Isolating human brain functional connectivity associated with a specific cognitive process

Michael A. Silver*[a,b,c], Ayelet N. Landau[c,d], Thomas Z. Lauritzen[a,b,c,e],
William Prinzmetal[d], Lynn C. Robertson[a,c,d]

[a]Helen Wills Neuroscience Institute, University of California, Berkeley, CA 94720
[b]School of Optometry, University of California, Berkeley, CA 94720
[c]Henry H. Wheeler Jr. Brain Imaging Center, University of California, Berkeley CA 94720
[d]Department of Psychology, University of California, Berkeley, CA 94720
[e]Redwood Center for Theoretical Neuroscience, University of California, Berkeley CA 94720

ABSTRACT

The use of functional magnetic resonance imaging (fMRI) to measure functional connectivity among brain areas has the potential to identify neural networks associated with particular cognitive processes. However, fMRI signals are not a direct measure of neural activity but rather represent blood oxygenation level-dependent (BOLD) signals. Correlated BOLD signals between two brain regions are therefore a combination of neural, neurovascular, and vascular coupling. Here, we describe a procedure for isolating brain functional connectivity associated with a specific cognitive process. Coherency magnitude (measuring the strength of coupling between two time series) and phase (measuring the temporal latency differences between two time series) are computed during performance of a particular cognitive task and also for a control condition. Subtraction of the coherency magnitude and phase differences for the two conditions removes sources of correlated BOLD signals that do not modulate as a function of cognitive task, resulting in a more direct measure of functional connectivity associated with changes in neuronal activity. We present two applications of this task subtraction procedure, one to measure changes in strength of coupling associated with sustained visual spatial attention, and one to measure changes in temporal latencies between brain areas associated with voluntary visual spatial attention.

Keywords: fMRI, functional connectivity, coherency, visual attention, voluntary attention, visual cortex, occipital cortex, parietal cortex

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) has proven to be a powerful method for identifying correlates of perceptual and cognitive processes in the human brain. Typically, fMRI time series are simultaneously acquired from many brain regions, and the correlation of these time series with stimulus presentation or performance of a cognitive task is assessed for each region. However, the computations that underlie perceptual and cognitive processes are not confined to a single brain area. Rather, these phenomena arise from interactions among brain areas within functional networks.

Measurement of functional connectivity among fMRI time series from multiple brain areas provides a method for identifying and characterizing these networks. However, the time series obtained with fMRI do not directly represent changes in neural activity. Rather, they are measurements of blood oxygenation level-dependent (BOLD) processes. Differences in the relative concentrations of oxygenated and deoxygenated hemoglobin in blood vessels in the brain are responsible for the basic fMRI contrast. These changes in blood oxygenation are often related to neural activity[1,2], but there is a sequence of several biochemical steps that link changes in neural activity to changes in BOLD signals.

* masilver@berkeley.edu; phone 1 510 642-3130; fax 1 510 643-5109; http://argentum.ucbso.berkeley.edu

Therefore, fMRI signals (and coupling of these signals) associated with a particular perceptual or cognitive state are a mixture of neural, neurovascular, and vascular phenomena, and the relative contributions of these phenomena are often unclear. At the neural level, simple anatomical connectivity can result in functional coupling of fMRI time series in two brain areas that is independent of a particular perceptual or cognitive state. For example, if area X has a strong direct projection to area Y, whenever X is active, Y will be active, due to the direct projection. Although this direct projection may well be functionally important for a particular cognitive operation, this cannot be determined with confidence if functional connectivity is only assessed in that cognitive state. It is also possible that the coupling due to anatomical connectivity is present for a wide variety of tasks and plays no significant role in the cognitive process of interest.

At the vascular level, if two brain areas share a common vascular input, the resulting BOLD signals will be highly correlated simply due to this fact. In this case, the relative temporal latencies of the BOLD signals will exhibit significant synchronization, and this synchronization will be largely independent of the coupling of neuronal activity in the two areas. Additionally, neuronal activity and BOLD responses are related by the hemodynamic response function (HRF). The shape of the HRF indicates that a brief impulse of neural activity will produce a BOLD response that is delayed by approximately 5 seconds. The HRF acts as a low-pass temporal filter with a cutoff frequency of approximately 0.15 Hz. As a result, functional connectivity analyses of fMRI time series typically focus on temporal frequencies below 0.15 Hz. There are a number of physiological signals that also exhibit significant fluctuations at these low frequencies, including arterial carbon dioxide[3] and changes in respiratory rate[4]. These global fluctuations will influence BOLD signals across the entire brain, thereby producing artifactual coupling of fMRI time series that are unrelated to neural activity.

Finally, the shape of the HRF varies significantly across brain areas of individual subjects[5]. The precise shape of the HRF can often be directly measured in sensory or motor areas. In these areas, stimuli and/or tasks can be designed that result in brief impulses of neural activity, allowing measurement of the resulting HRF. However, in higher level areas of the brain, it is often not possible to generate these impulses of activity, and the shape of the HRF in these areas in a given subject are generally unknown. This presents significant problems for interpretation of relative temporal differences in fMRI responses obtained from a pair of brain areas. On the one hand, neural activity could be highly synchronized in the two areas, but differences in HRF shape result in a temporal latency difference in the measured fMRI time series. On the other hand, there could be genuine latency differences between the areas at the neuronal level, but the HRF shape differences obscure or even reverse the sign of these latency differences in the BOLD time series.

The difficulty of isolating the neural component of fMRI functional connectivity in a single condition provides motivation to utilize a method for measuring changes in functional connectivity that modulate as a function of perceptual or cognitive state. We employed coherency analysis to compute coherency magnitude[6] and phase[7] for pairs of fMRI time series. Coherency is conceptually similar to temporal correlation analysis, but it is computed in the spectral domain[8]. For a given temporal frequency x, coherency is defined as the mean of the cross spectrum of two time series at that frequency divided by the product of the mean power spectra of the individual time series at that frequency.

The absolute value (modulus) of coherency corresponds to the magnitude of coupling between the two time series at frequency x. This value ranges from 0 to 1, with 0 indicating no functional coupling between the two time series at frequency x, and 1 indicating maximal coupling (identical relative power at frequency x). The complex argument, or polar angle, of coherency defines the phase difference between the two time series for that frequency, in units of radians. Dividing this phase difference by x yields the delay between the two time series in units of time (taking positive or negative values, depending on the direction of the delay). Thus, coherency analysis allows independent measurement of the strength of coupling between two time series as well as the temporal differences between them.

The critical feature of our analysis involves computing coherency magnitude and phase values for two distinct cognitive states and then subtracting the resulting values. This produces coherency magnitude and phase differences that specifically reflect changes in brain functional connectivity associated with the difference between the cognitive states. An assumption underlying this approach is that many of the functional connectivity artifacts described above (simple anatomical connectivity, global fluctuations in physiological signals, and differences in HRF shape across brain areas) will not vary as a function of cognitive state. The contribution of these artifacts to functional connectivity of fMRI signals will therefore be equivalent in the two tasks, and the subtraction procedure effectively eliminates them from the coherency magnitude and phase difference measures. Careful selection of task parameters and demands allows the differences between the two tasks to be associated with a particular cognitive process. The coherency magnitude and phase differences then isolate the brain functional connectivity that is specific to that cognitive process.

2. CHANGES IN THE STRENGTH OF COUPLING BETWEEN BRAIN AREAS

Here we present an example of the task subtraction approach to identify changes in coupling between brain areas associated with sustained visual spatial attention in the absence of visual stimulation[9]. The task has been described in a previous publication[10] and is summarized in Figure 1. Briefly, subjects performed a target detection task with a variable-duration delay period. They extensively practiced this task before any fMRI experiments were conducted. This allowed each subject to become familiar with the size, shape, location, and appearance of the visual target. Each trial began with an auditory cue that signaled subjects to begin attending in anticipation of target presentation. This cue was followed by a delay (between 2 and 16 seconds, duration randomly selected for each trial) during which no visual stimuli were presented. The target was then presented on 50% of the trials. When it was shown, the target was always presented at the end of the delay period. As a result, subjects had to maintain attention throughout the delay period at the visual field locations corresponding to the target while maintaining fixation at the center of the attended region. Each trial concluded with an 18 second intertrial interval during which the subjects simply fixated the central point.

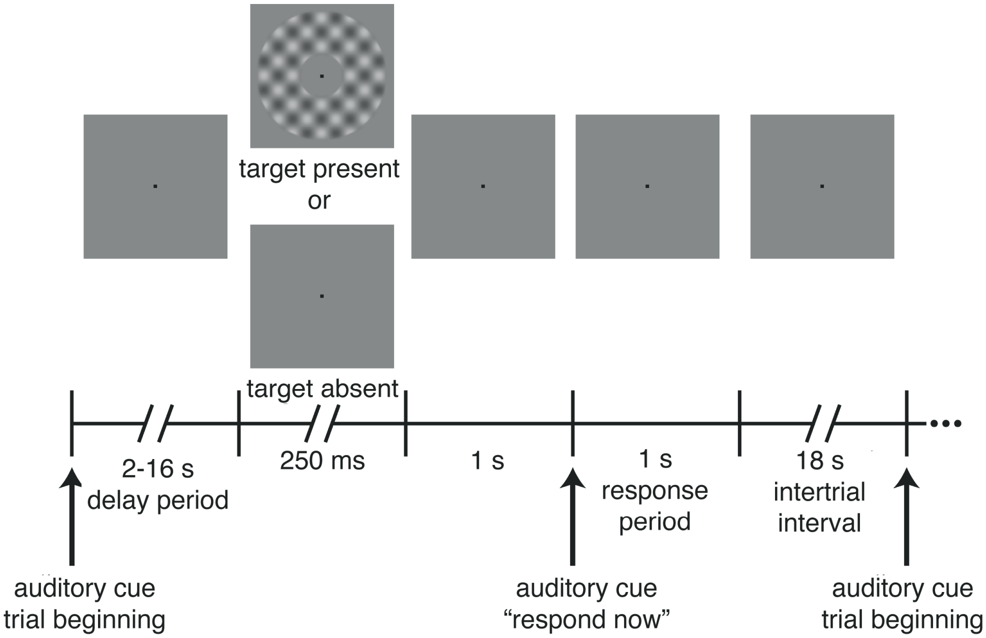

Figure 1. Sustained visual spatial attention task. Each trial began with an auditory cue followed by a delay period of variable duration. On half of the trials, a threshold-contrast target was presented at the end of the delay period. A second auditory cue instructed the subjects to respond whether they had seen the target. Each trial was followed by an 18-second fixation period. Reprinted with permission from Lauritzen et al. (2009) *Journal of Vision* 9(13):18, 1-14. Copyright held by The Association for Research in Vision and Ophthalmology.

In a previous publication[10], we described fMRI responses in visual cortex during performance of this task. Early visual cortical areas V1, V2, and V3 represent the first cortical stages in the visual processing hierarchy. Each of these areas contains a retinotopic map of the visual field. That is, stimuli presented to neighboring locations in the visual field will activate neighboring locations in the retinae. The representations of the visual field in areas V1, V2, and V3 preserve this retinotopic organization, with each cortical location within these areas representing a corresponding visual field location in the form of a continuous map of the visual field on the cortical surface.

In separate fMRI experiments, portions of the retinotopic maps of V1, V2, and V3 corresponding to the attended portion of the visual field in the sustained attention experiment were defined. Portions of these areas representing unattended locations in the peripheral visual field were also defined. For the regions representing attended visual field locations, fMRI signals were sustained for the duration of the delay period (corresponding to the period of time that subjects maintained attention)[10]. The regions representing unattended visual field locations exhibited a reduction in fMRI signal relative to the fixation baseline, and the duration of this negative BOLD response was also proportional to the duration of the delay period[10]. These findings demonstrate spatially-specific attention signals that are restricted to the portions of early visual cortical areas that represent attended visual field locations, even in the absence of visual stimulation.

These attention signals in early visual cortex presumably result from a volitional decision by the subject to allocate attention to the visual field locations where the target will appear. They therefore are unlikely to be generated by the early visual cortical areas themselves. Rather, they represent the “output” of the top-down visual spatial attention system. It is well known that allocating attention to a peripheral visual field location enhances perception of stimuli presented at that location[11,12], and this facilitation of perception by spatial attention may occur through changes in activity in early visual cortical neurons representing the attended visual field locations.

While the sources of these top-down attention signals are not fully characterized, other experiments suggest a role for IPS1 and IPS2, cortical areas in the intraparietal sulcus[13,14]. Like early visual areas, IPS1 and IPS2 contain topographic maps of visual field locations on their cortical surface. However, unlike early visual areas, IPS1 and IPS2 respond poorly to passive viewing of visual stimuli (i.e., unattended stimuli). Rather, activity in IPS1 and IPS2 seems to reflect the location of the “spotlight” of spatial attention, as a task that causes this spotlight of attention to systematically traverse the visual field results in traveling waves of activity across the topographic visual field maps in these areas[13]. The fact that these areas contain maps of top-down spatial attention raises the possibility that they transmit spatially-specific attention signals to early visual cortical areas that result in enhanced perception of visual stimuli presented at attended locations.

To test this hypothesis, we performed coherency analysis on fMRI time series obtained during performance of the sustained attention task. We defined visual cortical areas V1, V2, and V3, topographic parietal attention areas IPS1 and IPS2, and additional topographic visual cortical areas V3A, V3B, and V7 in the brain of each participant. Each of these areas was then restricted to the portion of the visual field map that represented the attended visual field locations.

Segments of the fMRI time series from the intertrial (fixation) period were excised and then concatenated into a single time series for each cortical area. A separate set of time series was generated by excising segments from the sustained attention portions of the trials. Coherency analysis was then performed on each pairwise combination of these time series, separately for fixation (Figure 2A) and for sustained attention during fixation (Figure 2B). The temporal frequencies used in this analysis ranged from 0.0625 to 0.15 Hz. The coherency magnitudes for the fixation time series provide a baseline measurement of functional connectivity, including contributions of anatomical connectivity, common vascular inputs, global fluctuations, etc. These artifacts are likely to be equally present in the attention coherency magnitude values (note that fixation was also maintained during the periods of sustained attention). Indeed, the coherency magnitude matrices in Figure 2A and 2B show very similar patterns. Specifically, cortical areas that are in close spatial proximity tend to have the highest amount of functional coupling of the fMRI time series. The spatial arrangement of these areas in a representative subject is shown in Figure 3.

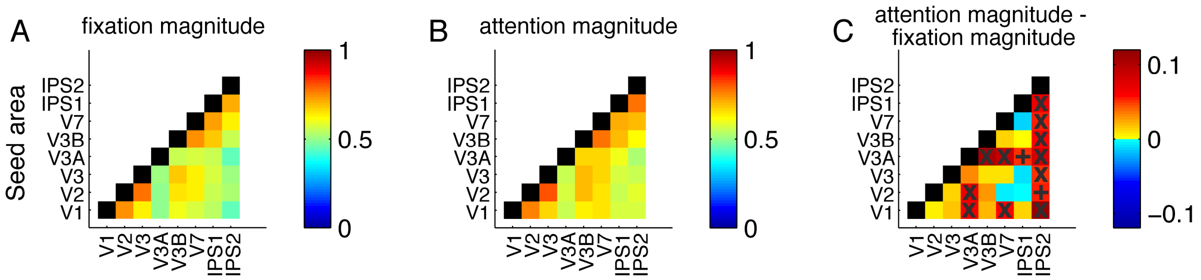

Figure 2. Coherency magnitude matrices for fixation and for sustained attention during fixation. Left column, functional connectivity associated with fixation. Middle column, functional connectivity associated with sustained attention during a period of fixation. Right column, differences in values between the first two columns, i.e., attention-specific connectivity. X, $p<0.01$; +, $p<0.05$ (corrected for multiple comparisons). Reprinted with permission from Lauritzen et al. (2009) *Journal of Vision* 9(13):18, 1-14. Copyright held by The Association for Research in Vision and Ophthalmology.

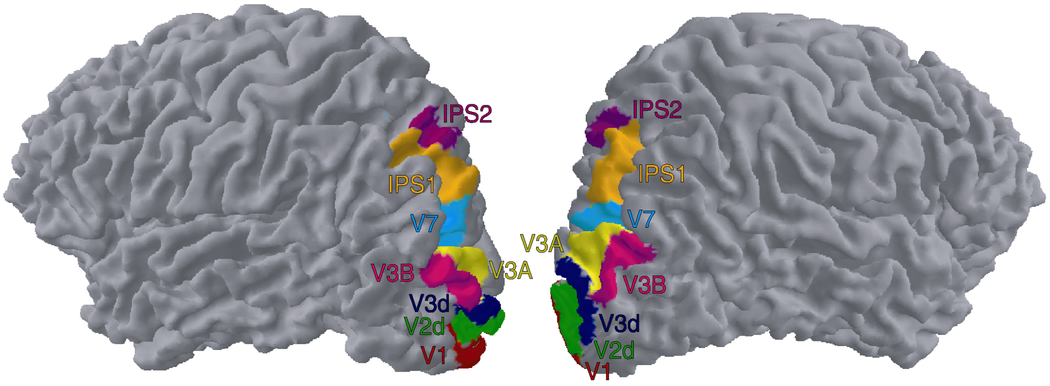

Figure 3. Spatial arrangement of areas V1 through IPS2 on the surface of the left and right hemispheres of a single representative subject.

Subtracting the fixation coherency magnitude values from those obtained during sustained attention while fixating results in coherency magnitude differences that are specific to top-down visual spatial attention (Figure 2C). The pattern of coherency magnitude differences is strikingly different from the corresponding pattern of coherency magnitudes obtained in a single cognitive condition (Figure 2A and 2B). Attention-specific coherency magnitude differences were detected for several pairs of cortical areas, including between the IPS areas and visual cortical areas, indicating greater coupling during sustained visual spatial attention relative to fixation.

These results demonstrate a pattern of attention-specific functional connectivity that would have been largely obscured if the analysis had only been conducted for fMRI responses measured during periods of sustained attention (Figure 2B). By including measurements from a fixation baseline condition that did not have explicit demands on attention, the differences between sustained attention and fixation coherency magnitudes could be computed. These differences reveal the component of functional connectivity that modulates as a function of task (in this case, engagement of the top-down attention system), removing several measurement artifacts that are common to both tasks. This method could easily be applied to a variety of task designs in which two conditions differ in a single cognitive operation. We now describe its application to study brain networks associated with the orienting of voluntary spatial attention.

3. CHANGES IN THE TEMPORAL LATENCIES BETWEEN BRAIN AREAS

Voluntary and involuntary attention represent two fundamentally different mechanisms for selecting visual locations for enhanced processing of stimuli at those locations[15]. Voluntary visual spatial attention can be allocated in a goal-oriented manner to locations containing behaviorally-relevant information. In contrast, involuntary attention is captured by salient events occurring at unexpected locations. These two forms of attention (voluntary and involuntary, or endogenous and exogenous) have been found to have different consequences for the processing of visual stimuli[16,17]. However, little is known regarding the functional connectivity of the brain networks that subserve voluntary and involuntary attention.

We employed a spatial cueing task[18] in which subjects reported the gender of a target face stimulus (Figure 4). Each trial began with a cue (red line segment) that appeared either to the left or right of the fixation point. The target face stimulus was then presented in either the left or right side of the display (no face was presented on 20% of the trials). Subjects responded as quickly as possible with a button press to indicate a male face target, female face target, or no target. Subjects were faster to respond (shorter reaction times, or RTs) when the cue and target appeared on the same side (cued trials) compared to when the cue and target appeared in opposite locations (uncued trials).

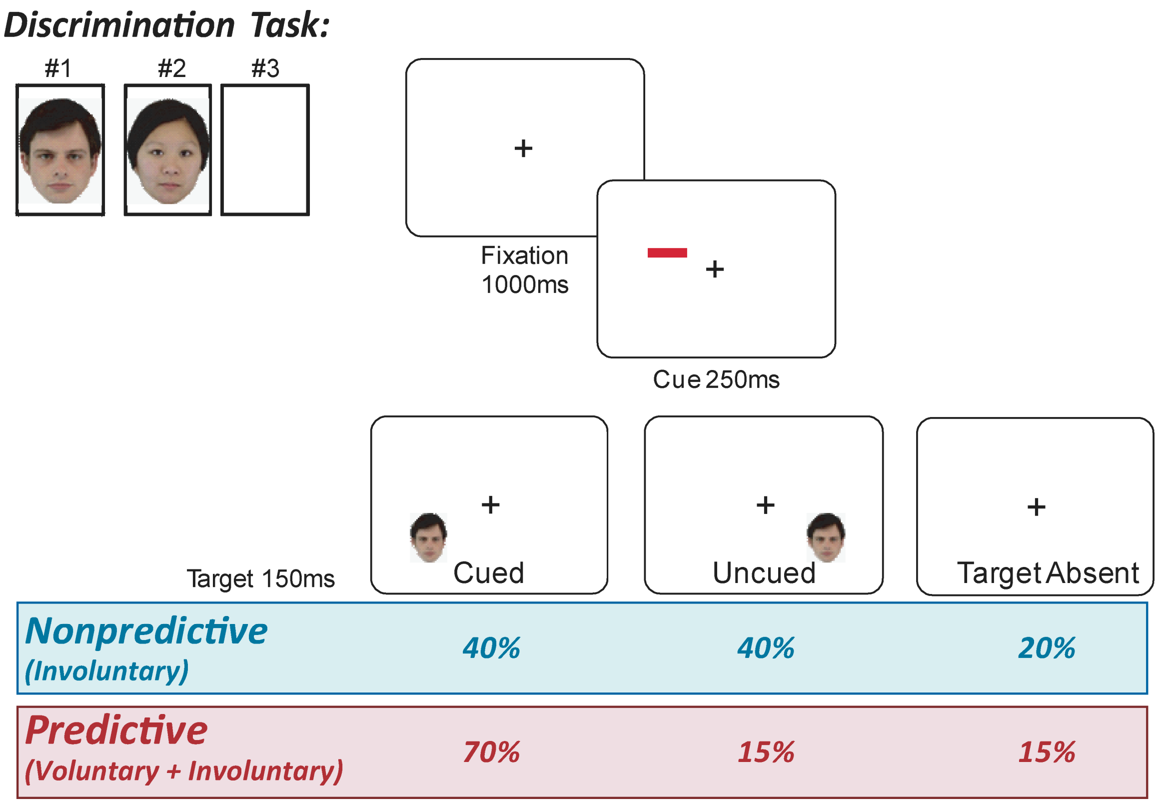

Figure 4. Spatial cueing face discrimination task. Each trial began with a 1000 ms fixation period followed by a cue that was presented either on the left or right side of the display. After a 250 ms interval, a target face stimulus was presented on either the left or right side (target-present trials), or no face was presented (target-absent trials). On target-present trials, the face appeared either on the same side as the cue (cued trials) or on the opposite side (uncued trials). In the nonpredictive condition, the face appeared on either the left or right side with equal probability, irrespective of cue location. In the predictive condition, the face appeared on the same side as the cue for 70% of trials.

Two blocked conditions allowed measurement of the effects of voluntary and involuntary attention on RTs to target presentation and analysis of functional connectivity associated with these two types of attention. In the nonpredictive condition, the location of the cue was random with respect to the location of the subsequent target. That is, the target face was equally likely to be presented in the left or right location, regardless of cue location. For these trials, the cue is a salient visual stimulus which captures involuntary attention, thereby reducing RTs for cued trials compared to uncued trials. In the predictive condition, presentation of the cue in a given location was followed by target presentation at the same location on 70% of trials (cued trials). For 15% of trials in the predictive condition, the cue and target appeared in opposite locations (uncued trials). Like the nonpredictive condition, the predictive condition results in capture of involuntary attention by the cue. However, in this case, because the cue also contains information regarding the location of the subsequent target face, subjects could guide voluntary attention to the cued location, as this was the most likely target location.

Critically, individual trials in the nonpredictive and predictive conditions contained identical cue and target stimuli, and the timing of the presentation of cue and target was also the same in the two conditions. The only difference between the conditions is that the cue carried information regarding target location in the predictive trials. Thus, the nonpredictive condition engages involuntary attention mechanisms, while the predictive condition engages both involuntary and voluntary attention. Nonpredictive and predictive conditions were blocked such that continuous fMRI time series were acquired in either the nonpredictive or predictive condition. Subjects were informed prior to the onset of fMRI recording whether the next block would contain predictive or nonpredictive cues, so they performed the task with the knowledge that the cue either did or did not provide information regarding target face location.

We report results on temporal differences in activity in two functionally identified brain regions. The first is area V1, or primary visual cortex. This is the earliest processing stage in the cortical visual pathways and is the recipient of top-down visual spatial attention signals[9,10,19,20]. The other area is IPS2, a region in the intraparietal sulcus that contains a topographic map of spatial attention signals and is thought to transmit spatially-specific top-down attention signals to early visual cortical areas[9,13].

Coherency between V1 and IPS2 was computed separately for the nonpredictive condition (involuntary attention) and for the predictive condition (combination of involuntary and voluntary attention). Coherency analysis was performed for temporal frequencies between 0.06 and 0.14 Hz. Analysis of coherency phase values in the nonpredictive condition revealed that V1 led IPS2 by 117 msec. For the reasons outlined above, this measured temporal difference in the fMRI time series is a function of not only latency differences in neuronal responses but also global fluctuations in physiological signals, differences between brain areas in the shape of the HRF, the amount of common blood supply to two regions, and other factors. Coherency phase values for the predictive condition showed that IPS2 led V1 by 676 msec.

However, as for the coherency magnitude values in the sustained attention experiment described in Section 2, subtraction of coherency phase values in the two conditions isolates the difference in functional connectivity between two cognitive states. In this case, because the nonpredictive blocks engage only involuntary attention but the predictive blocks engage both involuntary and voluntary attention, subtraction of coherency phase values from the nonpredictive condition from those from the predictive condition reveals the temporal dynamics of functional connectivity associated with voluntary attention. This subtraction (676 msec – (-117 msec)) yields a coherency phase difference of 793 msec, with IPS2 leading V1. That is, the direction of flow of signals between IPS2 and V1 associated with orienting of voluntary visual spatial attention is top-down.

We performed a similar task subtraction of coherency phase values for the sustained attention experiment described in Section 2. Here, the coherency phase values indicated that V1 led IPS2 by 109 msec for periods of sustained attention during fixation and that V1 led IPS2 by 515 msec for periods of fixation in the absence of attention demands. Thus, based on analysis of functional connectivity of only the single condition fMRI time series, one would conclude that V1 activity precedes that in IPS2.

However, subtraction of the phase values from the single conditions isolates the functional connectivity that is specific to top-down sustained spatial attention and reveals that, for the difference in coherency phase values, IPS2 leads V1 by 406 ms. This information was not available from analysis of functional connectivity of the fMRI time series acquired during sustained attention alone, presumably because the latency differences were dominated by non-neural factors unrelated to visual spatial attention circuits in the brain. The correct temporal differences between sustained visual spatial attention signals in IPS2 and V1 was only obtained by comparing the coherency phase values for the attention and fixation conditions.

The sustained visual spatial attention experiment and the spatial cueing face discrimination experiment isolate distinct aspects of top-down visual spatial attention. The first provides information regarding sustained spatial attention signals in the absence of visual stimulation (expectation signals), while the second describes brain mechanisms associated with shifting of voluntary attention in response to a peripheral cue that predicts target location. Nevertheless, task subtraction of coherency phase values from these two studies reveals very similar temporal dynamics in top-down visual spatial attention signals between parietal and occipital cortex. In the sustained attention study, IPS2 led V1 by 406 ms, while in the spatial cueing study, IPS2 led V1 by 793 ms.

These findings demonstrate that the temporal differences between areas detectable with this method are on the order of hundreds of milliseconds. Given the sluggishness of the HRF and the relatively slow sampling rate of fMRI measurements (1.5 seconds in the sustained attention study; 1.6 seconds in the spatial cueing study), it is commonly assumed that subsecond temporal information is not available from fMRI time series. However, although the temporal resolution of the fMRI signal is low, its temporal precision is not. As a result, fMRI has been used to estimate changes in cortical activity of approximately 100 ms[21,22]. The hundreds of millisecond difference between IPS2 and V1 activity roughly corresponds to the time course of some voluntary spatial cueing effects, as measured behaviorally[23]. Although the temporal differences obtained with coherency phase analysis are not a direct measure of temporal differences in neuronal activity, the task subtraction procedure removes any effects of differences in HRF shape between brain areas that do not modulate with task. Thus, it is possible that the temporal differences generated by this method are relevant for the time scales of neuronal interactions between brain areas.

4. CONCLUSIONS

The use of functional connectivity analysis of fMRI time series has tremendous potential for identification and characterization of the brain networks subserving specific perceptual and cognitive operations. However, because these analyses are conducted on blood oxygenation level-dependent signals, functional connectivity measures are subject to a number of potential artifacts, including common blood supply to nearby areas, global fluctuations in arterial carbon dioxide and changes in respiratory rate, and differences between brain areas in the shape of the HRF. These artifacts can have substantial effects on estimates of both the strength of coupling between two time series and on the measured temporal latency differences between the time series. We have described a task subtraction procedure to minimize the effects of these artifacts on functional connectivity data. In this procedure, functional connectivity is measured separately for two task conditions that differ by a single cognitive process. The functional connectivity values from one task are subtracted from the corresponding values in the other task, and any functional connectivity that does not modulate as a function of task is eliminated from the calculation. This procedure therefore allows measurement of the strength of coupling and temporal latency differences between two fMRI time series that is specific to a perceptual or cognitive process of interest.

REFERENCES

[1] Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann A., “Neurophysiological investigation of the basis of the fMRI signal,” Nature 412, 150-157 (2001).

[2] Logothetis, N. K. and Wandell, B.A., "Interpreting the BOLD signal," Annu. Rev. Physiol. 66, 735-769 (2004).

[3] Wise, R. G., Ide, K., Poulin, M. J., and Tracey, I., “Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal,” Neuroimage 21, 1652-1664 (2004).

[4] Birn, R. M., Diamond, J. B., Smith, M. A., and Bandettini, P. A., “Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI,” Neuroimage 31, 1536-1548 (2006).

[5] Handwerker, D. A., Ollinger, J. M., and D’Esposito, M., “Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses,” Neuroimage 21, 1639-1651 (2004).

[6] Sun, F. T., Miller, L. M., and D’Esposito, M., “Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data,” Neuroimage 21, 647-658 (2004).

[7] Sun, F. T., Miller, L. M., and D’Esposito, M., “Measuring temporal dynamics of functional networks using phase spectrum of fMRI data,” Neuroimage 28, 227-237 (2005).

[8] Rosenberg, J. R., Amjad, A. M., Breeze, P., Brillinger, D. R., and Halliday, D. M., “The Fourier approach to the identification of functional coupling between neuronal spike trains,” Prog. Biophys. Mol. Biol. 53, 1-31 (1989).

[9] Lauritzen, T. Z., D’Esposito, M., Heeger, D. J., and Silver, M. A., “Top-down flow of visual spatial attention signals from parietal to occipital cortex,” J. Vision 9(13):18, 1-14 (2009).

[10] Silver, M. A., Ress, D., and Heeger, D. J., “Neural correlates of sustained spatial attention in human early visual cortex,” J. Neurophysiol. 97, 229-237 (2007).

[11] Bashinski, H. S. and Bacharach, V. R., “Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations,” Percept. Psychophys. 28, 241-248 (1980).

[12] Posner, M. I., Snyder, C. R. R., and Davidson, B. J., “Attention and the detection of signals,” J. Exp. Psychol.: Gen. 109, 160-174 (1980).

[13] Silver, M. A., Ress, D., and Heeger, D. J., “Topographic maps of visual spatial attention in human parietal cortex,” J. Neurophysiol. 94, 1358-1371 (2005).

[14] Silver, M. A. and Kastner, S., “Topographic maps in human frontal and parietal cortex,” Trends Cogn. Sci. 13, 488-495 (2009).

[15] Prinzmetal, W. and Landau, A. N., “Dissecting spatial visual attention,” In [Tutorials in Visual Cognition], Coltheart, V., Editor, Psychology Press, Hove, United Kingdom, 43-66 (2008).

[16] Prinzmetal, W., McCool, C., and Park, S., “Attention: reaction time and accuracy reveal different mechanisms,” J. Exp. Psychol.: Gen. 134, 73-92 (2005).

[17] Prinzmetal, W., Zvinyatskovskiy, A., Gutierrez, P., and Dilem, L., “Voluntary and involuntary attention have different consequences: the effect of perceptual difficulty,” Q. J. Exp. Psychol. 62, 352-369 (2009).

[18] Posner, M. I., “Orienting of attention,” Q. J. Exp. Psychol. 32, 3-25 (1980).

[19] Gandhi, S. P., Heeger, D. J., and Boynton, G. M., “Spatial attention affects brain activity in human primary visual cortex,” Proc. Natl. Acad. Sci. USA 96, 3314-3319 (1999).

[20] Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., and Ungerleider, L. G., “Increased activity in human visual cortex during directed attention in the absence of visual stimulation,” Neuron 22, 751-761 (1999).

[21] Menon, R. S., Luknowsky, D. C., and Gati, J. S., “Mental chronometry using latency-resolved functional MRI,” Proc. Natl. Acad. Sci. USA 95, 10902-10907 (1998).

[22] Formisano, E. and Goebel, R., “Tracking cognitive processes with functional MRI mental chronometry,” Curr. Opin. Neurobiol. 13, 174-181 (2003).

[23] Posner, M. I., Cohen, Y., and Rafal, R. D., “Neural systems control of spatial orienting,” Philos. Trans. R. Soc. Lond. B Biol. Sci. 298, 187-198 (1982).