Published in final edited form as:
Neuroimage. 2012 August 15; 62(2): 1185–1189. doi:10.1016/j.neuroimage.2012.01.007.

Development of orthogonal task designs in fMRI studies of higher cognition: The NIMH experience

Susan M. Courtney
Department of Psychological & Brain Sciences, Johns Hopkins University, 204 Ames Hall, 3400 N. Charles Street, Baltimore, MD 21218, 410-516-8894

Susan M. Courtney: Courtney@jhu.edu

Abstract

This paper chronicles one researcher's journey at the National Institute of Mental Health, exploring ways to understand the neural systems responsible for the cognitive sub-processes of working memory tasks. Both the opportunities and the pitfalls with applying the idea of cognitive subtraction to neuroimaging data were well-known from studies using positron emission tomography. We took advantage of the improved temporal resolution of fMRI with a delayed-recognition task and identified the time-courses of the different stages of the task (encoding, memory delay, and recognition test) as predictor variables in a multiple regression analysis. Because these signals were temporally independent, individual components of tasks could be contrasted with one another, rather than entire tasks, reducing the problem of violations of pure insertion in cognitive subtraction. This approach enabled us to draw more detailed conclusions about the neural systems of higher cognition and the organization of prefrontal cortex than had been possible before fMRI. Further enhancements and innovations over the last 20 years by a multitude of researchers across the field have greatly expanded this knowledge, but this approach called "orthogonal task design" has remained a fundamental component of many of these modern studies.

Keywords

prefrontal cortex; working memory; fMRI; task design; data analysis; history

Motivation and exploration

I'll admit to being a bit naïve as a new postdoc entering the field of functional MRI soon after its birth. In 1992, as I was finishing my Ph.D. doing computational modeling of color perception with Leif Finkel and Gershon Buchsbaum, I saw Leslie Ungerleider give a talk at the annual meeting of the Association for Research in Vision and Ophthalmology, about a series of studies (Haxby et al., 1991, 1994) she did with Jim Haxby at intramural NIH demonstrating the existence of the "what" and "where" visual pathways in humans using positron emission tomography (PET). She presented the background, motivation, design and results extremely clearly and the neuroimaging results in humans fit nicely with what was

---

Correspondence to: Susan M. Courtney, `Courtney@jhu.edu`.

known about visual processing and attention in nonhuman primates. I was hooked. I wrote a letter to Leslie asking to do a postdoc with her. As I recall, her response was essentially "Why do you want to do a postdoc with me? You should stick with computational modeling. Besides, I'm not planning to do much more of this neuroimaging stuff…." Nevertheless, she kindly introduced me to Jim Haxby, who had been doing PET studies with Stanley Rappoport in the National Institute on Aging and was about to start his own independent lab section in the National Institute of Mental Health. Jim shared my enthusiasm for neuroimaging in general, and for this new thing called fMRI in particular. I was frustrated with modeling perception from a purely bottom-up perspective, but felt I didn't have enough data to constrain biologically-based models of higher cognitive influences. I wanted to study the neural systems of attention and memory experimentally. However, at that time I knew little about cognitive neuroscience, and less about neuroimaging. Amazingly, Jim and Leslie took me on anyway.

When I joined this new lab at the NIMH in 1993, the bread-and-butter was in the relatively proven technique of PET while we worked on figuring out how to do fMRI data collection and analysis. We worked during regular business hours collecting PET data, and on Friday evenings the entire lab would descend on the clinical MRI scanner. (The pizza and Chinese food delivery people got to know us well.) For each study, we had to manually plug in the high-speed gradients and reboot the console to run the fMRI software, crossing our fingers that the system wouldn't crash in the process. We also had to set up the screen and bring in the video projector in a metal shielding box on top of a wheeled cart, being careful not to get it too close to the bore of the magnet! Despite such practical and safety issues that are rare now, our biggest challenge was something that is still one of the main challenges in fMRI research: how to design a set of cognitive tasks and analyze the data in order to learn something new about the brain and cognition.

The starting points were Donders' method of "cognitive subtraction" (1868) and Sternberg's "additive factors" (1969), which had already been applied with some success to PET experimental design. Donders' method involved subtracting the reaction time for a task with cognitive processes A and B from the reaction time for a task with cognitive processes A, B, and C, in order to deduce the time required for process C. The method supposes that there are separate, cognitive processes that happen serially over time, between stimulus and response, and that each one is temporally independent of the others. The speed at which each of those processes, such as "detect a tone" and "discriminate whether it is high or low pitched," is presumed to occur would preclude being able to observe those processes separately in time with either O15 PET or fMRI. Therefore, the application of cognitive subtraction to neuroimaging data does not measure the absolute temporal duration of these cognitive processes, but rather the relative amount of neural activity associated with each process, integrated across both space (millimeters) and time (seconds for fMRI, minutes for PET).

Although the temporal resolutions of both fMRI and PET are too crude to resolve a series of cognitive processes when the duration of each is on the order of milliseconds, the temporal resolution of fMRI is much better than that of PET. With O15 PET, two or more tasks are compared, and data for each task must be collected in separate runs. One task is performed over the course of 1–4 minutes. Then it is necessary to wait a few minutes to allow the radiation to reach background levels before a different task can be performed. The standard analysis was, and to a large degree still is, essentially a t-test at each voxel, comparing data collected during performance of task 1 to data collected during task 2. Because signals were integrated across time, but fairly well localized to particular brain areas, cognitive subtraction in PET essentially assumed the spatial, rather than temporal, independence of cognitive processes within a task, each one localized to a different brain area. As with

Donders' original method, the PET signal related to all processes combined within one task had to be contrasted to the signal related to all processes combined in the other task. Therefore, the tasks had to be well-matched on every aspect, including sensory stimuli, motor actions, etc., etc., except the one particular process in which the researcher was interested.

We attempted to do this in my first PET study (Courtney et al., 1996), in which we had three tasks: a delayed-recognition task for faces, a delayed-recognition task for locations, and a sensorimotor control task. The two working memory tasks had exactly the same stimuli and motor response requirements, and presumably the same cognitive processes. The difference was that the instruction cue indicated that participants were to remember a different feature of the sample stimuli, the identity of the faces versus where they were presented on the screen. The sensorimotor control task had similar, but not identical, stimuli and motor requirements with the same timing. In the control task participants were told they did not need to remember anything. We then subtracted PET radiation counts collected during the sensorimotor control task from those during each of the memory tasks to identify areas generally involved in working memory. We also contrasted the two memory tasks directly to each other to identify regions differentially involved in selectively attending to, encoding, and maintaining the two different types of information. We found double dissociations in the activations related to the spatial versus nonspatial tasks in both sensory areas and prefrontal cortex.

These results indicated that, even with such carefully matched tasks, we hadn't really isolated any individual cognitive processes. Changing the type of information to be remembered changes selective attention and encoding in addition to memory maintenance, and because a single brain image was collected across all phases of the task (e.g. preparation, encoding, maintenance, and response) we couldn't separately observe these effects. The better temporal resolution of fMRI over PET promised to make this separation possible. Of course, simply knowing that this should be possible did not point directly to the best way for us to achieve our particular goal.

At this point we were still very much in an exploratory mode. I was concerned with whether we were even going to be able to observe different time courses for stimulus processing versus for memory maintenance at all. I started off by doing a simple t-test of all data collected during a face delayed-recognition task versus a control task, like the PET analysis, and just looked at the time courses of voxels within the regions that showed up in that contrast. Because many other groups had demonstrated responses to visual stimuli by then (e.g. Engel et al., 1994; Malach et al., 1995; Puce et al., 1996; DeYoe et al., 1996), I focused my attention on looking for memory delay activity in prefrontal cortex. Since I didn't know any better, I was looking at individual subjects' data and expecting to be able to see a time course as clean as those in V1 from flashing checkerboards. Instead, of course, the time courses looked terribly noisy. I thought there was something wrong with the data.

I kept working, collecting data from additional subjects and adjusting the preprocessing, and at this time in the history of fMRI, adjusting the preprocessing meant writing code. There were no publicly available, user-friendly fMRI data analysis software packages. An early version of SPM existed, but it was developed only for PET. Our lab at the NIMH had its own suite of analysis modules called FIDAP (Functional Imaging Data Analysis Platform), developed primarily by Joe Maisog, with some small contributions from me. For example, that little phrase in modern fMRI methods sections, "corrected for slice-acquisition-time," that now means you included a single command line during preprocessing, back then meant several days of debate regarding the best way to do this and a week writing and debugging the code.

Needless to say, I was not making very rapid progress and Jim inquired what was going on. When I told him I was having trouble finding any convincing sustained memory delay activations, he said, “Did you look at the visual stimulus time courses?” to which I replied, “Oh…well…I guess I could do that….” I soon learned not to make assumptions about what the data should look like, especially in prefrontal cortex. The time courses in early visual areas and in face-selective areas of the fusiform gyrus were beautiful! Moreover, the responses in early visual areas were clearly transient, without any sustained activation during the memory delay. They looked completely different from the time courses of activation in prefrontal cortex. It was then that I recognized the forest in the midst of the noisy trees. There were clearly functional differences between prefrontal cortex and other brain areas that could be seen with fMRI. Even though the time courses in prefrontal cortex were noisy, they were clearly more sustained during the memory delays than were the time courses in visual areas, whose activation increases occurred primarily during the initial stimulus-encoding interval. With this new confidence I was ready to attempt to statistically test these hypotheses and quantify the sustained versus transient characteristics of the time courses across different brain areas.

The Hard Part: The interplay of data analysis, task design, and psychology

The key realization was that a series of ones and zeros representing the times during which a single task component occurred throughout the scanning run (the faces to be remembered, the control task stimuli, the delay in the memory task, and the delay in the control task), convolved with a model of the hemodynamic response that would smooth and delay the fMRI signal, could be considered a predictor variable (a “regressor”) in a multiple regression analysis. This analysis evaluates each one of a set of predictor variables regarding whether it explains a significant portion of the variance in the observed data. It also provides a beta coefficient for each regressor, estimating the size of the effect. Thus, not only could we test whether a task component significantly contributed to activation within a single brain area, we could compare the *relative* contributions of all the different components from one brain area to another. Thus, we were able to demonstrate (Courtney et al., 1997), for example, that sensory areas responded primarily during stimulus presentation, posterior prefrontal regions responded to relevant stimuli and then sustained that response during the memory delay, and anterior prefrontal regions responded almost exclusively during the memory delay with very little response to the stimuli. This collection of different characteristic activation dynamics across the neural system suggested some sort of posterior-anterior functional organization within prefrontal cortex, which has been the focus of much recent research (e.g. Koechlin et al., 2003; Badre and D’Esposito, 2009).

While the results seemed robust and have since stood the test of time, doubts arose regarding whether we had completely separated working memory encoding from maintenance. In order to properly solve the system of equations in the multiple regression analysis, the regressors (the predictor variable time-courses) needed to be as orthogonal to each other as possible. Complete orthogonality means that the dot product of each of the pairs of vectors representing the time courses of each of your task components is equal to zero. A dot product of zero indicates that the two time courses are completely independent of each other. If they are not independent of each other, then a unique solution cannot be attained and it is impossible to unambiguously evaluate how to partition the data variance as being explained by one variable versus the other. This is the origin of the term “orthogonal task design” in fMRI.

The regressors for our original fMRI task design were not entirely orthogonal. We had a regular task sequence: 2 trials of memory task, 2 trials of control task, 2 trials of memory task again, and so on. The memory delay length and intertrial interval were always the same fixed amount of time. The stimulus to be encoded was always followed by the memory delay and the memory delay was always followed by the test stimulus. The time courses of these task components could not themselves be directly used as regressors for the fMRI data; we had to take into account that the fMRI activation changes due to the task components were smoothed and delayed in time due to the hemodynamic response. So, to create the regressors we had to convolve this regular task structure with a model of the hemodynamic response, resulting in partial colinearity. More sophisticated task designs were needed.

This imposition of independence in the time series of the variables driving the fMRI signal was also the idea behind what has been called “randomized event-related design”, which was being developed at about the same time (e.g. Clark et al., 1998; Dale and Buckner, 1997). Most of those early studies involved randomized presentation of multiple types of stimuli. There were still plenty of unanswered questions about the best way to make valid statistical inferences regarding a single stimulus-evoked response in an fMRI time series (Friston et al., 1995), and regarding comparisons between different stimulus event types. However, we impatiently jumped ahead to try to make multiple simultaneous inferences regarding multiple cognitive processing stages of a complex task. This impatience was typical for me, though. When I was working on my dissertation one of my committee members said, “You can’t model V4. We don’t even understand the retina!” Now, I wanted to understand prefrontal cortex. When the “events” of an event-related design are not stimuli, however, but rather are the components of a complex task, imposing an orthogonal temporal order can have greater consequences. We had to keep in mind what statistical contrasts would appropriately test our hypotheses and how changing the temporal order of task components might change the cognitive processes involved in the task.

With many randomized event-related designs, different types of stimulus events are directly contrasted with each other. When the events are task components, however, not all events are appropriate to contrast to any other event. In that first fMRI study (Courtney et al., 1997) described earlier, in which we had two tasks, a delayed-recognition working memory task, and a sensorimotor control task, we had “events” in which there was a stimulus (sample/test stimuli and control stimuli) and “events” in which there was not a stimulus (memory delays and control delays). Using multiple regression we could simultaneously estimate the magnitude of the BOLD response for each of these 4 events and contrast memory stimuli to control stimuli, and memory delays to control delays. This created a set of more well-matched cognitive subtractions, rather than lumping together all of the processes in the memory task and all of the processes in the control task. Events with stimuli were contrasted to other events with stimuli and events without stimuli were contrasted to other events without stimuli.

This approach of separating out different task components so that contrasts can be done across comparable, more well-matched components, rather than across entire tasks, attempts to reduce the problem of “pure insertion”. Both Donders’ original method and the neuroimaging version assume “pure insertion”, that one can add a cognitive process to a task without changing the other cognitive processes involved in the task. Sternberg (1969) attempted to eliminate this dubious assumption with his additive factors approach, parametrically manipulating a cognitive process rather than adding an entirely new process.

Manipulating only a single process in a task turns out to be extremely difficult in practice, however. Take, for example, an increase in memory load for a working memory task. In a Sternberg Item Recognition Task an increase in memory load will increase stimulus presentation or at least stimulus encoding, and memory scanning during the probe, as well as increasing memory representational load and maintenance processes. The popular n-back task attempts to manipulate memory load without changing stimulus presentation or probe response demands, but if one must compare all components of a 1-back task to all components of a 2-back task, one sees the effects of not just increased memory load for items, but also increased memory load for temporal order, and increased updating and selection of relevant versus irrelevant maintained items. Creating temporal structure in the fMRI data collected over the course of task performance was a key to beginning to address this problem.

At the same time that we were analyzing data dependent on the temporal structure of the components of a delayed-recognition task, Jonathan Cohen and his colleagues (Cohen et al., 1997) were using the n-back task, but with a long interstimulus memory delay inserted and an ANOVA analysis with memory load and time as factors. Their hypothesis was that areas with a main effect of time (the series of data points at each TR following the presentation of a stimulus) would be related to encoding, updating, responding, etc., and all of the other processes that occurred at the time of stimulus presentation but not during the delay. These areas would have a higher fMRI signal during the time points immediately following the stimulus presentation, but the signal would decline over time during the delay that followed. In contrast, areas that showed a main effect of memory load, but not an effect of time were proposed to reflect neural activity involved in working memory maintenance: high during the encoding of the stimulus, but then maintaining that high signal across the entire delay period, with greater activation elicited by a higher working memory load. Because of the long interstimulus interval, they could thus separately identify areas that showed load effects during encoding from those that showed load effects during memory maintenance. As luck would have it, both papers on our respective findings were submitted at about the same time to the same journal, so the two papers were ultimately published back-to-back in the same issue of *Nature.* Moreover, while the approaches and findings were similar, they were not identical, and their juxtaposition highlighted the advantages and disadvantages of each method.

Of course, adding temporal structure doesn’t eliminate the problem of pure insertion – adding a process can still alter other processes in the task even if they occur at different times. If the fMRI signals corresponding to those processes are temporally independent, however, then the changes in each process may be separately observable rather than being mixed together. For example, the effects of increasing memory load or changing the type of information to be remembered on encoding processes might be observable separately from the effects on maintenance processes during the delay. This was the goal of imposing additional temporal structure in the tasks in order to take advantage of the improved temporal resolution of fMRI.

Violations of pure insertion when adding a process to a task isn’t the only way that other task processes may be inadvertently changed. Imposing orthogonal temporal structure to the components of a task can also change the processes involved in that task. This is not necessarily a problem for all event-related designs. Changing the temporal order of the presentation of stimuli in a passive viewing task, for example, is not likely to result in changes to the neural activation in response to those stimuli. Event-related designs directly comparing equivalent stimulus types had, from the beginning, emphasized the need for counter-balanced event order, such that an event was equally likely to be followed by any of the possible event types. These kinds of studies also emphasized the need for “jittering,” the inclusion of a variable interstimulus interval, which essentially amounts to the inclusion of “blank” or “rest” events mixed into the sequence with the events of interest (Dale, 1999).

Many cognitive tasks, however, include components that typically proceed in a fixed order, such as the sample, delay, test sequence of a delayed-recognition task or the cue and detection phases of an attention task. In this case it is more difficult to design the experiment to create orthogonal time courses for the task components without changing the cognitive processes themselves or the strategies employed by the participants. One proposed solution (Ollinger et al., 2001) was to have some portion of the trials only contain the first of two task components. For example, if the task involved a cue indicating where to pay attention followed by a detection test, some trials (20–40%) would be “partial trials”, stopping after the presentation of the cue with no test event. The greater the number of partial trials relative to full task trials (up to 50%), the more orthogonal the time courses of the two task components will be, but the greater the risk that subjects cognitive processes and strategies will be different than if there were no partial trials. Similarly, it is possible to do a delayed recognition task with fully orthogonal sample, delay and test event timecourses by having each type of event equally likely to be followed by any other of the event types, with subjects instructed to respond to each test event based on the most recent sample stimulus presentation, even if there were multiple samples or multiple tests in a row, no memory delay or multiple delays (Ackerman & Courtney, under review). As with the partial trials method, however, there is a risk that the processes and strategies used by subjects may not be the same as with the standard order of the task components. So by imposing temporal orthogonality one might decrease the problem of pure insertion by being able to separately evaluate signals related to individual task components, but the problem might be exacerbated in other ways. As always, experimental design decisions are a trade-off.

What Orthogonal Task Design Alone Can’t Do

One problem with the assumption of pure insertion that orthogonal task design alone cannot fix is the issue of attention, motivation and task difficulty (e.g. Barch et al., 1997; D’Esposito et al., 1995; Postle et al., 2000). Adding a cognitive process, increasing memory load, or changing the temporal parameters of a task can change the task difficulty. It might increase uncertainty or make the task inherently more interesting. Thus, our bored experimental participant lying in the scanner may pay more attention to the stimuli, their responses, or the entire task. Enhanced overall vigilance can increase the gain of neural activation in general. Elevated selective attention can increase the gain for all task-relevant information during all temporal phases of the task: encoding, maintenance and response, for example. So, how does one separate increased activation due to the representation of a greater amount of information in working memory, or an added cognitive process, from increased activation due to greater attention to a harder task?

Our solution was to design a set of experimental conditions to test hypotheses that depended on double dissociations. Double dissociations are the ideal for answering many scientific questions. For example, if you lesion area X and your animals can do Task 1 (consisting of cognitive processes A and B), but cannot do Task 2 (consisting of processes A, B, and C), this is a single dissociation. You might conclude that area X is necessary for process C, and you would likely be right. However, what if your animal fails to perform any task that has any processes in addition to A+B? You cannot then conclude much of anything specific regarding area X and process C. However, let’s say you lesion area X and your animal can do a task involving A+B but has trouble with a task involving A+C. *Furthermore* you lesion area Y in a different animal and that animal can do a task involving A+C but has trouble doing a task involving A+B. Now, you can start to come to specific conclusions about the relationships among areas X and Y and processes A, B and C. In neuroimaging, if activation for Task 2 is greater than activation for Task 1, it might be simply because subjects are paying attention more in Task 2. However, if Task 2 produces greater activation in some brain regions, but Task 1 produces greater activation in a different set of brain regions, chances are reduced that the explanation lies solely in task difficulty.

The main constraints in demanding a double dissociation from an fMRI experiment in its early development era were the poor signal-to-noise and the lack of available scanner time. There was (and is) also a limit to the number of hours that a volunteer would tolerate doing a task over and over again in the uncomfortable scanner environment, even if the volunteer was also a co-investigator and author! Since a minimum number of data points per condition was necessary in order to have sufficient statistical power, and the total number of data points that could be collected in a single scanning session was limited, the number of task conditions that could be included in a session was also limited. After we were convinced that we could detect sustained activity related to working memory in prefrontal cortex, we set about trying to discover whether there was a double dissociation in prefrontal cortex for spatial and nonspatial information for working memory maintenance in humans similar to what Jim and Leslie had found for visual processing with PET. This double dissociation between visual streams is what I saw Leslie present that convinced me to do neuroimaging in the first place and there was recent evidence from nonhuman primates that there might be a similar dissociation in prefrontal cortex for working memory (Wilson et al., 1993). We had found a similar double dissociation with an earlier PET study (Courtney et al., 1996), but the poor temporal resolution did not allow us to know whether the dissociation was being driven by differences during stimulus encoding or by sustained working memory delay period activity. Furthermore, the spatial working memory area we identified in that study was very close to what had been previously identified as the frontal eye fields. Therefore, we needed the improved spatial resolution of fMRI with the ability to examine activations within individual subjects, as well as the improved temporal resolution (Clark et al., 1996).

Now knowing how noisy prefrontal fMRI signals could be, and hoping we could collect enough data in both the spatial and nonspatial memory tasks (as well as in the sensorimotor and eye-movement control tasks) to have sufficient statistical power, we decided that we would bring the same volunteers in for two separate, two-hour scanning sessions. Thank goodness for cooperative, motivated volunteers! We also spent a lot of time with behavioral pilot testing of the spatial and nonspatial tasks, titrating the various task parameters to make sure that performance was not at floor or ceiling and that performance was not significantly different between the two tasks. Keeping performance levels within these constraints was necessary so that any double dissociation in activation was not swamped by attention-related activation differences due to task difficulty. Our efforts were rewarded with a clear double dissociation: more activation in ventral frontal cortex for nonspatial working memory maintenance and more activation in dorsal frontal cortex for spatial working memory maintenance. Furthermore, analysis in single subjects revealed that the area preferentially activated during the spatial working memory delay (the posterior superior frontal sulcus) was anatomically distinct from the area showing saccade-related activity in the precentral sulcus (Courtney et al., 1998). Again, orthogonal task design played a big role in reaching these conclusions because it enabled separation of the transient activity during stimulus encoding, which overlapped with the saccade-related activity in the frontal eye field, from the sustained activity during the memory delay.

The Future of Orthogonal Task Designs in fMRI Research

Orthogonal task designs and multiple regression analysis of time series data have become staples in neuroimaging research. The field has become increasingly sophisticated over the last 20 years, however, regarding how the results of such experiments are interpreted. The temptation to assign tasks, or cognitive subprocesses, to brain areas with pretty, statistically thresholded pictures led to the criticism early on that fMRI was the new phrenology. Twenty years later the criticism is still voiced, but it is slowly fading due to improvements in neuroimaging experimental design, analysis, and interpretation. A few brain regions that do appear to be modules for the representation of particular kinds of information have been rigorously tested (Op de Beeck et al., 2008). Complex tasks and cognitive processes, however, are increasingly acknowledged to be accomplished by complex, dynamic interactions among multiple brain areas (see for example, Courtney 2004), not by individual modules. A key to this understanding has been the development of new, sophisticated techniques for data analysis. This includes the use of multivoxel pattern classification for probing information representation in the spatiotemporal activation of specialized neuronal populations (Mur et al., 2009). It also includes the large variety of approaches to evaluating functional and structural connectivity among brain regions, including network-network interactions (e.g. Bullmore and Sporns, 2009). Our interpretation of fMRI data has also been greatly enhanced by considering it in the context of other measures of and contributors to neural function, such as genetics and molecular processes, single-cell recordings and scalp EEG. Interpretation of all of this data, however, will continue to be primarily dependent on good experimental task design and, to this end, the inclusion of orthogonal temporal structure in the task components continues to play a leading role.

References

Ackerman, C.; Courtney, SM. Spatial Relations and Spatial Locations are Dissociated within Prefrontal and Parietal Cortex. Under review

Badre D, D'Esposito M. Is the rostro-caudal axis of the frontal lobe hierarchical? Nat Rev Neurosci. 2009; 10:659–669. [PubMed: 19672274]

Barch DM, Braver TS, Nystrom LE, Forman SD, Noll DC, Cohen JD. Dissociating working memory from task difficulty in human prefrontal cortex. Neuropsychologia. 1997; 35(10):1373–80. [PubMed: 9347483]

Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat Rev Neurosci. 2009; 10(3):186–98. [PubMed: 19190637]

Clark VP, Keil K, Maisog JM, Courtney SM, Ungerleider LG, Haxby JV. Functional magnetic resonance imaging (fMRI) of human visual cortex during face matching: A comparison with positron emission tomography. NeuroImage. 1996; 4:1–15. [PubMed: 9345493]

Clark VP, Maisog JM, Haxby JV. fMRI study of face perception and memory using random stimulus sequences. J Neurophysiol. 1998; 79(6):3257–65. [PubMed: 9636124]

Courtney SM, Ungerleider LG, Keil K, Haxby JV. Transient and Sustained Activity in a Distributed Neural System for Human Working Memory. Nature. 1997; 386:608–611. [PubMed: 9121584]

Courtney SM, Petit L, Ungerleider LG, Maisog JM, Haxby JV. An Area Specialized for Spatial Working Memory in Human Frontal Cortex. Science. 1998; 279:1347–1351. [PubMed: 9478894]

Courtney SM, Ungerleider LG, Keil K, Haxby JV. Object and Spatial Visual Working Memory Activate Separate Neural Systems in Human Cortex. Cereb Ctx. 1996; 6:39–49.

Courtney SM. Attention and cognitive control as emergent properties of information representation in working memory. Cognitive, Affective, & Behavioral Neuroscience. 2004; 4(4):501–516.

Dale AM. Optimal experimental design for event-related fMRI. Human Brain Mapping. 1999; 8:109–114. [PubMed: 10524601]

Dale AM, Buckner RL. Selective averaging of rapidly presented individual trials using fMRI. Hum Brain Mapp. 5(5):329–40. [PubMed: 20408237]

D'Esposito M, Detre JA, Alsop DC, Shin RK, Atlas S, Grossman M. The neural basis of the central executive system of working memory. Nature. 1995; 378(6554):279–81. [PubMed: 7477346]

DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, Neitz J. Mapping striate and extrastriate visual areas in human cerebral cortex. Proc Natl Acad Sci U S A. 1996; 93(6):2382–6. [PubMed: 8637882]

Donders, FC. On the speed of mental processes. In: Koster, WG., editor. Acta Psychologica 30, Attention and Performance II. North-Holland Publishing Company; Amsterdam: 1969. p. 412-431.Translated from the original, 1868

Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN. fMRI of human visual cortex. Nature. 1994; 369(6481):525. [PubMed: 8031403]

Friston KJ, Holmes AP, Poline JB, Grasby PJ, Williams SC, Frackowiak RS, Turner R. Analysis of fMRI time-series revisited. NeuroImage. 1995 Mar; 2(1):45–53. [PubMed: 9343589]

Haxby JV, Grady CL, Horwitz B, Ungerleider LG, Mishkin M, Carson RE, Herscovitch P, Schapiro MB, Rapoport SI. Dissociation of object and spatial visual processing pathways in human extrastriate cortex. Proc Natl Acad Sci U S A. 1991 Mar 1; 88(5):1621–5. [PubMed: 2000370]

Haxby JV, Horwitz B, Ungerleider LG, Maisog JM, Pietrini P, Grady CL. The functional organization of human extrastriate cortex: a PET-rCBF study of selective attention to faces and locations. J Neurosci. 1994; 14(11 Pt 1):6336–53. [PubMed: 7965040]

Koechlin E, Ody C, Kouneiher F. The architecture of cognitive control in the human prefrontal cortex. Science. 2003; 302(5648):1181–5. [PubMed: 14615530]

Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, Ledden PJ, Brady TJ, Rosen BR, Tootell RB. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. Proc Natl Acad Sci U S A. 1995; 92(18):8135–9. [PubMed: 7667258]

McCarthy G, Puce A, Constable RT, Krystal JH, Gore JC, Goldman-Rakic P. Activation of human prefrontal cortex during spatial and nonspatial working memory tasks measured by functional MRI. Cereb Cortex. 1996; 6(4):600–11. [PubMed: 8670685]

Mur M, Bandettini PA, Kriegeskorte N. Revealing representational content with pattern-information fMRI – an introductory guide. Social Cognitive and Affective Neuroscience. 2009; 4:101–109. [PubMed: 19151374]

Ollinger JM, Shulman GL, Corbetta M. Separating processes within a trial in event-related functional MRI: The Method. NeuroImage. 2001; 13:210–217. [PubMed: 11133323]

Op de Beeck HP, Haushofer J, Kanwisher NG. Interpreting fMRI data: maps, modules and dimensions. Nat Rev Neurosci. 2008; 9(2):123–35. [PubMed: 18200027]

Postle BR, Stern CE, Rosen BR, Corkin S. An fMRI investigation of cortical contributions to spatial and nonspatial visual working memory. Neuroimage. 2000; 11(5 Pt 1):409–23. [PubMed: 10806028]

Puce A, Allison T, Asgari M, Gore JC, McCarthy G. Differential sensitivity of human visual cortex to faces, letterstrings, and textures: a functional magnetic resonance imaging study. J Neurosci. 1996; 16(16):5205–15. [PubMed: 8756449]

Sternberg, S. The discovery of processing stages: Extensions of Donders' Method. In: Koster, WG., editor. Acta Psychologica 30, Attention and Performance II. North-Holland Publishing Company; Amsterdam: 1969. p. 276-315.