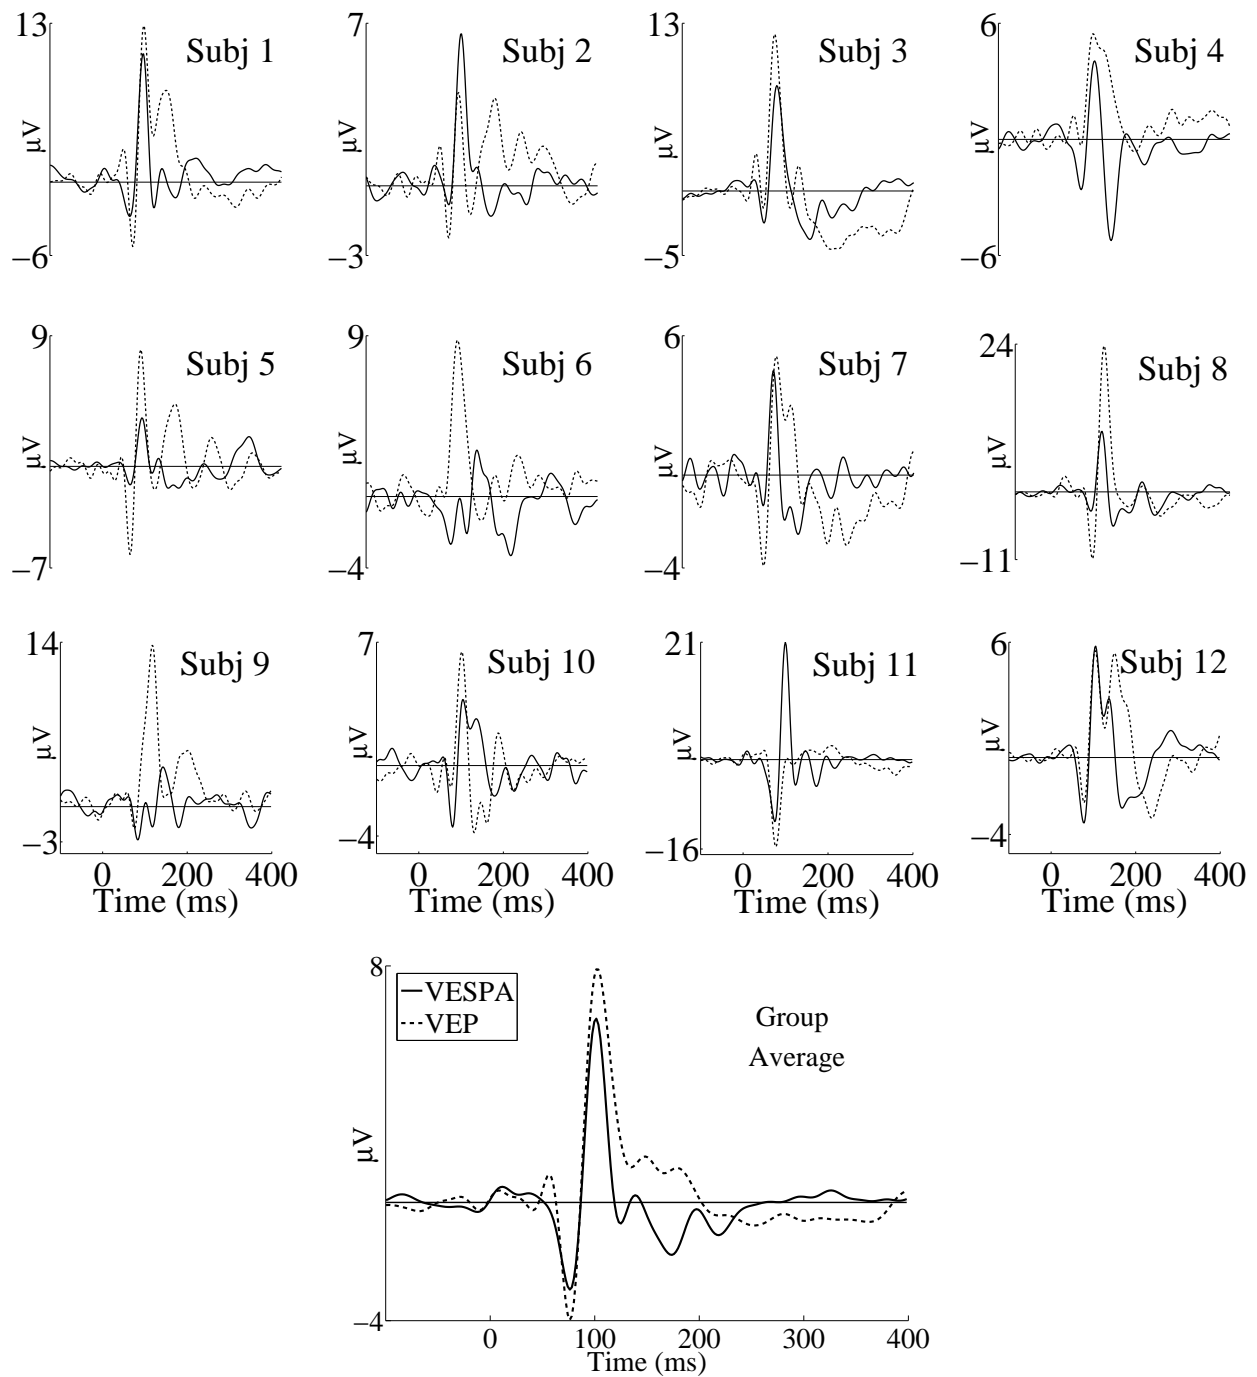


Figure 5: VEPs for all twelve subjects elicited by pattern reversal with checkerboard, and VESPA elicited by the constant mean luminance checkerboard stimuli. Group average also shown. Potentials measured at Oz. VESPAs plotted as solid lines, and VEPs as dashed lines.

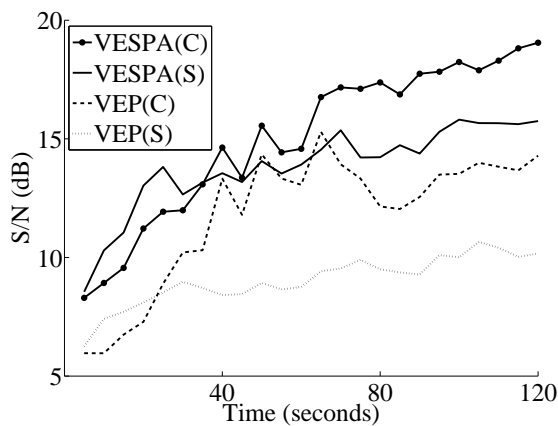


Figure 6: SNR achieved by each method at five second intervals averaged across subjects and sessions at electrode location Oz.

luminance checkerboard stimulus over time with error bars indicating one standard error for subject 7 session 1. The VESPA and error bars are plotted after 10, 20, 30 and 120 seconds to show the amount of time required to obtain a stable estimate.

3.4 Scalp Distribution

Scalp maps indicating the topography of the VESPA elicited by the constant mean luminance checkerboard stimulus at 75, 100, 110, 130 and 155 ms, averaged across subjects and sessions can be seen in Fig. 8. Also plotted are the topographies of the average VEP elicited by the pattern reversed checkerboard stimulus at the same time points.

3.5 Multiple Simultaneous Stimuli

Fig. 9 shows the grand averaged VESPA responses at electrode location Oz to the two simultaneous snowflake stimuli for both setups. The typical VESPA response with a negative peak around 75 ms, positive peak around 100 ms and another negative peak around 120–140 ms can be clearly seen for both stimuli for both setups. Furthermore there appear to be further positive peaks around 200–250 ms.

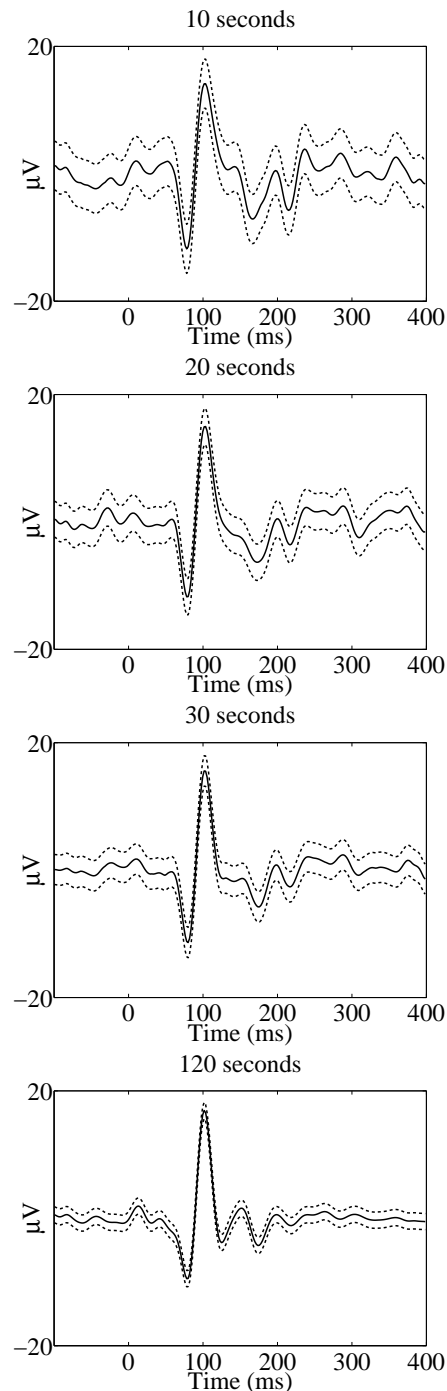


Figure 7: Estimated VESPA \pm one standard error, measured at Oz, for the constant mean luminance checkerboard stimulus, after 10, 20, 30 and 120 seconds. (Subject 7, Session 1.)

Table 1: Statistics showing that the VEP and VESPA are similar but not identical, with a similarity that varies by subject; and that the VEP and VESPA are both quite reproducible across sessions.

(a) Correlation between the VEP and VESPA for each subject, and between the group average VEP and VESPA, for both stimulus types. Data recorded at Oz. Asterisk indicates significance of $p < 0.001$.

Subject	Stimulus	
	Snowflake	Checkerboard
1	0.25	0.67*
2	0.17	0.14
3	0.51*	0.87*
4	0.53*	0.33
5	-0.58*	0.45*
6	0.55*	-0.41*
7	0.02	0.35
8	0.72*	0.68*
9	0.06	-0.24
10	0.40*	0.18
11	-0.08	0.33
12	0.41*	0.59*
Group Ave.	0.41*	0.91*

(b) Correlations showing the reproducibility of the VEP and VESPA for each subject and stimulus type. Numbers shown are the average of correlations of VEPs from pairs of disjoint 120-second sequences, and similarly for the VESPA.

Subject	Snowflake		Checkerboard	
	VEP	VESPA	VEP	VESPA
1	0.91	0.97	0.98	0.95
2	0.89	0.97	0.92	0.93
3	0.98	0.89	0.99	0.91
4	0.93	0.97	0.96	0.89
5	0.96	0.83	0.97	0.95
6	0.96	0.95	0.98	0.90
7	0.96	0.81	0.97	0.93
8	0.99	0.98	0.99	0.98
9	0.96	0.76	0.97	0.73
10	0.90	0.74	0.96	0.87
11	0.98	0.97	0.99	0.98
12	0.98	0.89	0.95	0.99
mean	0.95	0.89	0.97	0.92




Figure 8: Scalp maps showing the topographical evolution (at 75, 100, 110, 130 and 155 ms) of the VESPA (top) and standard VEP (bottom).

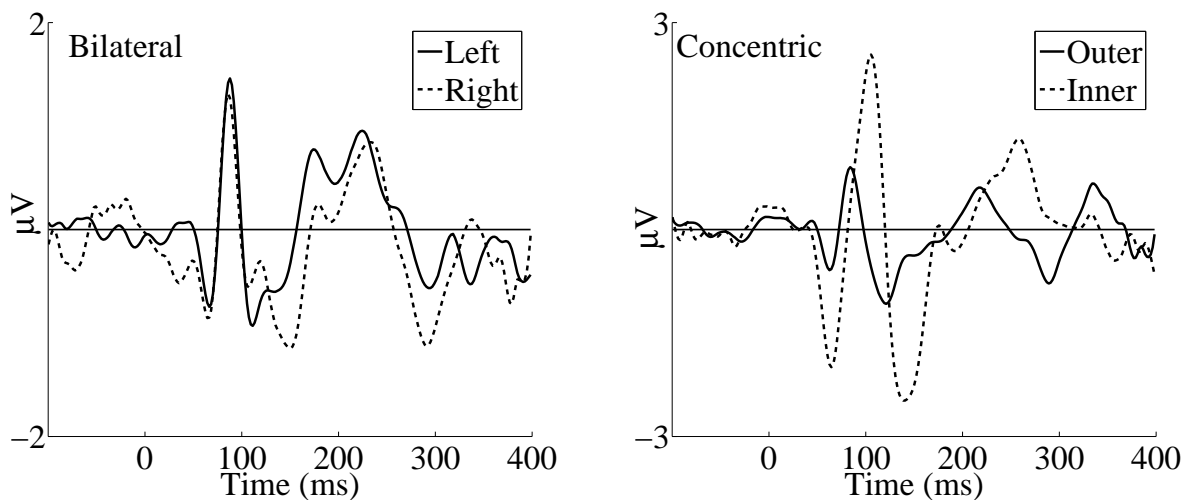


Figure 9: Grand average VESPAs elicited at Oz by the two simultaneous snowflake stimuli for both setups. The left panel plots the responses to the left and right snowflakes during the bilateral setup and the right panel plots the responses to the inner and outer snowflakes during the concentric setup.

In the case of the two bilateral snowflake stimuli setup, the VESPAs are comparable in both magnitude and latency as expected. The prestimulus noise is not quite as low as in the case of the single stimulus experiments described earlier, but this is to be expected given that the subjects were not fixating directly on the stimuli in this setup.

For the VESPAs elicited by the two simultaneous concentric stimuli for, as expected, the VESPA elicited by the central stimulus, on which the subject was instructed to fixate, is significantly larger than that elicited by the surrounding stimulus. However, the VESPA elicited by the surrounding stimulus is clear and, interestingly, some peaks appear to have considerably lower latency than those elicited by the central stimulus.

4 Discussion

We have developed a method for rapidly estimating an impulse response function, the VESPA, that is highly correlated with the visual evoked potential obtained by standard pattern reversal methods. This method uses a continuous stimulus, the luminance of which is unobtrusively mod-

ulated by an underlying spread spectrum waveform. The method may be useful in traditional VEP studies, studies of attention, brain-computer interface research, and the study of specific visual neural pathways.

4.1 Standard VEP as a Special Case

The VESPA is a strict generalization of the conventional VEP. In the special case of a conventional VEP stimulus waveform $x(t)$ consisting of a series of widely spaced impulses, processing through our acquisition and signal processing pipeline results $w(t)$ is precisely the conventional VEP. The conventional VEP can therefore be considered a computationally convenient simplification of this method, where the computational convenience arises because the covariance matrix of the input becomes a scaled identity matrix, $\langle \mathbf{x}_t \mathbf{x}_t^T \rangle = \alpha \mathbf{I}$.

Other systems that rely on binary pulses, such as M-sequences (Buracas and Boynton, 2002), are also a special case of the method proposed here. M-sequences became popular due to their computational efficiency, in that they also result in the input having a simple diagonal covariance ma-

trix. This is because M-sequences are orthogonal to shifted versions of themselves. Furthermore this allows for the simultaneous presentation of several stimuli modulated by different binary M-sequences. Because of the orthogonality of the sequences the responses to each sequence can be very easily determined afterwards. However modern computers are sufficiently powerful to generate and analyze non-binary stimulus waveforms with arbitrary covariance structure, including Gaussian processes. For this reason, we would argue that stimuli should be optimized for considerations like acquisition speed and unobtrusiveness, rather than for computational convenience.

4.2 Sensitivity of the VESPA

One can consider estimation of the impulse response of the EEG potential to the brightness of a visual stimulus as identification of the linear response properties of a noisy black box. Such identification is limited by the richness of the input. For instance, a constant input reveals nothing about the impulse response. In information-theoretic terms, we are attempting to identify the properties of a channel, and such identification is limited by the amount of information actually flowing into the channel. If we consider the power of various possible stimuli in the time-frequency domain, we note that regions of zeros in the power represent lost opportunities. In other words, they represent opportunities to drive the system and thus interrogate it that have been forgone. By putting power in all regions of the time-frequency domain, limited only by considerations of minimizing annoyance to the subject or avoiding saturation of the visual system, the VESPA is able to accelerate the process of acquisition as compared to conventional pulsed stimuli.

Gaussian stimuli have a potential advantage over binary stimuli, in that we would expect them to result in estimated impulse responses which are sharper and more sensitive to changes in brain state and function. The intuition behind this can be seen by consideration of a dramatically simplified system, a simple one-dimensional instantaneous response which saturates outside a region

of roughly linear response (Fig. 10). Let us compare two alternative sets of input stimuli used to linearly approximate the transfer function of this system. The first is a set of extreme inputs, chosen to give maximal (*i.e.* saturated) response (Fig. 10(a)). The second is a set of graded inputs (Fig. 10(b)). A linear fit (dashed line in Fig. 10(a)) to the curve measured using the saturating inputs gives a slope that is systematically much lower than the slope of the response curve at the centre of its dynamic range. This corresponds to systematic underestimation of the gain of the system, or in the non-instantaneous case, to a less peaky and smoother impulse response. The linear response estimated using the saturating inputs will also be insensitive to changes in the shape of the response curve that do not affect its values at those inputs. In contrast, the non-saturating stimuli will result in an estimated linear response (dashed line in Fig. 10(b)) which is sensitive to changes in the shape of the response curve. We might expect this to be reflected in the VESPA having narrower peaks, more structure, and higher between-subject variability than the conventional VEP—an expectation borne out in the results above. This would also lead us to predict that the VESPA will be a more sensitive assay of pathological conditions, and therefore a superior diagnostic tool even apart from considerations of ease of acquisition.

4.3 Comparison of VEP and VESPA

As shown in Fig. 4, Fig. 5, and Table 1(a), the group averaged VEPs and VESPAs are correlated at Oz for both stimulus types. This is particularly noticeable in the case of the checkerboard stimuli ($r = 0.91$, $p < 10^{-28}$). However, the figures and table also indicate a high level of individual variability, which warrants further discussion.

For subjects 3, 8 and 12 the VEP and the VESPA are significantly correlated ($p < 10^{-3}$) for both stimulus types.

For subjects 1, 4, 5, 6 and 10 the VEP and VESPA are significantly correlated for one stimulus type (with the VEP and VESPA being significantly negatively correlated for subjects 5 and 6 for the other stimulus type). In the case of the




Figure 10: Sigmoid curves representing the transfer function of a dramatically simplified one-dimensional instantaneous stateless noise-free visual system. The black dots indicate inputs to the system and the corresponding outputs. In the binary input case (left), there are only two possible inputs, at the extremes, whereas in the Gaussian input case (right), most inputs are far from either extreme. The dashed lines indicate linear approximations to the transfer function based on the samples. In the binary input case the system is saturated, leading to a less crisp linear approximation.

snowflake stimulus for subject 1 the lack of significant correlation seems to be due to different activity after the P1 (100 ms) peak. For subject 4 the VESPA to the checkerboard seems to have a more clearly defined N1 (75 ms) peak than the VEP which may account for the lack of significance of the correlation value. The VEP and VESPA are negatively correlated for the snowflake stimuli for subject 5 and the checkerboard stimuli for subject 6. This appears to be due to differing latencies in the responses. The responses to the snowflake stimuli for subject 10 are not correlated. This also appears to be due to differences in the response latencies.

Subjects 2, 7, 9 and 11 show no significant correlation between VEP and VESPA for either stimulus type. For subject 2 this appears likely to be due to the strong post-stimulus alpha activity in the case of the standard VEP. The responses for subject 7 appear to differ only after the P1 peak. Subject 9 exhibits the poorest quality VESPA in terms of SNR, and hence it is not correlated with the standard VEP. The reason for this is not clear. Subject 11 exhibits the highest quality VESPA in terms of SNR but due to the unusual profile of the standard VEP to the checkerboard stimulus, in particular the lack of a P1, the VEP and VESPA

are not correlated. In this case it appears that the VESPA gives more information than the standard VEP.

The large degree of variability between VEP and VESPA correlations for individual subjects combined with the fact that the group averages are highly correlated suggests that the VESPA gives largely the same information on average as the standard VEP and that the method may in fact be *more* sensitive to individual differences and as a result may be more useful in clinical and research settings. As suggested in section 4.2, this increased sensitivity may be due to the properties of the stimulus. Furthermore the VESPA provides information about the response of the visual system to rapidly changing stimuli as well as to stimuli whose power is spread over specific frequency ranges, which cannot be ascertained by standard VEP methods.

The higher correlations between the group average responses obtained using the checkerboard than the snowflake stimuli and the higher SNR performance of the checkerboards may be due to the fact that the higher spatial frequencies in the stimuli result in them simply being more effective in evoking responses from the visual system.

While a high correlation between the VEP and VESPA for each subject would very simply highlight the usefulness of the method outlined in this study as an alternative for obtaining a standard VEP, a low correlation between the two does not necessarily imply that the VESPA would not be a very useful signal for both clinical and research applications. A crucial factor in assessing its potential usefulness is its reproducibility. Table 1(b) details the reproducibility of the VEP and the VESPA for each subject and for each stimulus type. Clearly, the VESPA is a very reproducible signal. For most subjects the reproducibility of the VEP and the VESPA are comparable, however for subjects 5, 7, 9 and 10 at least one of the stimuli resulted in a noticeably lower reproducibility measure than the measures for their VEPs. Given the small number of sessions performed per subject this may simply be down to artefacts, however other possible factors can be considered. As discussed in a later section (4.6) it is possible that the rapidly changing spread spectrum stimuli ac-

tivate a distinct subpopulation of cells from those activated by the VEP. For some subjects, the measure of the activity as detected at the scalp for this group of cells may not be as good as that relating to the activity of the cells activated by a standard stimulus.

4.4 Required Estimation Time

The VESPA obtained using constant mean luminance checkerboards has a significantly higher SNR for a given time than any of the other methods. Given this finding, it could be suggested that using such a stimulus may not only be a more unobtrusive alternative to standard approaches, but also a superior method in terms of the time taken to obtain responses with specific SNRs. It should be noted that the pattern reversals were carried out every 1000 ms so it would be possible to increase the SNR by reversing the pattern more frequently. However, if the inter-stimulus interval is decreased too much the VEPs begin to overlap and the information inherent in the separate components of the VEP is lost.

It is also important to note that trials with large artifacts were rejected in the case of the pattern reversal method, whereas no artifact rejection of any kind was performed on the spread spectrum data. Further by including parameters corresponding to the 100 ms pre-stimulus, which is necessarily noise, we are degrading the quality of the estimate of the impulse response in the interval of interest. These considerations mean that the VESPA SNR curves are pessimistic and may, perhaps, be improved by using methods such as independent component analysis and by restricting the estimation to only the interval of interest.

The plots in Fig. 7 for the subject with best SNR give another picture of how rapidly the VESPA can be obtained. After just 10 seconds the response is quite clear and after just 20 seconds it is very stable indeed.

4.5 Scalp Distribution

One intriguing possibility for application of this method might be in the isolation of different visual neural pathways by altering the statistical prop-

erties of the input waveform. Scalp topographic mapping of the VESPA under the present stimulus parameters revealed highly specific scalp topographies, quite distinct from that obtained using the standard VEP. The abiding characteristic of the early VESPA maps Fig. 8 was the persistently delimited focus over midline occipital scalp without any evidence for the characteristic early bilateral spread over lateral occipital scalp regions that is consistently seen for the standard VEP (Gomez-Gonzalez et al., 1994; Foxe and Simpson, 2002). This pattern suggests that the VESPA may well have a distinct cellular activation pattern from that of the VEP, favoring midline structures such as striate cortex and neighboring retinotopically mapped extrastriate regions as well as regions in the dorsal visual stream, activation of which are known to produce midline scalp topographies. Previous studies have shown that the bilateral maps found during the early components of the VEP represent, in large part, activation of structures in the ventral visual stream such as the lateral occipital complex (Doniger et al., 2001; Murray et al., 2004).

The first evidence of bilateralization for the VESPA is seen as a robust negativity around 155 ms that is notably different to the VEP topography in the same timeframe. Given that bilateral distributions are first seen for the VEP in the timeframe of the C1 and P1 components more than 80 ms earlier (beginning at just 70 ms), and given what is known about the very rapid spread of activity throughout the dorsal and ventral visual streams (Foxe and Simpson, 2002; Schroeder et al., 1998), this represents a considerable delay in the spread of activity into the ventral stream for the VESPA. One implication of this delay is that the parvocellular system, which provides the major input to the ventral stream, is not effectively activated by the VESPA. Using spread-spectrum stimulation, which in the present implementation was driven mainly at high temporal frequencies, it may be the case that the parvocellular system simply cannot follow the stimulating frequency given the slower response properties of this system (Maunsell et al., 1999).

This could be investigated further in an fMRI study.

4.6 Multiple Simultaneous Stimuli

As can be seen from Fig. 9 the VESPA can be obtained from two simultaneously presented stimuli. These responses can be obtained in both the case where the modulating waveforms have the same statistics and the case where the statistics differ. In the case of the two bilateral stimuli, as expected, the VESPA to the right and left stimuli are very similar in both magnitude and the latency at electrode location Oz. In the case of the concentric stimuli, as expected the VESPA to the centrally located stimulus is significantly larger than that to the surround stimulus, as the subject was instructed to fixate on the centre of the centrally located snowflake. No significant differences were observed between the responses in either set-up in the cases where the waveforms had different properties and the cases where they had the same properties, although further investigation is deemed necessary to fully explore this.

By comparing the latencies of the VESPA in Fig. 9 with that of the VESPA in Fig. 3, it is clear that the latency of the response to stimuli outside the fovea is faster. This may be partly due to the particularly rapid response times of cones in the periphery (Tyler, 1985). This may be further accounted for by the greater ratio of M-cells to P-cells in the periphery than the fovea (Malpeli et al., 1996). M-cells have much more rapid response times and much greater sensitivity to changes in luminance and signals relayed through magnocellular pathways travel with faster conduction speeds than those through parvocellular pathways (Maunsell et al., 1999).

The results of these two simultaneous stimuli experiments suggest the possibility of utilizing these unobtrusive spread spectrum stimuli in the implementation of visual spatial attention studies including visual attention based BCIs (Kelly et al., 2005). In order to further improve the SNR of the VESPA for such an application further signal processing, including higher order modeling and independent component analysis might be helpful.

4.7 Clinical Application

Recent research has identified changes in certain components of the standard visual evoked poten-

tial (Foxe et al., 2005) in patients with schizophrenia that may occur as a result of deficits specific to the dorsal visual stream. It is also suggested that certain ventral stream processes are contingent on inputs from the dorsal stream and as a result failure in these 'higher-level' ventral stream processes may be caused by these underlying dorsal stream deficits. Given that the spread spectrum stimuli outlined in the present study seem to effectively isolate the dorsal visual stream and that they allow for the estimation of VESPA with distinct components, the VESPA may have a significant and immediate impact on research into schizophrenia.

5 Conclusions

The VESPA is a robust visual evoked response that can be rapidly and continuously obtained using stimuli that are less aversive than those used for standard VEPs. It is hoped that it will prove a useful tool for research into schizophrenia and other areas as well as in clinical practice. The ease with which multiple spread spectrum stimuli can be incorporated into more natural experimental paradigms should render them very useful in studies which require short time monitoring of visual attention as well as in the design of brain-computer interfaces.

Acknowledgements

Support for this work was provided by the Irish Higher Education Authority (grant HEA-r9310 to RBR), Science Foundation Ireland (grant 00/PI.1/C067 to the Hamilton Institute at NUI Maynooth), and the US National Institute of Mental Health (grant NIMH-MH65350 to JJF). We thank Dr. Patrick Brennan for valuable assistance and Santiago Jaramillo for helpful discussions.

References

Alexander, K., Rajagopalan, A., Seiple, W., Zemon, V., and Fishman, G. (2005). Contrast response properties of magnocellular and parvocellular pathways in retinitis pigmentosa assessed by

- the visual evoked potential. *Invest Ophthalmol Vis Sci*, 46(8):2967–2973.
- Buracas, G. and Boynton, G. (2002). Efficient design of event-related fMRI experiments using M-sequences. *NeuroImage*, 16(3):801–813.
- Cobb, W. and Dawson, G. (1960). The latency and form in man of the occipital potentials evoked by bright flashes. *J. Physiol*, 152:108–121.
- Coppola, R. (1979). A system transfer function for visual evoked potentials. In Lehmann, D. and Callaway, E., editors, *Human Evoked Potentials: Applications and Problems*, NATO III: Human Factors, Volume 9, pages 69–82. Plenum Press.
- Delorme, A. and Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods*, 134(1):9–21.
- Doniger, G., Foxe, J., Schroeder, C.E. and Murray, M., Higgins, B., and Javitt, D. (2001). Visual perceptual learning in human object recognition areas: A repetition priming study using high-density electrical mapping. *NeuroImage*, 13(2):305–313.
- Fotiou, F., Fountoulakis, K., Iacovides, A., and Kaprinis, G. (2003). Pattern-reversed visual evoked potentials in subtypes of major depression. *Psychiatry Res*, 118(3):259–271.
- Foxe, J., Murray, M., and Javitt, D. (2005). Filling-in in schizophrenia: a high-density electrical mapping and source-analysis investigation of illusory contour processing. *Cereb Cortex*, 15:1914–1927.
- Foxe, J. and Simpson, G. (2002). Flow of activation from v1 to frontal cortex in humans: A framework for defining ‘early’ visual processing. *Experimental Brain Research*, 142(1):139–150.
- Gomez-Gonzalez, C., Clark, V., Fan, S., Luck, S., and Hillyard, S. (1994). Sources of attention-sensitive visual event-related potentials. *Brain Topogr*, 7(1):41–51.
- Gruber, T., Müller, M. M., Keil, A., and Elbert, T. (1999). Selective visual-spatial attention alters induced gamma band responses in the human EEG. *Clin Neurophysiol*, 110(12):2074–2085.
- Halliday, A., McDonald, W., and Mushin, J. (1972). Delayed visual evoked response in optic neuritis. *Lancet*, 6(1):982–985.
- Hubel, D. and Wiesel, T. (1959). Receptive fields of single neurones in the cat’s striate cortex. *J. Physiol.*, 148:574–591.
- Kelly, S. P., Lalor, E. C., Reilly, R. B., and Foxe, J. J. (2005). Visual spatial attention tracking using high-density SSVEP data for independent brain-computer communication. *IEEE Trans. Neu. Syst. & Rehab. Eng.*, 13.
- Kemner, C., Verbaten, M., Cuperus, J., Camfferman, G., and Van Engeland, H. (1994). Visual and somatosensory event-related brain potentials in autistic children and three different control groups. *Electroencephalogr Clin Neurophysiol*, 92(3):225–237.
- Kupersmith, M., Siegel, I., Carr, R., J., R., Flamm, E., and Shakin, E. (1981). Visual evoked potentials in chiasmal gliomas in four adults. *Arch. Neurol.*, 38(6):362–365.
- Malpeli, J., Lee, D., and Baker, F. (1996). Laminar and retinotopic organization of the macaque lateral geniculate nucleus: magnocellular and parvocellular magnification functions. *J Comp Neurol*, 375(3):363–377.
- Markey, H. K. and Antheil, G. (1942). Secret communication system. US Patent 2,292,387.
- Marmarelis, P. Z. and Marmarelis, V. Z. (1978). *Analysis of Physiological Systems: The White Noise Approach*. Plenum Press.
- Matthews, W., Small, D., Small, M., and Pountney, E. (1977). Pattern reversal evoked visual potential in the diagnosis of multiple sclerosis. *J Neurol Neurosurg Psychiatry*, 40(10):1009–1014.
- Maunsell, J., Ghose, G., Assad, J., McAdams, C., Boudreau, C., and Noerager, B. (1999). Visual response latencies of magnocellular and parvocellular lgn neurons in macaque monkeys. *Vis Neurosci*, 16:1–14.
- Müller, M. M., Gruber, T., and Keil, A. (2000). Modulation of induced gamma band activity in the human EEG by attention and visual information processing. *Int. J. Psychophysiol.*, 38(3):283–299.
- Müller, M. M. and Hillyard, S. A. (2000). Concurrent recording of steady-state and transient event-related potentials as indices of visual-

- spatial selective attention. *Clinical Neurophysiology*, 111:1544–1552.
- Müller, M. M., Malinowski, P., Gruber, T., and Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature*, 424(7):309–312.
- Murray, M., Foxe, D., Javitt, D., and Foxe, J. (2004). Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. *J Neurosci*, 24(31):6898–6903.
- Regan, D. (1989). *Human Brain Electrophysiology: Evoked Potentials and Evoked Magnetic Fields in Science and Medicine*. Elsevier.
- Schroeder, C., Mehta, A., and Givre, S. (1998). A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cerebral Cortex*, 8(7):575–592.
- Shibata, T., Shimoyama, I., Ito, T., Abia, D., Iwasa, H., Koseki, K., Yamaouchi, N., Sato, T., and Nakajima, Y. (1999). Attention changes the peak latency of the visual gamma-band oscillation of the EEG. *Neuroreport*, 10(6):167–170.
- Tyler, C. (1985). Analysis of visual modulation sensitivity. ii. peripheral retina and the role of photoreceptor dimensions. *J. Opt. Soc. Am.*, 2:393–398.
- Vaughan Jr., H. and Hull, R. (1965). Functional relation between stimulus intensity and photically evoked cerebral responses in man. *Nature*, 206(985):720–722.
- Yamaguchi, S., Tsuchiya, H., and Kobayashi, S. (1995). Electrophysiologic correlates of visuospatial attention shift. *EEG and Clinical Neurophysiology*, 94(6):450–461.
- This continuous-time equation must be discretized and put in standard form for a least-squares fit. To this end we create a column vector consisting of the sampled points of the response function, $\mathbf{w} = (w(F_0), w(F_0 + \Delta t), \dots, w(F_0 + n_w \Delta t))$, and a column vector consisting of windows of sampled points of the modulating stimulus, $\mathbf{x}_t = (x(t + F_1), x(t + F_1 - \Delta t), \dots, x(t + F_1 - n_w \Delta t))$, where F_0 and F_1 are the limits of the region of support, $w(\tau)$ is allowed to be non-zero only for $F_0 \leq \tau \leq F_1$. The number of samples in the window is $n_w + 1$ where $n_w = N(F_1 - F_0)$, where N is the sampling rate and therefore $\Delta t = 1/N$. For the case of the snowflake stimulus the values of $x(t)$ are simply the luminance values of the displayed snowflakes, while in the case of the constant mean luminance checkerboards the difference in luminance between the light and dark checks of each checkerboard is used.
- Rewriting the model in matrix notation in discrete time, and inserting the result into a formula for the mean squared error, gives:

$$y_t = \Delta t \mathbf{w}^T \mathbf{x}_t + noise \quad (3)$$

$$E = \langle |\Delta t \mathbf{w}^T \mathbf{x}_t - y_t|^2 \rangle \quad (4)$$

where $\langle \cdot \rangle$ indicates an average over t . We wish to fit the column vector \mathbf{w} to a set of input column vectors \mathbf{x}_t and corresponding output scalars y_t so as to minimize E . Expanding $dE/d\mathbf{w} = 0$ yields the linear system

$$\Delta t \langle \mathbf{x}_t \mathbf{x}_t^T \rangle \mathbf{w} = \langle \mathbf{x}_t y_t \rangle. \quad (5)$$

This can be solved for \mathbf{w} either in closed form or by sample-based stochastic methods.

A Estimation of the VESPA

The VESPA is estimated using a linear least squares fit. We use the response model

$$y(t) = w(\tau) * x(t) + noise \quad (2)$$

where $y(t)$ is the measured EEG response, $x(t)$ is the intensity waveform of the stimulus, the symbol $*$ indicates convolution, $w(\tau)$ is the impulse-response function to the brightness of the stimulus, and $noise$ is assumed to be Gaussian.

B Regularization

In order to improve the quality of the estimate of $w(\tau)$, a regularization term was added. This reduces the variance of $w(\tau)$ at the expense of a bias, with a net decrease in the off-sample mean squared error. Standard ridge regression incorporates a quadratic penalty term, $\lambda \int_{F_0}^{F_1} w(\tau)^2 d\tau$, which in discrete time comes to $\lambda \sum_{i=0}^{n_w} \Delta t w_i^2 = \lambda \mathbf{w}^T \mathbf{M} \mathbf{w}$ where the matrix \mathbf{M} , introduced for generality, is here merely a scaled identity matrix,

