

Decoding and Reconstructing the Focus of Spatial Attention from the Topography of Alpha-band Oscillations

Jason Samaha¹, Thomas C. Sprague², and Bradley R. Postle¹

Abstract

■ Many aspects of perception and cognition are supported by activity in neural populations that are tuned to different stimulus features (e.g., orientation, spatial location, color). Goal-directed behavior, such as sustained attention, requires a mechanism for the selective prioritization of contextually appropriate representations. A candidate mechanism of sustained spatial attention is neural activity in the alpha band (8–13 Hz), whose power in the human EEG covaries with the focus of covert attention. Here, we applied an inverted encoding model to assess whether spatially selective neural responses could be recovered from the topography of alpha-band oscillations during spatial attention. Participants were cued to covertly attend to one of six spatial locations arranged concentrically

around fixation while EEG was recorded. A linear classifier applied to EEG data during sustained attention demonstrated successful classification of the attended location from the topography of alpha power, although not from other frequency bands. We next sought to reconstruct the focus of spatial attention over time by applying inverted encoding models to the topography of alpha power and phase. Alpha power, but not phase, allowed for robust reconstructions of the specific attended location beginning around 450 msec postcue, an onset earlier than previous reports. These results demonstrate that posterior alpha-band oscillations can be used to track activity in feature-selective neural populations with high temporal precision during the deployment of covert spatial attention. ■

INTRODUCTION

The visual system is composed of interacting populations of neurons that are specialized for processing distinct stimulus features, such as orientation or spatial location (Riesenhuber & Poggio, 1999). An emerging hypothesis is that oscillatory mechanisms have a role in routing information through networks of these feature-selective neural populations to effect goal-directed behavior (Akam & Kullmann, 2014; Zumer, Scheeringa, Schoffelen, Norris, & Jensen, 2014; Saalmann, Pinsk, Wang, Li, & Kastner, 2012; Jensen & Mazaheri, 2010; Fries, 2005). During covert spatial attention, the scalp distribution of power in the alpha-band (8–13 Hz) of human M/EEG recordings has been observed to covary with attended locations (Capotosto, Babiloni, Romani, & Corbetta, 2009; Van Gerven & Jensen, 2009; Rihs, Michel, & Thut, 2007; Kelly, Lalor, Reilly, & Foxe, 2006; Sauseng et al., 2005; Worden, Foxe, Wang, & Simpson, 2000). Furthermore, stronger phase synchronization between alpha oscillations over frontal sensors and posterior sensors contralateral to the attended hemisphere has also been observed (Sauseng et al., 2005), suggesting that specific topographic phase relationships may code the attended location.

Given the tight correspondence between these alpha-band dynamics and the focus of spatial attention, we in-

vestigated whether topographic patterns of alpha power and/or phase could be used to reconstruct activity in neural populations tuned to spatial location using a multivariate inverted encoding model (IEM). Multivariate IEMs model the relationship between neural activity and stimulus features using hypothesized response profiles, which can then be used to reconstruct the neural representation of novel stimuli that vary with respect to the trained feature (Sprague & Serences, 2015). Previous work applying IEMs to fMRI data has been able to reconstruct the perception of and STM for basic visual features, such as color, orientation, and spatial location (Ester, Sprague, & Serences, 2015; Sprague, Ester, & Serences, 2014; Ester, Anderson, Serences, & Awh, 2013; Kok, Brouwer, van Gerven, & de Lange, 2013; Sprague & Serences, 2013; Ho et al., 2012; Scolari, Byers, & Serences, 2012; Brouwer & Heeger, 2009, 2011). Applying this approach to EEG provides estimates of neural representation on a finer temporal scale (Garcia, Srinivasan, & Serences, 2013) and can be used to assess hypotheses about the functional specificity of oscillatory activity within particular frequency bands.

To test whether neural activity in the alpha-band tracks the focus of spatial attention, we took two approaches. First, we trained a linear support vector machine (SVM) to decode which of six possible spatial locations participants were attending using the topography of different frequency bands. Decoding was only successful for the

¹University of Wisconsin, Madison, ²University of California, San Diego

alpha band. Next, using the time-resolved scalp topography of alpha power (but not phase), we were able to reconstruct the temporal evolution of the deployment of selective attention to six different visual field locations. These results demonstrate the utility of IEMs to measure feature-selective neural representations encoded over multivariate population level activity in near real time.

METHODS

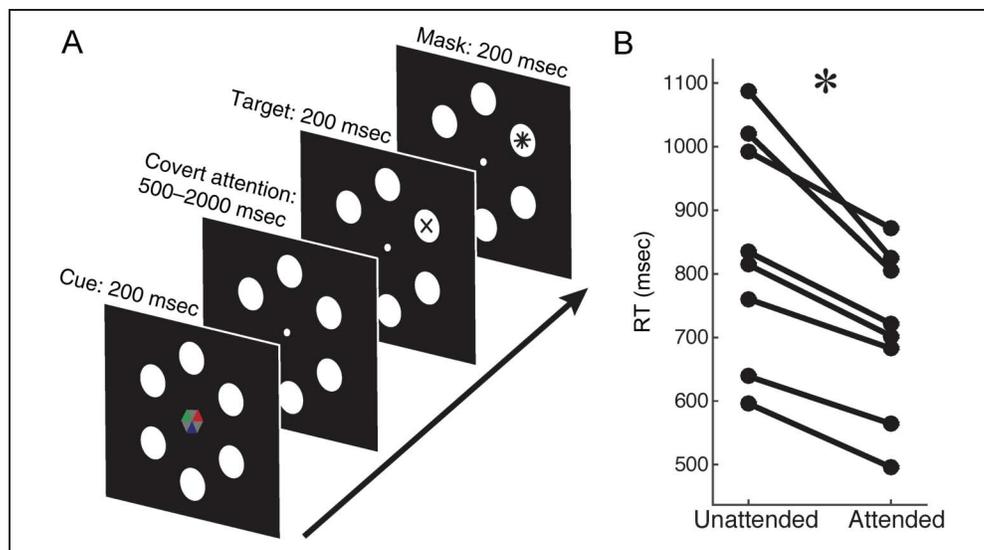
Participants

Data were taken from the publicly available BNCI Horizon 2020 research consortium database (www.bnci-horizon-2020.eu) and were first reported by Treder, Bahramisharif, Schmidt, van Gerven, and Blankertz (2011). Eight healthy volunteers (seven men; aged 18–27 years) participated in this study. All had normal or corrected-to-normal vision. All participants gave written consent, and the study was performed in accordance with the Declaration of Helsinki.

Stimuli and Procedures

Each trial began with a central fixation dot surrounded by six white disks (eccentricity: 9° , size: 3.27° of visual angle) lasting 1000 msec. A cue indicating the location to attend appeared in the center of the screen for 200 msec. After a variable duration of 500–2000 msec, a 200-msec target appeared in either the attended circle (80% of trials) or a different, randomly selected location (20% of the trials). To increase difficulty, a 200-msec mask (“*”) was presented after each target (see Figure 1). Stimuli were presented on a 24-in. thin-film transistor LCD screen with a refresh rate of 60 Hz and a resolution of 1920×1200 pixels. The experiment was implemented in Python using the open-source BCI framework Pyff (Venthur et al., 2010) with Pygame (www.pygame.org).

Figure 1. Task design and behavioral data. (A) Covert spatial attention was directed with an omnidirectional symbolic cue (80% valid) that indicated in which of six spatial locations an upcoming target was likely to appear. Participants were instructed to attend to the cued location while maintaining fixation and then decide if the target was a “x” or a “+.” (B) RT decreased when targets appeared in cued, as contrasted with uncued locations ($t(1, 7) = 5.62$, $p < .001$), suggesting that participants were indeed attending the cued location. Each line represents one participant.



Targets were either a “x” or a “+” and were determined randomly on each trial. To reduce lateralized preparatory motor activity, participants had to indicate which target they saw via a thumb press on a button box in their right or left hand. To ensure that the cue did not have an orientation that covaried with the orientation of the attended location (otherwise decoding and reconstructions could be based on cue orientation, rather than spatial attention, per se), a symbolic, omnidirectional color cue was used. This was a hexagon with each of the six faces pointing to one of the target discs. Three of the faces were gray, and the other three were colored blue, red, and green. One of these colors was assigned as the cue (color counterbalanced across participants). In 50% of trials, the cue–target interval was 2000 msec; it varied randomly between 500 and 2000 msec on the remaining trials. This helped ensure that participant’s attention was sustained throughout the trial. The task was 600 trials spread over six blocks of 100 trials each, with ~2-min rests in between blocks. Participants were instructed to maintain fixation throughout the trial and to respond as quickly and accurately as possible.

EEG Acquisition and Analysis

EEG was recorded at 1000 Hz from a Brain Products (Munich, Germany) 64-channel actiCAP. Electrodes were referenced against a nose electrode, and impedances were kept below 20 k Ω . Eye movements were monitored with an EOG electrode placed below the right eye. Referencing Fp2 against the EOG and F10 against F9 created vertical and horizontal bipolar EOG channels, respectively. Data were downsampled offline to 200 Hz.

Only trials where the cue–target interval was 2000 msec were analyzed. Data from these trials were cut into –800:2500 msec epochs, relative to cue onset. Each epoch was carefully screened for eye-related or muscle-related

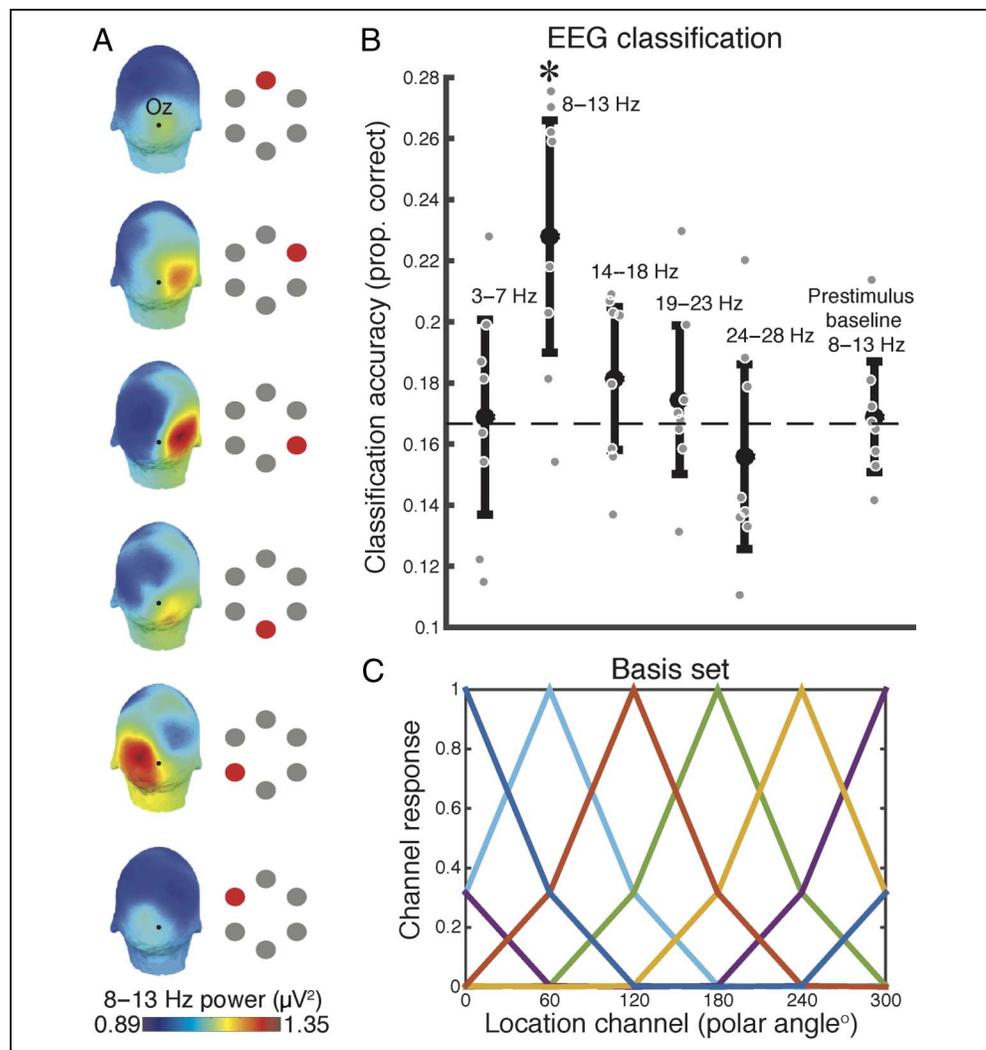
artifacts, leaving an average of 241 artifact-free trials per participant, and an average of 40 observations for each attended location per participant. Time–frequency decomposition was performed using EEGLAB (Delorme & Makeig, 2004) and custom code, running in a MATLAB environment (The Mathworks, Natick, MA). Single trial data were convolved with a family of wavelets, spanning 4–30 Hz, in 1-Hz steps, with wavelet cycles increasing linearly between three and eight cycles as a function of frequency. Power and phase were extracted from the resulting complex time series by squaring the absolute value of the time series (μV^2) in the case of power and by taking the phase angle (Matlab function *angle.m*) in the case of phase. Alpha power was defined by averaging power between 8–13 Hz. The circular average of this frequency range created the phase time series. Power data were normalized by dividing power at each time point by the average power from the precue baseline period of –350 to –50 msec. Because the general linear model used in the IEM may not optimally capture relationships with circular variables-like phase, we also conducted the same IEM phase analysis (see below Multivariate IEM) using the complex values from the wave-

let decomposition after normalizing the amplitude components by dividing each complex number by the absolute value of the real component of the number. This analysis resulted in channel responses virtually identical to those we report based on phase angles only (see Results).

Multivariate Classification

To assess whether the spatial distribution of power during the cue–target interval could be used to decode the attended location, we applied an SVM classifier to our data (implemented in the MATLAB Machine Learning Toolbox). SVMs were constructed with a linear kernel and default parameters supplied by the software (box constraint = 1, sequential minimal optimization separating algorithm). For each electrode, data were averaged into different frequency bands between 3 and 28 Hz (see Figure 2B) and then averaged over the cue–target interval 1000–1900 msec. We focused on this interval to avoid cue-evoked responses and temporal smearing of target-evoked responses because of wavelet decomposition. This procedure resulted in one power value per channel, per frequency

Figure 2. Classification and IEM. (A) The topography of alpha-band power (8–13 Hz) over posterior electrodes during sustained attention (1000–1900 msec) shows a clear spatial correspondence with the attended location (indicated by a red dot). (B) A linear SVM classifier trained to decode the attended location using the topography of power from 1000 to 1900 msec could successfully classify the attended location, but only from activity in the alpha band ($t(1, 7) = 3.46$, $p = .010$). As a control, classification did not differ from chance (dashed line) when precue alpha power was used ($t(1, 7) = 0.29$, $p = .774$). Error bars represent 95% confidence intervals. Each gray point represents a participant's average classification accuracy. (C) The basis set used for the IEM contained six tuning functions, one for each spatial location, that represent the idealized response of the neural population tuned to each location.



band, as input to the classifier. Leave-one-trial-out cross-validation, iterated over every trial, ensured that training and testing data were independent. We evaluated the statistical significance of classification against theoretical chance (1/6), but also because this comparison can overestimate significance with a small number of observations (Combrisson & Jerbi, 2015), we also compared performance against classification obtained from a prestimulus baseline (−350 to −50 msec).

Multivariate IEM

Following similar approaches (Garcia et al., 2013; Brouwer & Heeger, 2009, 2011), we modeled the response in each electrode as a linear sum of six information channels (one per polar angle location) composed of half-rectified sinusoids raised to the fifth power and centered on each polar angle (0°, 60°, 120°, etc.; Figure 2C). In the first step, “training data” from all but one trial (“test data”) is used in a general linear model of the form:

$$B_1 = WC_1$$

where B_1 (m electrodes $\times n$ trials) is the observed signal at each electrode (either power values or phase angles) for each training trial, C_1 (k channels $\times n$ trials) is a matrix of predicted responses for each information channel on each trial, and W (m electrodes $\times k$ channels) is a weight matrix that characterizes the mapping from “channel space” to “electrode space.” The weight matrix W (m electrodes $\times k$ channels) can be derived using ordinary least-squares regression as follows:

$$W = B_1 C_1^T (C_1 C_1^T)^{-1}$$

Next, the model is inverted to transform the observed test data B_2 (m electrodes $\times 1$ trial) into a set of estimated channel responses, C_2 (k channels $\times 1$ trial), using the weights W derived from the training data, via the equation:

$$C_2 = (W^T W)^{-1} W^T B_2$$

This procedure was iterated until every trial served as a testing set (i.e., leave-one-trial-out cross-validation). The estimated channel responses were then circularly shifted to a common center (0°) and averaged across trials. Note that the response in the most distant orientation channel (180°) in Figures 3 and 4 was duplicated for display symmetry. To construct tuning functions over time (Figures 3A and 4A), this procedure was performed for every sample of phase or power, after downsampling to 100 Hz.

To statistically assess the robustness of the reconstructed tuning functions, we compared the amplitude of the on-channel response (i.e., the attended location) averaged over the covert attention interval 1000–1900 msec to both the amplitude of the response during a prestimulus baseline (−350 to −50) and to the amplitude of the

most distant nonattended channel (i.e., 180°) during attention using a nonparametric permutation procedure. Condition labels were randomly shuffled, and condition differences were computed and saved. This was repeated 20,000 times, forming a distribution of difference scores under the null hypothesis that there is no amplitude difference between conditions. If the true difference between conditions exceeded the 95 percentile of this distribution, it was considered statistically significant at $\alpha < 0.05$. The selectivity of the tuning functions was also analyzed across time (Figures 3A and 4A) by using linear regression to estimate the slope of responses of channels tuned to equidistant locations at each time point (Foster, Sutterer, Serences, Vogel, & Awh, 2015). More positive slopes indicate greater selectivity to the attended location, whereas a large negative slope would indicate selectivity to the location opposite the one cued. 99% confidence intervals were generated empirically with a bootstrap analysis in which data from all participants were pooled and sampled with replacement 80,000 times, generating a distribution of slopes around the group mean slope. If zero (the null hypothesis of no selectivity) fell outside the lower 99th percentile of this distribution, it was considered statistically significant at $p < .01$, one-tailed.

Mutual Information Analysis

To quantify which electrodes contained information about the attended location, we computed the mutual information (MI) between each electrode and the cued location using code provided in Cohen (2014). MI describes the extent to which uncertainty about the state of one variable (e.g., cued location) is reduced by knowledge of the state of another variable (e.g., electrical brain activity; Shannon, 1948). When MI equals 0 bits, the two variables are statistically independent. MI was computed for each electrode using a number of bins determined by the rule of Freedman and Diaconis (1981).

RESULTS

Behavior

RTs to targets were significantly lower on validly versus invalidly cued trials. ($t(1, 7) = 5.62, p < .001$; see Figure 1B), which also held true when using a log transform ($t(1, 7) = 5.92, p < .001$), and when no trials were excluded based on oculomotor artifacts ($t(1, 7) = 4.49, p < .01$). Mean discrimination accuracy was 95% correct ($SEM = 1.8\%$) following valid cues and 79% correct ($SEM = 11\%$) following invalid cues. Although this difference did not reach statistical significance ($p = .23$), it suggests that there was no speed–accuracy tradeoff. These results indicate that participants were indeed allocating attention to the cued location.

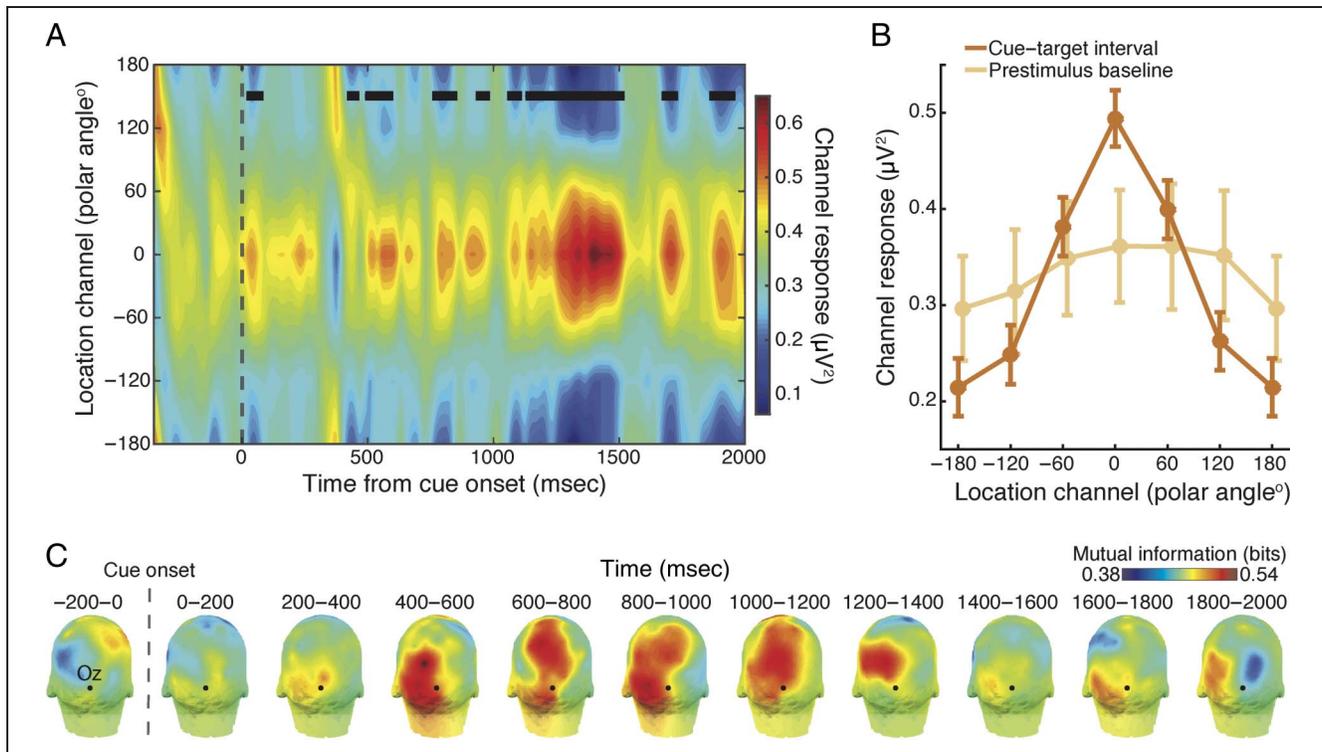


Figure 3. Reconstructed tuning functions from alpha power. (A) Using the topography of 8–13 Hz power over time, we were able to reconstruct the spatial focus of attention. Tuning functions for each polar angle location were circularly shifted and averaged so that 0° corresponds to the response of the channel representing the attended location on that trial. The reconstructed tuning functions demonstrate a characteristic Gaussian shape, with a maximal response at the attended location (0°) and a fall off in response for channels tuned to locations further away. Time points with significant selectivity ($p < .01$), as measures by the slope of the tuning function at each time point, are indicated with black squares. (B) Comparison of channel tuning during the prestimulus baseline (-350 to -50 msec) and covert attention interval (1000 – 1900 msec). The amplitude of the attended channel response (0°) was significantly higher compared with both baseline ($M_{\text{diff}} = 0.125 \mu\text{V}^2$, $p = .017$) and to the channel tuned to the most distant location (i.e., 180° ; $M_{\text{diff}} = 0.279 \mu\text{V}^2$, $p = .001$), indicating that the reconstructions depended on the allocation of spatial attention. Error bars indicate \pm within-subject *SEM*. (C) As a metric of which regions were most informative for the reconstructions, we computed MI between alpha power and attended location for each electrode for 11 time windows. This analysis reveals a clear posterior distribution of electrodes that carried the most information about the attended location, beginning around 400 msec.

Decoding Attended Location

The topography of alpha power varied systematically with attended location (Figure 2A), which allowed for significant classification, as compared against both theoretical chance (i.e., $1/6$, $t(1, 7) = 3.80$, $p = .006$) and an empirical baseline ($t(1, 7) = 3.46$, $p = .010$; Figure 2B). Extending prior work (Treder et al., 2011; Bahramisharif, van Gerven, Heskes, & Jensen, 2010; Van Gerven & Jensen, 2009), we found that successful decoding was specific for the alpha band; no other frequency range analyzed allowed for significant classification ($ps > .177$).

Reconstructing the Focus of Attention

To determine whether patterns of alpha-band oscillations could be used to reconstruct population level, feature-selective tuning profiles, we implemented two IEMs, one based on the topography of alpha power and one based on alpha phase topographies. IEMs can be thought of as a form of targeted data reduction that allows for quantification of characteristics of the underlying feature-specific

neural population, such as its response selectivity (Sprague & Serences, 2015). Figure 3A shows the temporal evolution of reconstructed channel tuning functions, with time points with significant selectivity marked with black squares. The response shows a characteristic Gaussian profile, with maximal amplitude for the attended location (0°) and a steady reduction in responses from neural populations tuned to locations further away, with this profile beginning around 450 msec and persisting throughout the sustained attention interval. To further quantify these reconstructions, we compared the amplitude of the response in the attended channel with the response in that channel during the baseline interval using a permutation test (see above Multivariate IEM). This revealed significantly higher amplitude responses from neural populations tuned to the attended location during attention as compared with baseline ($M_{\text{diff}} = 0.125 \mu\text{V}^2$, $p = .017$) as well as significantly higher responses in the attended channel than the channel tuned 180° away, ($M_{\text{diff}} = 0.279 \mu\text{V}^2$, $p = .001$), indicating that the reconstructions tracked the deployment and the location of spatial attention (Figure 3B). An MI analysis revealed a posterior

distribution of informative electrodes (Figure 3B), consistent with the topographies in Figure 2A.

Because spatial attention may also be supported by long-distance phase synchronization in the alpha band (Sauseng et al., 2005), we determined whether phase topographies could be used to reconstruct the attended location. If, for example, the phase relationship between frontal and posterior regions reliably covaries with attended location, reconstructions should be successful. In contrast to this proposal, however, we did not observe robust reconstructions based on alpha-band phase topography. Figure 4A shows the reconstructed tuning functions over time, which suggests a transient, cue-evoked response but no sustained tuning during the attention interval. This was supported statistically via comparison with both the baseline attended channel response ($p = .54$) and the response in the channel tuned 180° away from the attended location ($p = .78$; Figure 4B). These results were virtually identical when using amplitude-normalized complex values as input to the phase-based IEM (see Methods). As anticipated, the MI analysis revealed no clear scalp topography of location-informative electrodes.

To assess whether eye movements may have been a confound in these results, we tested for a correspondence between EOG activity and cued location in several ways. Horizontal EOG amplitudes during sustained attention did not differentiate attention to horizontal from ver-

tical locations ($t(1, 7) = 1.8, p = .12$). SVM classifiers trained on horizontal and vertical EOG activity at each sample during the epoch resulted in above-baseline classification at only one sample (out of 200) during the sustained attention interval. Because this sample was not neighbored by other significant samples, it is likely spurious. Lastly, EOG-based classification accuracy did not reliably predict EEG-based alpha decoding accuracy ($r = 0.37, p = .36$), in line with a previously published decoding analysis of this same data set (Treder et al., 2011) that also found no correlation ($r = 0.029, p = .75$) when examining EEG and EOG accuracies to the best classifiable pairs of locations. These analyses suggest that eye movements do not explain the EEG results.

DISCUSSION

We present evidence that the topography of alpha-band power allows for tracking the focus of spatial attention with relatively fine temporal resolution. Our decoding analysis revealed that spatial location information was predominantly carried in the alpha band. Our IEM revealed that population level neural activity in the alpha band showed robust spatial tuning during covert attention. Mechanistically, it has been proposed that alpha oscillations play a role in coordinating information flow between feature-selective neural populations (Saalmann

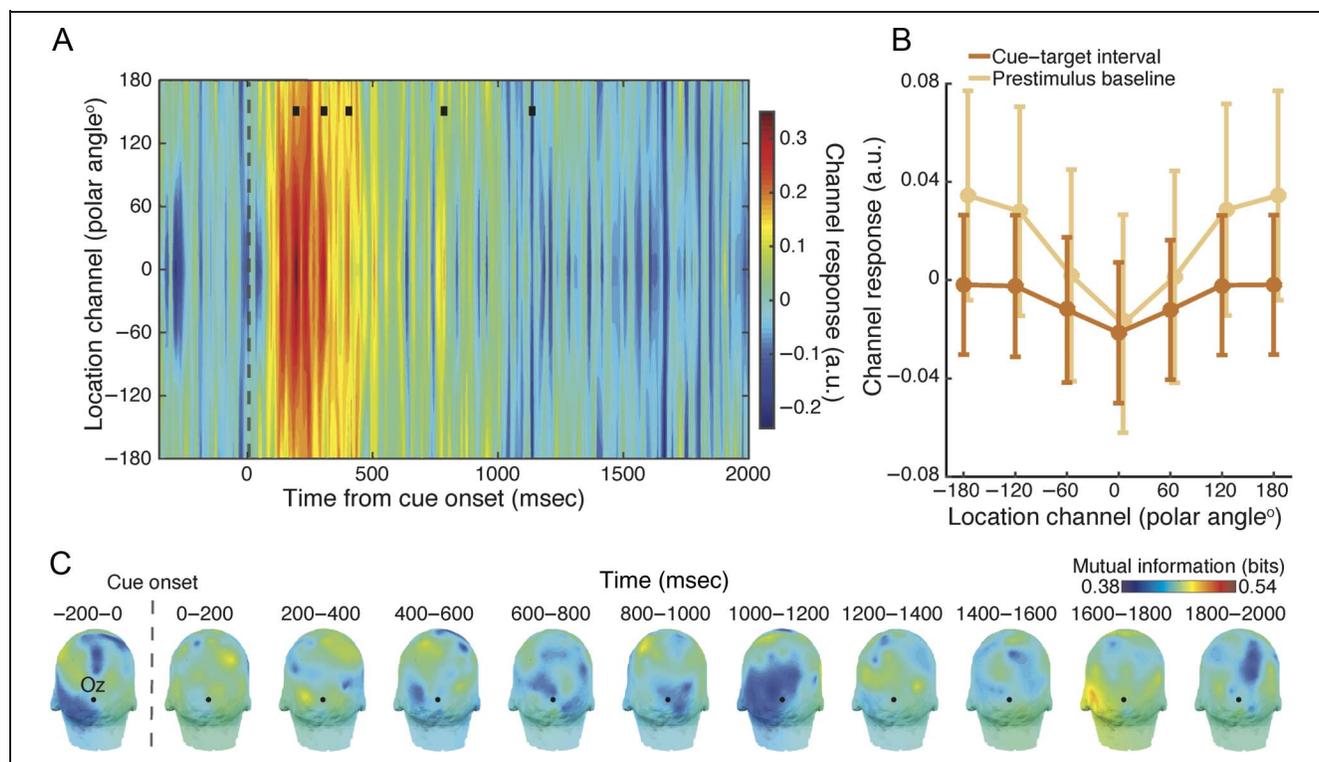


Figure 4. Reconstructed tuning functions from alpha phase. (A) In contrast to the alpha power results, reconstructions based on the topography of phase angles did not reveal sustained location tuning during spatial attention (black squares indicate $p < .01$). A transient response was observed immediately following cue onset, likely reflecting phase locking (e.g., an evoked response) from the cue. (B) The attended-channel response during attention was statistically indistinguishable from baseline and from the response in the channel tuned 180° away ($ps > .54$). Error bars indicate \pm within-subject SEM. (C) MI analysis shows no clear topography of information between phase and attended location.

et al., 2012; Siegel, Donner, & Engel, 2012; Jensen & Mazaheri, 2010). The fact that the topography of our MI analysis was maximal over occipital and parietal sensors is consistent with the idea that alpha rhythms may be responsible for coordinating activity between spatial receptive field maps in visual and parietal cortex (Saygin & Sereno, 2008; Silver, Ress, & Heeger, 2005), rather than within only a single feature-selective region.

That the IEM based on alpha-band phase topography failed to reconstruct the focus of sustained attention suggests that phase relationships across electrodes may not contain feature specific information. Prior findings of long-distance phase coupling during spatial attention (Sauseng et al., 2005) may reflect a more general signature of the engagement of top-down control mechanisms in frontoparietal attentional systems (Kundu, Johnson, & Postle, 2014; Capotosto et al., 2009). Although alpha phase may be controlled in other ways (Samaha, Bauer, Cimaroli, & Postle, 2015; Bonnefond & Jensen, 2012), these findings suggest a privileged role for the dynamics of alpha-band power in the implementation of covert spatial attention.

Multivariate, information-based techniques, such as IEMs, are offering new insight about the mechanisms that underlie core cognitive functions (Naselaris & Kay, 2015; Serences & Saproo, 2012). For covert spatial attention, results from univariate analyses have been interpreted as evidence for multistage accounts of attentional control, with early bilateral “alpha desynchronization” (emerging within 700 msec postcue) reflecting a re-orienting operation and a later pattern of sustained asymmetric “alpha synchronization,” beginning around 700 msec postcue, corresponding to sustained attention (Rihs, Michel, & Thut, 2009). Results from our IEM, however, do not follow this pattern. Instead, the time course of the power-based reconstructions in Figure 3A showed a short-lived window of significant selectivity concurrent with cue onset (likely driven by a cue-evoked response, rather than spatial attention) with relatively sustained tuning emerging around 450 msec. This estimate of 450 msec is congruent with the well-characterized behavioral time course of cued shifts of spatial attention, which reveals benefits around 400–500 msec postcue (Müller, Teder-Sälejärvi, & Hillyard, 1998; Cheal, Lyon, & Gottlob, 1994). Future work will be needed to reconcile the interpretations arising from univariate versus multivariate analyses. These results demonstrate that IEMs provide a useful new technique to track the deployment of spatial attention in near real time from extracranial signals.

Acknowledgments

We would like to thank Dr. Barry Giesbrecht and Dr. Javier Garcia for code and analytic guidance. This study was funded by MH095984 to Bradley R. Postle.

Reprint requests should be sent to Jason Samaha, 1202 W. Johnson Street, Madison, WI 53703, or via e-mail: jsamaha@wisc.edu.

REFERENCES

- Akam, T., & Kullmann, D. M. (2014). Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nature Reviews Neuroscience*, *15*, 111–122.
- Bahramisharif, A., van Gerven, M., Heskes, T., & Jensen, O. (2010). Covert attention allows for continuous control of brain-computer interfaces. *European Journal of Neuroscience*, *31*, 1501–1508.
- Bonnefond, M., & Jensen, O. (2012). Alpha oscillations serve to protect working memory maintenance against anticipated distracters. *Current Biology*, *22*, 1969–1974.
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, *29*, 13992–14003.
- Brouwer, G. J., & Heeger, D. J. (2011). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, *106*, 2108–2119.
- Capotosto, P., Babiloni, C., Romani, G. L., & Corbetta, M. (2009). Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. *Journal of Neuroscience*, *29*, 5863–5872.
- Cheal, M. L., Lyon, D. R., & Gottlob, L. R. (1994). A framework for understanding the allocation of attention in location-precued discrimination. *Quarterly Journal of Experimental Psychology Section A*, *47*, 699–739.
- Cohen, M. X. (2014). *Analyzing neural time series data: Theory and practice*. Cambridge, MA: MIT Press.
- Combrisson, E., & Jerbi, K. (2015). Exceeding chance level by chance: The caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. *Journal of Neuroscience Methods*, *250*, 126–136.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21.
- Ester, E. F., Anderson, D. E., Serences, J. T., & Awh, E. (2013). A neural measure of precision in visual working memory. *Journal of Cognitive Neuroscience*, *25*, 754–761.
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*, *87*, 893–905.
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2015). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*, *115*, 168–177.
- Freedman, D., & Diaconis, P. (1981). On the histogram as a density estimator: L2 theory. *Probability Theory and Related Fields*, *57*, 453–476.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*, 474–480.
- Garcia, J. O., Srinivasan, R., & Serences, J. T. (2013). Near-real-time feature-selective modulations in human cortex. *Current Biology*, *23*, 515–522.
- Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E.-J., & Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *Journal of Neuroscience*, *32*, 7992–8003.
- Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: Gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 186.
- Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active

- retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of Neurophysiology*, *95*, 3844–3851.
- Kok, P., Brouwer, G. J., van Gerven, M. A. J., & de Lange, F. P. (2013). Prior expectations bias sensory representations in visual cortex. *Journal of Neuroscience*, *33*, 16275–16284.
- Kundu, B., Johnson, J. S., & Postle, B. R. (2014). Trait-like differences in underlying oscillatory state predict individual differences in the TMS-evoked response. *Brain Stimulation*, *7*, 234–242.
- Müller, M. M., Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature Neuroscience*, *1*, 631.
- Naselaris, T., & Kay, K. N. (2015). Resolving ambiguities of mvpa using explicit models of representation. *Trends in Cognitive Sciences*, *19*, 551–554.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019–1025.
- Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by α -band EEG synchronization. *European Journal of Neuroscience*, *25*, 603–610.
- Rihs, T. A., Michel, C. M., & Thut, G. (2009). A bias for posterior α -band power suppression versus enhancement during shifting versus maintenance of spatial attention. *Neuroimage*, *44*, 190–199.
- Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X., & Kastner, S. (2012). Pulvinar regulates information transmission between cortical areas based on attention demands. *Science*, *337*, 753–756.
- Samaha, J., Bauer, P., Cimaroli, S., & Postle, B. R. (2015). Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. *Proceedings of the National Academy of Sciences*, *112*, 8439–8444.
- Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., et al. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, *22*, 2917–2926.
- Saygin, A. P., & Sereno, M. I. (2008). Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. *Cerebral Cortex*, *18*, 2158–2168.
- Scolari, M., Byers, A., & Serences, J. T. (2012). Optimal deployment of attentional gain during fine discriminations. *Journal of Neuroscience*, *32*, 7723–7733.
- Serences, J. T., & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, *50*, 435–446.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*, 379–423.
- Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience*, *13*, 121–134.
- Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology*, *94*, 1358–1371.
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, *24*, 2174–2180.
- Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*, 1879–1887.
- Sprague, T. C., & Serences, J. T. (2015). Using human neuroimaging to examine top-down modulation of visual perception. In B. U. Forstmann & E. J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 245–274). New York: Springer.
- Treder, M. S., Bahramisharif, A., Schmidt, N. M., van Gerven, M. A., & Blankertz, B. (2011). Brain-computer interfacing using modulations of alpha activity induced by covert shifts of attention. *Journal of NeuroEngineering and Rehabilitation*, *8*, 24.
- Van Gerven, M., & Jensen, O. (2009). Attention modulations of posterior alpha as a control signal for two-dimensional brain-computer interfaces. *Journal of Neuroscience Methods*, *179*, 78–84.
- Venthur, B., Scholler, S., Williamson, J., Dähne, S., Treder, M. S., Kramarek, M. T., et al. (2010). Pyff—A pythonic framework for feedback applications and stimulus presentation in neuroscience. *Frontiers in Neuroscience*, *4*, 179.
- Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific α -band electroencephalography increases over occipital cortex. *Journal of Neuroscience*, *20*, RC63.
- Zumer, J. M., Scheeringa, R., Schoffelen, J.-M., Norris, D. G., & Jensen, O. (2014). Occipital alpha activity during stimulus processing gates the information flow to object-selective cortex. *PLoS Biology*, *12*, e1001965.