Neural correlates of temporality: Default mode variability and state-dependent temporal awareness

Abstract: The continual background awareness of passing time is an essential structure of consciousness, conferring temporal extension to the evanescent sensory present. Seeking the possible neural correlates of ubiquitous temporal awareness, this article reexamines fMRI data from off-task "default mode" (DM) periods in 25 healthy subjects studied by Grady et al. ("Age-related Changes in Brain Activity across the Adult Lifespan," JOCN 18(2), 2005). "Brain reading" using Support Vector Machines detected information specifying elapsed time, and further analysis specified distributed networks encoding implicit time. These networks fluctuate; none are continuously active during DM. However, the aggregate regions of greatest variability closely resemble the default mode network. It appears that the DM network has an important role as a state-dependent monitor of temporality.

Introduction: Timing and temporality

Time is important in human behavior, cognition, and consciousness. Yet the role of time is differently conceived by different disciplines. In cognitive science, time is a perceptual dimension exploited as needed in time-dependent tasks, such as temporal order and simultaneity judgments, or duration estimation and reproduction tasks. Time in these contexts is attended time as such, time in the foreground of perception and behavior – *timing*, in short. As such it is one prominent stimulus feature among many. In contrast, in philosophical phenomenology time is the foundation of every aspect of consciousness. It is a structural feature of consciousness, among "those properties that are common to most or all conscious experiences" (Seth, 2009). This is time as *temporality,* "the infrastructure of reality" (Zahavi, 1999). Its richest description can be found in the works of Husserl (esp. (Husserl, 1966 (1928)) whose account of temporal awareness has remained a centerpiece of phenomenology ever since. Although capacities for judging time and coordinating behavior in time are essential to human cognition, timing is nonetheless only a subset of phenomenal temporality. The latter concept, as developed in phenomenology, encompasses time as a fundamental dimension of all percepts and all behavior, a dimension which may or may not be a center of explicit attention.

We encounter temporality when background expectations of duration are violated. For example, delays in conversational turns are highly significant; apparent haste or tardiness in a reply is quickly read as an indicator of one's attitude toward the speaker or the exchange. This is common at various time scales, from short ("Do you love me?" coupled with "Of course," as compared to "[pause] Of course.") to long, as in, for example, perceived slights in tardy email responses. One's attention need not be cued to notice these timing anomalies. Thus, it seems that duration is continuously monitored at many scales in a host of human activities. The phenomenologists augment examples with foundational arguments. Husserl maintained that our fundamental distinction between changing and unchanging (enduring) objects was essentially a temporal distinction. Is a pendulum at the bottom of its arc moving or still? The distinction could not even be conceived without some structure in which the present impression of the pendulum was coordinated with its immediate past (and its anticipated future). These short-term "retentions" and "protentions" are accordingly necessary components of every moment of consciousness. We can readily

update Husserl's argument to bear on cognitive capacities in general, with the observation that the most useful model of the perceptual world is not so much a construct of objects in scenes as trajectories through space/time (Grush, 2004). (Mastery of trajectory estimation confers obvious adaptive advantages.)[1]

The explicit cognition of time is a lively area of research in cognitive neuroscience and functional neuroimaging. Temporality, in contrast, is infrequently studied (Grush, 2004; Lloyd, 2002; Yoshimi, 2007). There are several reasons contributing to this comparative neglect. First, since temporal awareness is continuous but non-sensory, it is the ambient background to explicit, articulated sensory experiences. Informal phenomenologies, of the sort that occupy brief chapters in books by philosophers in the analytic tradition, overlook temporality (along with other structural features of consciousness). This bias toward the sensory and atemporal present is also reflected in many aspects of 20th Century cognitive science. Marr, for example, classically regarded the problem of perception as the extraction of three-dimensional structure from the occurrent retinal image, and a huge research literature in psychophysics and perception operates "in the present tense," addressing perception of present objects and events (Marr, 1982). These interests have been readily translated into experiments in functional neuroimaging.

This comparative neglect is not merely a contingent preference, however. Temporality eludes experiment by its very nature. If the phenomenologists are right that temporality is a continuous aspect of awareness, it can't be turned on and off by experimental manipulation. Moreover, in any experiment in which time is a variable, subjects are likely to notice. Time, then, will move from background to foreground, and the general structural scaffold of temporality will yield to explicit timing. These obstacles are compounded by the constraints of neuroimaging, especially fMRI. "Raw" fMRI images reveal continuous metabolic activity nearly everywhere in the brain, overlaid with a great deal of noise. Unlike other medical and scientific imagery, single images reveal little or nothing. Interpretation thus often rests on contrasts (to detect small changes between all-over activation patterns) and averages (to smooth away some of the noise). Both are inimical to detecting temporality. If temporality is continuous, it is present on both sides of any contrast, and won't appear in a contrastive image. If, on the other hand, temporality involves change that is independent of the time course of conditions in an experiment, averaging will eliminate those changes along with the noise.

These considerations impose severe constraints on any scientific study of temporality, especially research based in neuroimaging. Nonetheless, this paper will re-examine data from one fMRI study as a case

[1] It's arguable that temporality is what makes the "hard problem of consciousness" seem so hard. In the standard formulation of the hard problem, one is asked to weigh the phenomenal sensation of red (or some other quale) against the current picture of visual cognition in the brain, usually presented as an elaborate cascade of hierarchical detections of stimulus properties (Chalmers, 1996). The resulting intuition is that all this visual processing is not enough to give us phenomenal red. I suggest that in part this intuition arises (if indeed it does arise) because all experience, qualia included, occurs in temporal contexts. Consciousness is temporally thick, but this thickness is not accounted for in hierarchical models of sensory awareness.

study in temporality (Grady, Springer, Hongwanishkul, McIntosh, & Winocur, 2006). We will consider several questions: 1) Can we demonstrate that specific images encode temporal information? 2) Can we identify regions of the brain that are particularly involved in encoding temporal information? 3) Can we begin to characterize mechanisms supporting the awareness of temporality?

II. Encoding temporality

Temporality may be implicit, but if it is an aspect of consciousness then temporal properties must affect the contents of consciousness from moment to moment. That entails (as Husserl realized) that all perception is in continuous flux. Even a static scene is always changing as its objects progressively extend their perceived duration. This in turn implies that the information encoding objects of perception is time-inflected. Even if no other stimulus property is changing, the brain activity that supports awareness must also support its temporal progression.

Lloyd (2002) reexamined five neuroimaging experiments involving several aspects of perception and cognition to confirm that (in 89% of subjects) continuous change is a feature of fMRI image series. That is, as the lag between images increased, the measured differences between images also increased, whether this is measured as correlation between images or as Euclidean distance. Although this is consistent with a Husserlian conception of temporality, other processes might give rise to continuous "drift" in image series. What is needed, therefore, is a way to test for the presence of explicitly temporal information.

Because temporality is implicit, ubiquitous, and continuous, it is not subject to experimental manipulation in normal, awake human subjects. Nonetheless, temporal information must be present in the brain and neuroimages may, possibly, detect it. Helpfully, if temporality is the continuous background for consciousness, it ought to be present in every conscious subject, regardless of the "official" task of interest. One might especially look for signs of temporality during moments when subjects are performing no task, and their environment is not changing, or at least is without markers of passing time. Time may be unavoidably noticed by subjects in the noisy and confined space of a scanner, but when there is no need to judge or reproduce durations, temporality should appear in awareness similarly to its appearances in other perceptual environments.

Functional MRI experiments typically contain many blocks off-task and without markers of time, often intervals when the baseline images are collected. Cognition during these "off-task" intervals has been termed "default mode" (DM) activity, following the discovery that, regardless of the target tasks, the off-task brain exhibits spontaneous activity in a network of regions that is quite stable from experiment to experiment (Buckner, Andrews-Hanna, & Schacter, 2008; Mason, et al., 2007; McKiernan, Kaufman, Kucera-Thompson, & Binder, 2003; Raichle, et al., 2001). The default network is generally depicted as a network engaged in continuous activity throughout off-task intervals, but this may be an artifact of averaging the images collected throughout the interval. Temporality, in contrast, implies an awareness of change, even in the absence of changing tasks or contexts. At a minimum, one is aware that time is passing, despite the recurring sounds and unchanging visual scene. Among the possible expressions of temporality, here we explore the encoding of elapsed time. That is, we ask whether brain images include information that specifies how much time has elapsed during the default interval. In an experiment where there are multiple off-task blocks, this can be generalized: Does the brain encode elapsed duration information *in the same way* at comparable time points in each DM interval?

The evidence bearing on this question will be indirect, using a "brain reading" method. Employing pattern recognition methods described below, we construct multivariate mathematical functions that can determine serial positions of unlabelled images taken from the DM blocks. The success of this machine learning strategy (compared to appropriate control conditions) provides evidence that the images indeed encode temporal information. Following this analysis a separate analysis can approach the question of how this information is encoded.

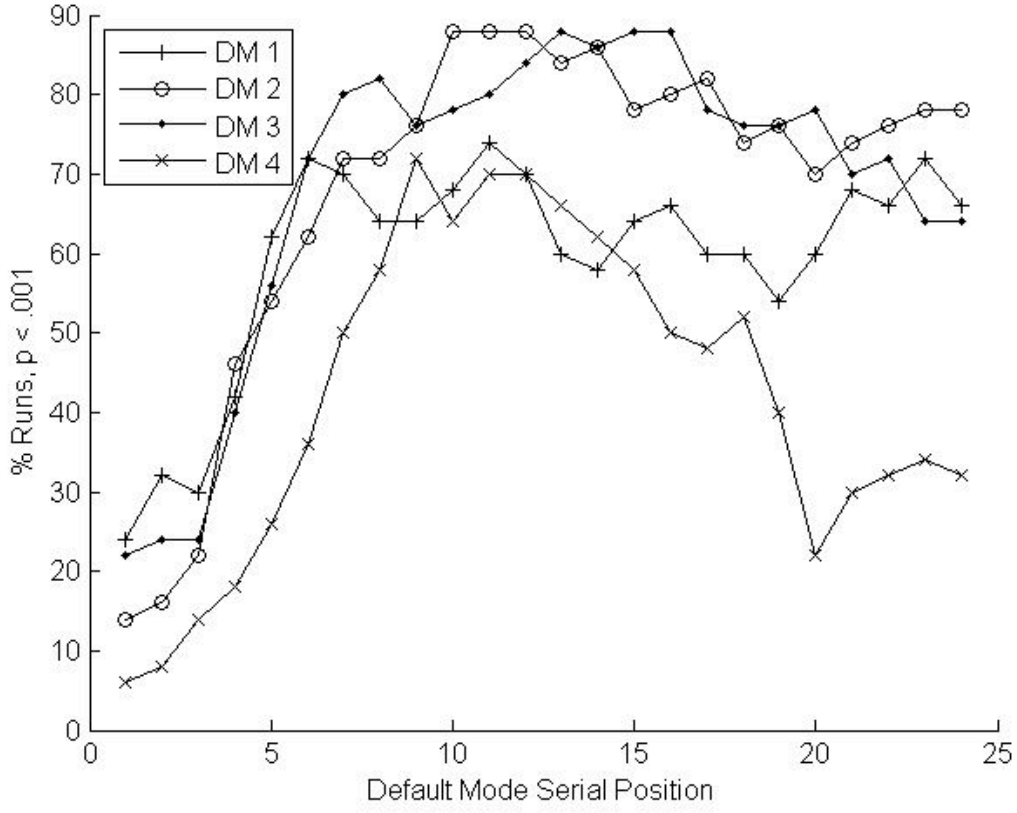
The present analysis is based on data from Grady et al., "Age-related Changes in Brain Activity across the Adult Lifespan (Grady, et al., 2006), archived at the fMRI Data Center (www.fmridc.org , fMRIDC Accession Number 2-2005-119CE). As the title implies, Grady was specifically interested in memory and cognition in different age groups. For present purposes, the four tasks of the experiment are not of interest, but rather the default mode intervals that interweave among them. We analyzed two runs of 384 seconds each in twenty five subjects (20 to 87 years old). Each run comprised eight alternating blocks of tasks and baseline conditions, each twenty four seconds long. During baseline conditions, subjects pushed a button each time a fixation cross appeared on screen. Each image series was preprocessed using Independent Component Analysis (ICA), a statistical method for reducing the dimensionality of data while preserving most of the variance (Calhoun, Adali, Pearlson, van Zijl, & Pekar, 2002). ICA separated twenty components for each run; each component can be understood as a temporally coherent distributed network of activity, an ensemble of voxels that are active or inactive in unison. The figures below use individual and composite component images. Importantly, ICA is a "data driven" method, which identifies activity without reference to the time-course of experimental conditions. While some components approximate the experimental model, most do not. Whole-brain images were collected every 2.5 seconds; for the purpose of this analysis, component images were interpolated to produce a pseudo-image series separated by one second intervals.

The image series for each subject was used to train a series of pattern recognizers, using half of each image series for training and the remaining images for testing. Each subject data set was analyzed separately. The task for each detector was to identify the elapsed duration from the DM block onset to the current time point, using just the current image as input. (Elapsed duration was encoded using a binary scheme, where images before the current time point were binned in one category, and images after the current time point were binned in another. This was repeated at each time point.) Success in this training implies that the pattern recognizers detect information in the images that encodes the relative elapsed time within each DM block. We tested this separately for each of the twenty four possible serial positions, using Support Vector Machines (SVMs) (Cox & Savoy, 2003; Haynes & Rees, 2006; LaConte, Strother, Cherkassky, Anderson, & Hu, 2005; Ma, Zhao, & Ahalt, 2002; Norman, Polyn, Detre, & Haxby, 2006). Additionally, we encoded the basic distinction between task blocks (regardless of task) and DM blocks. To validate the training results, we contrasted them with the same training regime using surrogate data generated by randomly permuting the original image series and executing the same training and testing procedure (Schreiber & Schmitz, 1996, 2000). The point of this comparison was to control for "spurious success" due to the likelihood that some test outputs will match their targets by chance alone (Lloyd, 2002). (For example, if the relative proportion of one target value exceeds another, a pattern recognizer could simply produce that target value at all time points.) At each time point, we calculated a p value for the contrast between the actual SVM output and a set of one hundred surrogate outputs. The figures

below are based on these comparisons. Accordingly, they are statistical parametric maps in the temporal domain, or T-SPMs.

Figure 1 displays the percentage of runs in our analysis in which serial position (elapsed time) could be recovered ($p < .05$, determined by a one-tail t-test, corrected for multiple comparisons), at each time point in the default mode blocks reserved for testing (that is, the latter four of the eight DM blocks). In general, temporal information was recovered for most runs, with better detection after about five elapsed seconds. The figure thus confirms that temporal information is encoded in the ensemble of independent components. Time in this situation was not implicated in any particular behavior during these blocks, nor was time signaled by any environmental feature. (Subjects could count scanner pulses or their own rhythmic button presses, but there would be little motivation to do so.) Rather, passing time is a spontaneous aspect of evolving brain activity in these subjects, and in general is quite accurate as a relative measure of elapsed duration within each block.

Figure 1. Pattern recognition after SVM training, at each time point during off-task blocks. The X-axis shows elapsed time in seconds from the onset of each off-task interval. The Y-axis shows the percentage of sessions in which elapsed duration information could be recovered by SVM pattern detection after training, as compared to control trials using surrogate data. Four 24-second blocks were in the test set for each session.



III. Neural correlates of temporality

This result above naturally raises the question of which brain regions might be particularly involved in encoding temporal information. Something in the distributed pattern of component activity encodes elapsed time. The encoding could be embedded in a single component, activated at a different intensity at each time point, or it could be encoded in distinct patterns involving all components, or it could be somewhere in the middle of a spectrum from localized to distributed. To explore these possibilities, the contribution of each component to temporal representation was calculated using a "leave-one-out" strategy (Chapelle & Vapnik, 2000; Opper & Winther, 2000), as follows: The SVM pattern learning analysis was repeated on each subject image series, however with each repetition one component was omitted from the training and test set. That is, the analysis repeatedly used 19 of 20 components, recording the impact of omitting each component on the overall capacity of the image to represent elapsed time. (This impact was measured as a change in p value from the twenty-component baseline to each of the nineteen-component alternatives. The greatest increase in p signaled the component whose omission had the greatest negative effect on learning performance. This component would accordingly be the largest contributor to the distributed pattern encoding temporal information.) The component of greatest impact was then removed, and the leave-one-out process repeated on the remaining nineteen components, reducing them to eighteen, and so on until the six most important contributors to the encoding of temporal information had been identified.

The leave-one-out strategy offers a systematic way to explore the joint contribution of multiple components to the representation of information. The analysis revealed that, on average, three components generally accounted for the encoding of temporal information. However, the intersubjective variation was large. After one component was removed, 92% of component series still achieved significant pattern detection ($p < .05$, corrected for multiple comparisons); minus 2 components: 84%; minus 3 components: 68%; minus 4 components: 50%; minus 5: 36%; minus 6: 30%. Thus, no one component explained the recoverable temporal information, except in a few cases (4 runs out of 50). But neither is it the case that the information was spread evenly across all components: 35 out of 50 runs embedded the significant temporal information in six or fewer components. Figure 2 displays aggregate results for all subjects, partitioning the default mode blocks into eight three-second intervals. For each session, the ICA magnitudes of significant components at each time point were used as coefficients, and the significant components combined as a single image representing the distributed brain regions active during that three second window. Note that the brain images combined in this figure were those found significant in the leave-one-out analysis; the figure does not show evolution of the default mode overall, just the evolution of areas involved in encoding temporal information. Significant voxels were probed by comparing image series for all sessions. Dark areas in figure 2 represent brain areas with increased activity, relative to the mean, in all subjects ($p < .0001$). Light areas represent brain regions where activity was significantly diminished, relative to the mean.

Figure 2. Dynamic evolution of temporal information during off-task time periods. Aggregate "glass brain" images showing independent components that are most significantly involved in encoding

temporal information ($p < .0001$), combined in three-second blocks spanning the off-task intervals. Dark areas indicate positive component activation; light areas indicate decreased activation.



This "glass brain" presentation affords an overview of the dynamic evolution of brain activity relevant to temporality throughout the default mode blocks in the Grady study. In general, each aggregate image is more like its near neighbors than those at greater lags. But overall, the series displays broad changes over time, involving broadly distributed brain areas. No region is either continuously active or suppressed during the DM encoding of temporality.

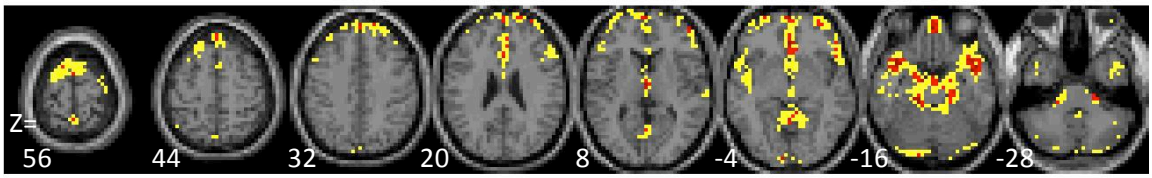
Anatomical region identification is difficult in Figure 2. Accordingly, figure 3A displays the standard deviation of the DM temporality image series, over time. That is, the figure shows areas of the brain where change in temporal information is greatest. Once again, individual variation is large, but some commonalities emerge, including changing activity in superior and medial frontal regions, anterior and posterior cingulate, hippocampal regions, cerebellum, and others. While figure 3A may look like a statistical parametric map of activated regions following an experimental time course, it is not. These are regions whose *changes* encode temporal information during periods when no other external conditions are changing. With the passing seconds during each DM block, these areas are waxing and waning in complex configurations, as implied by figure 2, above. Reflecting the broad distribution of the involved areas, this will be referred to as a Dynamic Temporality Network.

Figure 3. Changing regions underlying temporal information encoding, compared with aggregate Default Mode network in 25 subjects, 2 runs each.

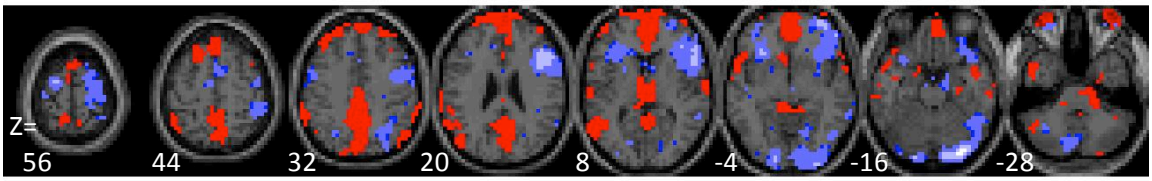
A. Standard deviation of dynamic change in temporal components in fifty runs. For each run (two per subject), voxels above the 95th percentile in standard deviation were noted. The aggregate image here displays voxels above the 95th percentile from the first pass. Thus, voxels in the image represent the top .25% of standard deviation values. Voxels in yellow fell in the top .25% standard deviation in one to fifteen runs. Voxels in red were in the top .25% in more than fifteen runs.

B. Aggregate best-match Default Mode components in fifty runs, showing voxels above the 95th percentile, activated (red) and deactivated (blue).

A.



B.



IV. Temporal information and the Default Mode Network

The Default Mode network generally encompasses the medial prefrontal cortex (MPFC), posterior cingulate cortex (PCC), inferior parietal cortex (IPC), lateral temporal cortex (LTC), and hippocampal regions, among others (Buckner, et al., 2008; Mason, et al., 2007; McKiernan, et al., 2003; Raichle, et al., 2001; Sheline, et al., 2009). What we have called the Dynamic Temporality Network (DTN) overlaps with the "traditional" Default network in several areas, most prominently the anterior DM regions: MPFC, LTC, and hippocampal regions. Figure 3B shows a more detailed image of DM components in the Grady et al. study. For each run, we determined the component whose time course most correlated with the sequence of off-task default mode blocks. The figure is a composite of these, showing the top 5% of voxels common to the subjects and runs studied. Figure 3B generally fits within the outlines of a "classical" default mode network. More important, the DM network in this study can be compared to the sites of dynamism associated with temporally specific information. This match is fairly close, save for the weaker involvement of posterior areas (PCC and IPC; however, the DT network trends toward involvement of those areas, just slightly below the threshold for figure 3A). The comparison suggests a functional role for the DM network that incorporates the function of the Dynamic Temporal network, namely, the encoding of elapsed time.

Before considering the positive implications of this hypothesis, some potential objections should be entertained. First, by basing the analysis on off-task blocks, we already restrict consideration to regions that are active off-task, which is the DM network itself. Accordingly (by this objection), temporality will be found in this network "by default." In reply to this objection, we note that the SVM analysis was not limited to DMN regions, and that any supposition that only DMN components are active in off-task blocks is false. Independent Component Analysis reveals complex oscillations in every component. Even the component most significantly correlated with the default mode is only weakly associated with it.

A second objection could challenge the identification of the Dynamic Temporal Network with an aspect of conscious experience. Specifically, it might be that the time-determining information driving SVM pattern detection is merely the intrinsic oscillation of the DMN (Greicius & Menon, 2004), and presumably not a correlate to any feature of consciousness. Greicius and Menon reported intrinsic DMN oscillations at approximately .005, .015, and .035 Hz, which oscillations were unrelated to the task wave-form. Neither observation characterizes the time course of the information available to temporal pattern detection. The variation within the DTN does specify time points in the experiment. If this specificity is due to a spontaneous oscillation, then that oscillation is at exactly (some multiple of) .0417 Hz and consistently entrained to the task wave form over more than six minutes. It seems highly unlikely that the machine learning analysis could exploit intrinsic oscillations that are unrelated to cognitive brain activity.

But what of the more specific identification of the DTN with the conscious awareness of time? The involved areas are implicated in a wide range of tasks, and default mode activity has been linked to an increase in task-independent thoughts (daydreaming) and inversely related to task difficulty (Mason, et al., 2007; McKiernan, et al., 2003; Sheline, et al., 2009). These "task independent thoughts" have been characterized as internal mental simulations, including planning, auto-biographical memory, "theory of mind" reasoning, and other forms of self-reflection (Addis, Wong, & Schacter, 2007; Amodio & Frith, 2006; Andreasen, et al., 1995; Buckner, et al., 2008; Greicius & Menon, 2004; Saxe & Kanwisher, 2003; Svoboda, McKinnon, & Levine, 2006).[2] Alternatively, the DM network may serve to monitor the environment while focal attention is relaxed (Gilbert & Wilson, 2007; Gusnard & Raichle, 2001; Hahn, Ross, & Stein, 2007; Shulman, et al., 1997). Here the most telling considerations are already in figure 1. The SVM analysis is not indexed to the total temporal sweep of the experiment, but instead assumes that each off-task block will be accompanied by a dynamically similar awareness of time spent during that 24-second block. For example, the analysis assumes that subjects will be in similar states of temporal awareness at five seconds, ten seconds, and so forth during each off-task block. This cross-block similarity was confirmed. It is unlikely, then, that episodes unrelated to time (like daydreaming) could be mapped so specifically to the same time point in each of eight blocks. The involved brain regions are certainly associated with conscious cognition in their other functions; it would be arbitrary to deny that this temporal mapping is not also an aspect of consciousness.

Temporal information does not exclude other functions expressed within a single complex state, even if only part of the brain is considered. The temporal analysis here does not contradict the other research on the cognitive functions of the default mode. Temporality also reconciles the tension between inner

[2] Several recent articles speculate that DM activity embodies William James' "stream of consciousness" (Buckner, et al., 2008; Fransson, 2006; Greicius & Menon, 2004; McKiernan, et al., 2003). This is a mistaken reading of James' metaphor (James, 1890), which was intended to capture a feature of *all* conscious thought:

> Consciousness, then, does not appear to itself chopped up in bits. Such words as 'chain' or 'train' do not describe it fitly as it presents itself in the first instance. It is nothing jointed; it flows. (p. 239)

Most of James' examples are perceptual and active. The DM authors are more likely thinking of the post-Jamesian literary meaning of stream of consciousness, as in the novels of Dorothy Richardson, VirginiaWoolf, Katherine Anne Porter, William Faulkner, James Joyce, and many others. These streams do have the disengagement and inner focus of default cognition.

mentation and environmental monitoring. In this experiment, elapsed time is not signaled by the environment, entailing that it must arise endogenously. But it is a property of environmental objects – an endogenous construct applied to the outside world.

V. Conclusion

The successive analyses in this paper lend support to a new interpretation of the function of the Default Mode Network, relating it to the continuous monitoring of passing time, the structural feature of consciousness known as temporality. The emerging picture of default mode function, and the specific function of the overlapping Dynamic Temporal Network, conforms nicely to the phenomenology of time. In Husserl's classic description (and in the phenomenological tradition since), objects are perceived as an amalgam of their static sensory properties and their dynamic non-sensory temporal properties. One cannot perceive an apple without also perceiving its continually growing duration; even if the apple is immobile, one can tell that it has been immobile for this long, now this, now this. The 3-dimensional object is inflected by the ever moving fourth dimension, even if the "object" is itself a psychological state, insofar as these too have readily perceptible durations. The mental functions of the default mode seem also to be inflected by time. If the hypothesis proposed here is confirmed in other experiments and other analyses, time may take its place as an omnipresent contextual dimension of states of conscious awareness, adding to an emerging science of consciousness an account of the thickness and effervescence of the extended Now.

Acknowledgments:

Thanks for Brian Castelluccio for data analysis and to the Trinity College Interdisciplinary Science Program for research assistance support. An earlier draft of this paper was presented at a workshop on temporality sponsored by the European Platform for Life Sciences, Mind Sciences, and the Humanities at the University of Turku, Finland; I thank the workshop hosts and participants for their comments. Independent Component Analysis was implemented using the ICA Toolbox for Matlab (http://icatb.sourceforge.net/), developed by Vince Calhoun. SVM was implemented using the OSU-SVM toolbox (Ma, et al., 2002).

REFERENCES

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia, 45*(7), 1363-1377.

Amodio, D. M., & Frith, C. D. (2006). Meeting of minds: the medial frontal cortex and social cognition. *Nat Rev Neurosci, 7*(4), 268-277.

Andreasen, N. C., O'Leary, D. S., Cizadlo, T., Arndt, S., Rezai, K., Watkins, G. L., et al. (1995). Remembering the past: two facets of episodic memory explored with positron emission tomography. *Am J Psychiatry, 152*(11), 1576-1585.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci, 1124*, 1-38.

Calhoun, V. D., Adali, T., Pearlson, G. D., van Zijl, P. C., & Pekar, J. J. (2002). Independent component analysis of fMRI data in the complex domain. *Magn Reson Med, 48*(1), 180-192.

Chalmers, D. J. (1996). *The conscious mind : in search of a fundamental theory*. New York: Oxford University Press.

Chapelle, O., & Vapnik, V. (2000). Model selection for support vector machines. In S. A. Solla, T. K. Leen & K.-R. Muller (Eds.), *Advances in Neural Information Processing Systems* (Vol. 12). Cambridge, MA: MIT Press.

Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage, 19*(2 Pt 1), 261-270.

Fransson, P. (2006). How default is the default mode of brain function? Further evidence from intrinsic BOLD signal fluctuations. *Neuropsychologia, 44*(14), 2836-2845.

Gilbert, D. T., & Wilson, T. D. (2007). Prospection: experiencing the future. *Science, 317*(5843), 1351-1354.

Grady, C. L., Springer, M. V., Hongwanishkul, D., McIntosh, A. R., & Winocur, G. (2006). Age-related changes in brain activity across the adult lifespan. *J Cogn Neurosci, 18*(2), 227-241.

Greicius, M. D., & Menon, V. (2004). Default-mode activity during a passive sensory task: uncoupled from deactivation but impacting activation. *J Cogn Neurosci, 16*(9), 1484-1492.

Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. *Behav Brain Sci, 27*(3), 377-396; discussion 396-442.

Gusnard, D. A., & Raichle, M. E. (2001). Searching for a baseline: functional imaging and the resting human brain. *Nat Rev Neurosci, 2*(10), 685-694.

Hahn, B., Ross, T. J., & Stein, E. A. (2007). Cingulate activation increases dynamically with response speed under stimulus unpredictability. *Cereb Cortex, 17*(7), 1664-1671.

Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nat Rev Neurosci, 7*(7), 523-534.

Husserl, E. (1966 (1928)). *Zur Phänomenologie des inneren Zeitbewusstseins (Phenomenology of Inner Time Consciousness)* (Vol. 10). The Hague: Martinus Nijhoff.

James, W. (1890). *Principles of Psychology*. New York, NY: Henry Holt & Co.

LaConte, S., Strother, S., Cherkassky, V., Anderson, J., & Hu, X. (2005). Support vector machines for temporal classification of block design fMRI data. *Neuroimage, 26*(2), 317-329.

Lloyd, D. (2002). Functional MRI and the Study of Human Consciousness. *Journal of Cognitive Neuroscience, 14*(6), 818-831.

Ma, J., Zhao, Y., & Ahalt, S. (2002). OSU-SVM Toolbox for Matlab.

Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.

Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. *Science, 315*(5810), 393-395.

McKiernan, K. A., Kaufman, J. N., Kucera-Thompson, J., & Binder, J. R. (2003). A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. *J Cogn Neurosci, 15*(3), 394-408.

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci, 10*(9), 424-430.

Opper, M., & Winther, O. (2000). Gaussian processes and SVM: Mean field results and leave-one-out. *Advances in large margin classifiers* (pp. 43-65). Cambridge, MA: MIT Press.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proc Natl Acad Sci U S A, 98*(2), 676-682.

Saxe, R., & Kanwisher, N. (2003). People thinking about thinking people. The role of the temporo-parietal junction in "theory of mind". *Neuroimage, 19*(4), 1835-1842.

Schreiber, T., & Schmitz, A. (1996). Improved Surrogate Data for Nonlinearity Tests. *Phys Rev Lett, 77*(4), 635-638.

Schreiber, T., & Schmitz, A. (2000). Surrogate time series. *Physica D: Nonlinear Phenomena, 142*(3-4), 346-382.
Seth, A. (2009). Explanatory Correlates of Consciousness: Theoretical and Computational Challenges. *Cogn. Comput., 1*, 50-63.
Sheline, Y. I., Barch, D. M., Price, J. L., Rundle, M. M., Vaishnavi, S. N., Snyder, A. Z., et al. (2009). The default mode network and self-referential processes in depression. *Proc Natl Acad Sci U S A, 106*(6), 1942-1947.
Shulman, G. L., Corbetta, M., Buckner, R. L., Raichle, M. E., Fiez, J. A., Miezin, F. M., et al. (1997). Top-down modulation of early sensory cortex. *Cereb Cortex, 7*(3), 193-206.
Svoboda, E., McKinnon, M. C., & Levine, B. (2006). The functional neuroanatomy of autobiographical memory: a meta-analysis. *Neuropsychologia, 44*(12), 2189-2208.
Yoshimi, J. (2007). Mathematizing Phenomenology. *Phenomenology and the Cognitive Sciences, 6*(3), 271-291.
Zahavi, D. (1999). *Self-Awareness and Alterity*. Evanston, IL: Northwestern University Press.