Tremotopic Mapping of the Human Thalamic Reticular Nucleus

Joseph D. Viviano

A Thesis Submitted to the Faculty of Graduate Studies in Partial Fulfillment of the Requirements for the Degree of Master of Science

Graduate Program in Biology
York University
Toronto, Ontario
December 2013

The thalamic reticular nucleus is an important structure in the mammalian brain, participating in the coordination of large-scale processes such as sleep and attention. To date, this structure has not been investigated in the human brain. I developed a series of methods for anatomically and functionally localizing the visual regions of the thalamic reticular nucleus in the human brain using magnetic resonance imaging and the presentation of various flicker frequencies. First, I describe the results obtained from a modified retinotopy analysis. I next apply network theory to the data in an attempt to localize the TRN is a data-driven way. Third, I describe a lateral-inhibitory network the TRN participates in. I conclude the TRN plays a role in regulating interhemispheric activity in the brain, and that flicker can be used to probe the resonance properties of neural populations with magnetic resonance imaging.

For Greg Brophy

Acknowledgements

None of this work would have been possible without Dr. Keith Schneider. He offered me an opportunity at a time when it seemed no one else was willing to take a chance on me. His tutelage over these last few years has shaped me immensely, and sparked my curiosity in ways that led to many of the ideas presented here, but most importantly it has given me a focus I have never known before. I know that I will forever be in his debt for steering me in the right direction, wherever that may lead me.

I would like to thank Dr. Kari Hoffman, Dr. Gary Turner, and Dr. Dale Stevens, who have all provided excellent advice and guidance beyond the call of duty. I would also like to thank Joy Williams, who was not only a welcome comrade in the trenches of science, but spent many hours teaching me the details of MRI, and was instrumental in the collection of this most unique data set.

I would be remiss to neglect thanking my lab mates Kevin DeSimone for teaching me the basics as I struggled at the beginning and Larissa McKetton for assisting with the collection of these stunning anatomical images. Thank you to Kyle for his unorthodox caffeinated beverages, and my supportive friends: Devon Murphy, Brendan Hurley, Adam Wray, Danielle Zies, Andrew Edmonds, Sarah Briggs-Jude, Benjamin Cassidy, Timothy Leonard, and Rachel Ganaden.

I would also like to thank (soon to be Dr.) Matthew Lovett-Barron. His friendship through the last decade of my life was instrumental to the development of these ideas. His passion and tenacity is a constant source of inspiration for me, and his kindness has seen me through countless sessions of turmoil. I would not have written this thesis had I not known him.

Finally, I would like to thank my partner, Denise Wong. Her passionate support was instrumental to my well-being through all of this, and her relentless personal ambitions have served as an inspiration throughout this entire project. I must also thank her for being patient with my workaholic tendencies. I would not have written this thesis had I not known her.

Table of Contents

List of Figures

Chapter 1

"The most merciful thing in the world,
I think, is the inability of the human
mind to correlate all its contents."

-H.P. Lovecraft

An Introduction to the Problem

The thalamic reticular nucleus (TRN) is an integral structure in the brain that has not, to the best of the author's knowledge, yet been functionally examined in the human. This is unfortunate, as the TRN is believed to play an important role in the sleep-wake cycle, the allocation of attention, and the coordination of large-scale processes in the brain (Jones, 2009; McAlonan et al., 2006; Steriade and Timofeev, 2003; Steriade et al., 1993). A reliable procedure for identifying the TRN in the human would therefore be beneficial to basic and clinical researchers alike.

In Section **1.1**, I give a brief overview of the current theories surrounding the role of oscillations in brain function and the functional properties of the TRN. In Section **1.2**, I outline the rationale behind the proposed procedure for identifying the TRN using magnetic resonance imaging.

1.1 Structure and Proposed Function of the Thalamic Reticular Nucleus

A well-functioning brain's primary task is to efficiently comprehend the external world and generate appropriate responses. One of the greatest technical challenges the brain faces is reducing the dimensionality of incoming sensory information to a relatively small number of informative features. The labelled lines of sensation theory, supported by a large volume of current evidence, suggests that the brain partially accomplishes this goal by maintaining many parallel streams of information across the bodily senses, each sensitive to particular aspects of the environment (see Norrsell et al., 1999 for a review). These streams are composed of neurons with similar receptive fields

(Hubel and Wiesel, 1959; Sherrington, 1906). For example, points on the retina are organized as a map of visual space that is preserved through the early sensory thalamic and cortical regions. This organization allows the brain to internally represent a large percentage of the sensory world in cases where the most important aspects are not known, and to efficiently discard unimportant processes in a single pass as soon as the system has converged on a probable point of interest. The complexity of receptive field representations increase dramatically when one begins to observe small networks of neural activity. Sensory neurons with well-defined receptive fields in the presence of simple stimuli display higher-order responses to complex stimuli with statistical regularities, which requires information from complementary neurons (Chacron et al., 2003; Estebanez et al., 2012). Therefore, the mechanisms whereby the brain is able to generate functional cell assemblies of tuned neurons is integral to understanding how the brain able to effectively understand and react to the environment (Hebb, 1949).

A generic mammalian neural sensory system can be crudely defined as the following: a sense organ is composed of neurons tasked with the decomposition of a particular physical feature into an array of channels that comprise a topographic map. These channels carry sensory information via precisely-timed action potentials along axons to a topographically-organized first order sensory thalamic nucleus. This nucleus then sends orderly inputs to various cortical regions, including a primary sensory cortex that follows a relatively regular topographic organization (Sherman and Guillery, 1996). Cortical regions communicate with one another via second order relays in the thalamus, thereby making the thalamo-cortical loop ideally suited for multisensory communication, integration, and modulation (Guillery and Sherman, 2002).

These thalamo-cortical and cortico-thalamic interfaces all pass through and are regulated by the thalamic reticular nucleus (TRN), a thin layer of γ-Aminobutyric acid (GABA) releasing cells wrapping the dorsolateral and anterior segments of the thalamus (Jones, 1975). The TRN receives topographic glutamatergic inputs from fibres traveling in both the thalamo-cortical and cortico-thalamic directions via metabolic glutamate receptors (Jones, 2007; Paz et al., 2011) and drives inhibitory feedback to thalamic relay

cells via $GABA_A$ and $GABA_B$ receptor activation (Kim et al., 1997), generating synchronous activity in the thalamo-cortical system (Halassa et al., 2011; Huntsman et al., 1999). Increased TRN activity produces topographically-matched inhibitory post-synaptic potentials (IPSPs) in the thalamus, transiently reducing tonic and bursting activity (Funke and Eysel, 1998). The TRN itself is divided into modality-specific sectors roughly preserving the topography seen in each sectors associated first order thalamic relay and sensory cortex, and receives inputs from every sensory system, the brainstem, the motor system, and limbic system in multiple mammalian species (Conley et al., 1991; Crabtree, 1992a, 1992b, 1996; Guillery and Harting, 2003; Künzle, 1976; Lozsádi, 1994; Montero et al., 1977; Pinault et al., 1995a, 1995b; Wang et al., 2001; Zikopoulos and Barbas, 2006, 2012).

These properties leave the TRN well-equipped for the task of synchronizing the activities of multiple disjoint thalamic nuclei, serving as a key node in the sensory-motor network (Crabtree and Isaac, 2002). A long line of research has established the role of neural oscillations in the coordination of neural activity at long distances through low frequency coherence, and local computations at higher frequencies (for a review, see Buzsáki and Draguhn, 2004). Neural oscillations are typically defined as comprising multiple frequency bands, and can be categorized into the following (sometimes overlapping) frequency bins: slow-wave [0.025–1.4 Hz], delta [δ; 1.5–4 Hz], theta [θ; 4–10 Hz], alpha [α; 8–13 Hz], beta [β; 10–30 Hz], gamma [γ; 30–80 Hz], and fast/ripple [γ_H; 80–600 Hz]. The communication-through-coherence hypothesis proposes that interactions between groups of neurons are facilitated by coherent oscillations in the LFP in distant cortical regions, creating a flexible and transient neural assembly of highly correlated action-potentials quantized by the assembly oscillation-phase (Canolty et al., 2010; Fries, 2005; Womelsdorf et al., 2007), allowing for reliable signal transmission in a noisy and unpredictable environment (Montemurro et al., 2008; Wang et al., 2010) and facilitation of effective spike-timing-dependent plasticity (Markram et al., 1997; Song et al., 2000; Thivierge and Cisek, 2011; Wespatat et al., 2004).

There is evidence that low-frequency thalamo-cortical oscillations coordinate local cortical processes. Thalamic stimulation at α-band frequencies trigger cortical activity that is tightly coupled to the phase of the thalamic driver (MacLean et al., 2005). Magnetoencephalography (MEG) shows that γ-band oscillations are phase-locked to ongoing α-band oscillations in the human primary visual cortex (V1; Osipova et al., 2008). Furthermore, electrode recordings in primates show cross frequency coupling between γ-band local field potentials (LFPs) in the superficial layers (associated with local intra-cortical connectivity), and α-band LFPs in the deeper layers of V1 (associated with long-range cortico-thalamic and cortico-cortical connectivity), suggesting that α-band oscillations control the precise timing of local γ-band dependent cortical processing (Spaak et al., 2012). This implies that thalamic drivers might modulate the coherence of disparate regions of sensory cortex, facilitating communication between them (Fries, 2005; Womelsdorf et al., 2007). At the cellular level, a Ca^{2+}-dependent dendro-somatic transfer of local, presynaptic, voltage fluctuations propagate along the axon, with strong transfer of frequencies below 20 Hz, producing mild hyper and de-polarizations of the axon (Shu et al., 2006). Thus, neurons themselves act as low-pass filters, and promote graded (analog) potentials that favour triggered (digital) action-potentials at a particular phase.

The organization of electrical and chemical communication within the TRN suggests that the structure acts to mediate low-frequency oscillations within the thalamo-cortical loop. Connexin36 (Cx36) dependent electrical gap junctions are found between neighbouring TRN neurons (Landisman et al., 2002). Bath-application of metabolic glutamate receptor agonists in TRN slices produces locally-synchronized 5–15 Hz rhythms (less than 25–40 μm from the soma) via Cx36 gap junctions (Long et al., 2004). The TRN is also known to produce 7–12 Hz spindle-type oscillations during alert behaviour and non-rapid eye movement slow-wave sleep, similar in many ways to the α-band oscillation observed throughout the thalamo-cortical system (Bazhenov et al., 2000; Contreras et al., 1997; Steriade, 1997). The similar response properties of Cx36 connections suggest that they drive local synchronization of spindle-frequency

oscillations within the TRN in response to glutamate (Bal and McCormick, 1993; Halassa et al., 2011).

Approximately 7–25% of TRN synapses are GABAergic (Liu and Jones, 1999; Williamson et al., 1994), which receive afferents from the substantia nigra (Paré et al., 1990), basal forebrain (Jourdain et al., 1989), and other TRN neurons (Asanuma, 1994); these synapses drive intra-TRN IPSPs. Fast-acting reversal potentials are mediated by ionotropic $GABA_A$ receptors, producing spindle oscillations (Bazhenov et al., 1999). In contrast, metabotropic $GABA_B$ receptor activation results in slow IPSPs (~ 300 ms); pharmacological blockade of $GABA_A$ receptors results in the normal α-frequency spindle wave being replaced by a paroxysmal slow oscillation in which many thalamic and TRN neurons discharge synchronously within the δ-band frequencies also seen during non-rapid eye movement slow-wave sleep (Sanchez-Vives and McCormick, 1997; Sanchez-Vives et al., 1997; Shu and McCormick, 2002). $GABA_A$ receptors in the TRN have also been found to initiate postsynaptic depolarizations of other GABA neurons in the mature brain due to the low concentration of the Cl^- transporter KCC2, leading to the activation of postsynaptic T-type Ca^{2+} channels, subsequent action potentials in TRN neurons, and finally powerful inhibitory inputs to thalamic relay neurons (Sun et al., 2012). The $GABA_B$-mediated slow wave can be conceived of as a modulator of the $GABA_A$-mediated spindle-frequency carrier wave. The slow wave is regulated and amplified through competition between T-channel mediated Ca^{2+} intake and the subsequent activation of small-conductance-type outward K^+ currents, which produce a post-burst hyperpolarization in TRN dendrites at δ-band frequencies (Cueni et al., 2008).

Brain-wide oscillations are apparent at these frequencies during sleep, as populations of neurons vacillate between α-frequency "up" states resembling attentive dynamics, and a δ-frequency "down" states, generating a superimposed slow-wave oscillation less than 1 Hz in frequency (Steriade et al., 1993, 2001). Indeed, dendritic $GABA_B$ receptor-mediated reciprocal connectivity within the TRN is believed to underlie large-scale thalamic synchronization across the topographic map by inhibiting action-potentials over the majority of a slow oscillation (Breton and Stuart, 2012; Juhász

et al., 1994) via presynaptic autoinhibition of IPSCs in both TRN and relay cells (Ulrich and Huguenard, 1996). In contrast, intra-TRN $GABA_A$-mediated transmission acts to decrease the number of times a TRN neuron will spike per oscillation, regulating general excitability by reducing the number of times a TRN cell will burst in response to excitatory inputs from the thalamus and cortex (Sohal and Huguenard, 2003; Warren et al., 1994).

Knockout of the α_3 or β_3 subunits of the $GABA_A$ receptor in the rat results in a decrease of $GABA_A$-mediated intra-TRN inhibition, as well as a dramatic increase in δ-frequency oscillatory synchrony (Huntsman et al., 1999; Schofield et al., 2009). These subunits are known to uniquely coexist on receptors that mediate phasic inhibition on TRN neurons in the rat, and in the monkey to a lesser extent (Browne et al., 2001; Huntsman et al., 1996). Furthermore, presynaptic $GABA_A$ receptor activation has an inhibitory influence among synchronously firing TRN neurons, which could act to maintain distinct synchrony within a restricted region of the topographic map (Bazhenov et al., 1999; Lam et al., 2006).

The electrical and chemical connections in the TRN are distinctly organized. Photolysis of caged glutamate in coronal and horizontal slices of the rat thalamus reveal a higher incidence of chemical connectivity in the anteroposterior plane (60 % of neurons), and a higher incidence of Cx36 gap junction connectivity in the dorsoventral plane (47 %; Deleuze and Huguenard, 2006). Considering that the topographic maps of the TRN are organized across the anteroposterior plane (Lam and Sherman, 2010) and the laminar structure arises from communication with first and second order sensory-motor structures (Lam and Sherman, 2007), it appears that electrical communication in the TRN is biased to synchronize topographically-specific information between lower and higher-order intra-modal structures, while chemical communication in the TRN is biased toward regulating the spatial extent of synchronization across multiple topographic maps at once (Lam et al., 2006). This organization accounts for the TRN's apparent ability to dynamically influence receptive field properties of thalamic neurons (Cotillon-Williams et al., 2008).

To summarize, evidence suggests that $GABA_A$-mediated connections act as intra-TRN desynchronizers, while $GABA_B$-mediated connections act as intra-TRN recruiters, creating emergent patterns of phase-synchrony within the thalamo-cortical loop over time. As the TRN switches from low-frequency tonic to burst firing mode, thalamic neurons change from $GABA_A$ mediated low-amplitude to high-amplitude IPSP-dependent activity, synchronizing disparate regions of the thalamo-cortical loop via coherent α-band oscillations. Cortico-thalamic volleys drive sustained hyperpolarizing potentials in the TRN, producing α-frequency spindles; prolonged spindle activity leads to the recruitment of metabotropic $GABA_B$ receptors, resulting in δ-frequency $GABA_B$ mediated IPSPs that act to modulate the α-band carrier wave (Fuentealba et al., 2004; Kim et al., 1997; Zhang and Jones, 2004). $GABA_B$-mediated recruitment of α oscillations across the topographic map is therefore resilient to noisy inputs to the TRN, and might therefore function as a robust coincidence detector across multiple oscillatory processes.

Some of the few *in vivo* visual experiments on TRN suggest it has change detection properties: neurons transiently respond as a moving object enters a given receptive field, and give strong transient on-off responses to flashing stimuli (Dubin and Cleland, 1977; Funke and Eysel, 1998; Yu et al., 2009). These sharp responses would allow for the detection of phase alignment of multiple stimuli as they would produce the maximum response in the TRN, driving optimal suppression of the primary sensory relay at the appropriate topographic coordinates (Sillito and Jones, 2008). The subsequent synchronized disinhibition across the topographic map would produce a phase-aligned oscillation in the thalamo-cortical loop, facilitating object segmentation using temporal evidence (Castro-Alamancos, 2000). The same mechanism might be used to facilitate selective attention by down-regulating the transmission of unimportant sensory events to the cortex. There is single unit evidence in the monkey of a modality-specific spike-rate increase within the visual sector of the TRN during an audio-visual attention task when attention was directed within the receptive field of the recorded neuron, and evidence that this increase modulates the visual primary sensory relay, the

lateral geniculate nucleus (LGN) during a unimodal selective-attention task (McAlonan et al., 2006, 2008). Due to the restricted spatial resolution of this recording technique, however, it is unclear how attention modulates the activity of the TRN across the topographic map.

1.2 Approach for Mapping the Thalamic Reticular Nucleus in the Human

Of the methods available to study the brain non-invasively, only magnetic resonance imaging (MRI) allows us to directly sample signals from the thalamus. MRI techniques can be roughly classified as either structural approaches, allowing for high-resolution static imaging of the bodily tissue, and functional imaging, allowing for lower-resolution imaging of the body's metabolic demands over time. Pioneering work has shown that specially-tuned MRI sequences are sensitive to the relative proportion of deoxyhemoglobin and oxyhemoglobin within a three-dimensional region of the body, or voxel, allowing one to track the metabolic demands of the body with relatively high precision over time using the so-call blood oxygen level dependent (BOLD) signal (Ogawa et al., 1990). The fine microvasculature of the brain allows us to track neural activity indirectly via the BOLD signal, but the direct relationship between the metabolic demands of ongoing neural processes and the recorded BOLD signal remain unclear. There is a good deal of evidence that the BOLD contrast is representative of LFP power in the α , β, and γ-frequency range, and therefore is representative of inputs and local processes rather than spiking outputs from a region (Goense and Logothetis, 2008; Logothetis, 2008; Logothetis et al., 2001; Magri et al., 2012). Therefore, the oscillatory function of the nervous system is observable with functional MRI, although it is very challenging, if not impossible, to determine the frequency of the underlying process using BOLD contrasts alone.

Functional MRI experiments have historically relied on a model of the expected BOLD response, or canonical hemodynamic response function (HRF; Friston et al., 1998; Glover, 1999). The expected BOLD response is traditionally computed as a linear convolution of this canonical HRF with a time series representing the onset and offset

of particular stimuli and/or behavioural events, and then the error of a fit between this theoretical time series and the actual BOLD response is used to determine the likelihood that a particular region of the brain was responsive to the experimental paradigm. The HRF is known to be highly variable across subjects, and across brain regions within a subject, as well as across time points within a subject, leading to a problem of canonical subjects and a bias towards detecting effects in particular brain regions (Handwerker et al., 2004, 2012). With respect to the present problem of developing a localizer for the human TRN, these biases are highly undesirable. In particular, the metabolic demands of the TRN, composed solely of interneurons, are very different and generally less than those seen in the typically-imaged cortex, where multiple cell types contribute to the neuroimaging signal (Buzsáki et al., 2007). It is for this reason that I propose one should use a model-free analysis, which is unbiased towards detecting any particular neural population, and has previously shown to detect brain-wide responses to sensory stimulation when used in conjunction with time-locked signal averaging (Gonzalez-Castillo et al., 2012).

The pioneering work in high-resolution thalamic imaging made use of such an analysis. Multiple studies with increasing spatial precision have made use of a phase-encoded stimulus design to detect voxels that respond at the fundamental frequency of the stimulus over the course of the experimental run (Engel et al., 1997; Schneider and Kastner, 2005; Schneider et al., 2004). These studies made use of a simple visual stimulus that would travel through space repeatedly at a constant speed, producing traveling waves of BOLD activity across the retinotopic map in various visual structures. Determining the retinotopic representation of a given voxel therefore only requires the detection of BOLD modulation at the fundamental harmonic of the stimulus repetition period, where the phase of this modulation denotes the approximate retinotopic population receptive field of a given brain region. Typical stimuli make use of a rotating wedge and expanding ring to map polar angle and eccentricity, respectively, and have been successfully used to map large swaths of the human cortex (see Wandell et al., 2007 for a review).

This method produces maps of voxels representing various visually responsive structures. Analysis of signals in the thalamus have traditionally made use of a functionally-defined region-of-interest (ROI) approach for the mapping of the LGN, superior colliculus (SC), and pulvinar (Kastner et al., 2004; Schneider and Kastner, 2005; Schneider et al., 2004). This requires the collection of a high-resolution anatomical reference image and the manual definition of the locations of the structures of interest. With respect to the LGN, SC, and pulvinar, this approach was tenable using a T1-weighted sequence, however the exact boundaries between the structures remained ambiguous from the contrast contained within the anatomical images alone. T1-weighted images are optimal for differentiating between fat and water in the body, providing excellent contrast between the grey and white matter of the cortex. In the thalamus, however, there less spatial segregation of myelinated axons and the thalamic nuclei of interest at the resolution attainable using MRI, rendering very low-contrast definition of the thalamus with reasonable scan lengths. The human TRN is never more than 3mm in diameter, and often close to 1mm in diameter, rendering it undetectable with a T1-weighted anatomical image (Mai et al., 2007). Since these previous experiments discarded functional activity outside of their roughly-defined ROIs, it remains possible that they missed TRN responses to their stimulus, or included TRN responses within their ROIs. For these reasons, this project made use of proton density (PD) weighted anatomical images, as described in Devlin et al., 2006. The density of protons is approximately 20% greater in grey matter than in white matter, allowing one to distinguish between adjacent thalamic nuclei (Fujita et al., 2001; Wood et al., 1993).

Historically, the stimuli used to map the visual regions of the brain were high-contrast patterns which would contrast-reverse at a steady flicker frequency. Due to the poor signal-to-noise ratio of functional MRI, most retinotopy studies were tuned to use an ~8 Hz flicker frequency, which was found to produce the most robust BOLD response in the LGN and V1 (Chen et al., 1999; Engel et al., 1997; Schneider et al., 2004; Singh et al., 2000; Thomas and Menon, 1998, although see Ozus et al., 2001). The fact that these flicker frequencies, which fall in the α range, preferentially drive early visual

areas known to preferentially oscillate at α-band frequencies (Hughes et al., 2004), suggest that visual flicker may be a way of probing the resonance properties of various neural populations non-invasively (Hutcheon and Yarom, 2000). There is electroencephalography (EEG) evidence that the human visual cortex resonates at α-band frequencies when presented with an α-band visual flicker during variable flicker frequency presentation (Herrmann, 2001), and complimentary LFP evidence from the cat V1 during variable flicker presentation (Rager and Singer, 1998), but to date, no study has systematically investigated whether the use of variable flicker frequency can be used to entrain distinct neural populations in the human brain. Baring in mind that the BOLD contrast reflects the presence of oscillatory activity, but not its frequency, it seemed probable that the visual areas of the TRN associated with cortico-thalamo-cortical communication might appear silent with the presentation of only α-band flicker frequencies, but will be driven at higher frequencies, as pulvinar-mediated inter-cortical communication is associated with β-band oscillations (Wang et al., 2012; Wróbel et al., 2007), and the TRN is proposed to mix signals arising from first order and second order thalamic nuclei (Lam and Sherman, 2007).

To address this, I propose a novel dual-frequency phase-encoded stimulus, henceforth referred to as the tremotopic stimulus. A tremotopic stimulus consists of a phase-encoded retinotopic stimulus with a superimposed variable flicker frequency presentation. These two features must complete an integral number of periods per experiment at different frequencies. If a neural population is resonant when presented with a particular flicker frequency, the retinotopic carrier wave should be amplitude-modulated according to a voxels population-level flicker frequency preference. I hypothesize that this approach will allow me to identify visually-responsive areas of the human thalamus with greater precision than standard retinotopy, including the TRN, which should express unique temporal properties when analyzed in conjunction with the nuclei it acts to modulate.

Furthermore, in order to avoid experimenter bias in selecting inappropriate ROIs that are merely noise that happen to fall on top of the anatomically defined TRN, this

work attempts to use the manually-defined ROIs to verify the data-driven analysis. Since the TRN maintains representations of every sensory modality, and it is unclear where the visual sector of the TRN would be in the human, and ROI-based approach seems likely to miss interesting areas of activation in the thalamus. Instead, I propose a principled data-driven approach for the identification of the visually-responsive thalamic nuclei, including the TRN. This analysis is then followed by an ROI based time-series analysis of the identified TRN regions.

Specifically, I hypothesize that one should be able to identify the TRN as a visually driven cluster of voxels falling in an anatomically defined region that expresses a flicker frequency tuning in correspondence with the adjacent thalamic nuclei it acts upon. This cluster of voxels should also show functional relationships with multiple first and second order thalamic structures, as found by (Lam and Sherman, 2007). If multiple neural populations express different flicker frequency preferences, the expected connectivity of the TRN should be evident from a graph-theoretic analysis of the visually driven thalamic voxels. Specifically, I expect the TRN to act as a hub in the thalamic network, connecting these first and second order nuclei.

Chapter 2

An Overview of the Experiment Design

Here I present the details of the experiment design used to test my hypothesis using high-resolution structural and functional MRI. Analysis and results are detailed in Chapter 3. All data was collected on a 3 T Trio Tim Siemens scanner using a 32 channel head coil at the York University Neuroimaging Center. Six people (one female, one left-handed) participated in the full experiment, all with normal or corrected to normal vision. Informed consent was collected from all participants in line with the York University Ethics Board. The data was collected over three separate one-hour sessions.

In Section **2.1**, I give a detailed account of the generation of detailed proton density weighted images of the thalamus. In Section **2.2**, I give an overview of the stimulus design and presentation. In Section **2.3**, I present the high-resolution functional MRI imaging protocols designed for this experiment.

2.1 Anatomical Imaging of the Human Thalamus

High quality PD scans were collected from all participants during the first session to anatomically define thalamic structures. Multiple runs were collected (94 runs for S1, 40 runs for S2–6). Each run consisted of 30 coronal slices covering the entire thalamus at high resolution (512×512 matrix for S1, 256×256 matrix for S2–6 with a 192 mm field of view), 1 mm thickness, and no gap between slices leading to an in-plane resolution of 0.375 mm and 0.75 mm for S1 and S2–6, respectively [repetition time (TR) = 3000 ms; echo time (TE) = 22 ms; flip angle, = 120°]. These images were co-registered and interpolated to 2× resolution in one step using align_epi_anat.py program and the local Pearson correlation cost function from the AFNI software package (built Aug. 9th, 2012; http://afni.nimh.nih.gov/afni; Cox, 1996; Saad et al., 2009). Spatial interpolation was accomplished using the minimum side lobe three-term cosine sinc function tapered across 7 voxels (Nuttall, 1981). Finally, a mean PD image was produced for each subject.

As both the high-resolution PD and functional scans imaged only a small portion of the brain, T1-weighted MPRAGE images were collected from each subject to facilitate co-registration [TR = 1900 ms; TE = 2.52 ms; flip angle = 9°; 256×256 resolution matrix with a 256 mm field of view, slice thickness of 1 mm and slices gap of 50% between slices leading to a voxel resolution of 1 mm]. Volumetric segmentation was performed on these images using Freesurfer, which is documented and freely available for download online (version 5.1.0; http://surfer.nmr.mgh.harvard.edu/). Briefly, this processing includes the removal of non-brain tissue using a hybrid watershed/surface deformation procedure (Segonne et al., 2004), automated Talairach transformation, segmentation of the subcortical white matter and deep gray matter volumetric structures (Fischl et al., 2002, 2004), intensity normalization (Sled et al., 1998), and surface deformation following intensity gradients to optimally place the gray/white and gray/cerebrospinal fluid borders at the location where the greatest shift in intensity defines the transition to the other tissue class (Dale and Sereno, 1993; Dale et al., 1999; Fischl and Dale, 2000). Freesurfer morphometric procedures have been demonstrated to show good test-retest reliability across scanner manufacturers and across field strengths (Han et al., 2006).

These averaged PD images provide objectively greater contrast in the thalamus, as shown in **Figure 1**. Here, we provide the reader with a detailed hand-labelled thalamic atlas and comparison images from a T1-weighted scan. As can be seen, PD imaging allows for the clear differentiation of the LGN, medial geniculate nucleus (MGN), SC, TRN, pulvinar nuclei, and putamen. Of note, the boundaries between the TRN and surrounding structures are apparent. While the 512 matrix scan collected from S1 offers superior definition of the smallest structures, the 256 matrix scans collected for S2–6 provide sufficient visualization of the thalamus for approximating the location of the TRN. I therefore chose to rely mostly on these lower resolution scans, as they can be collected in one eighth the scan time.

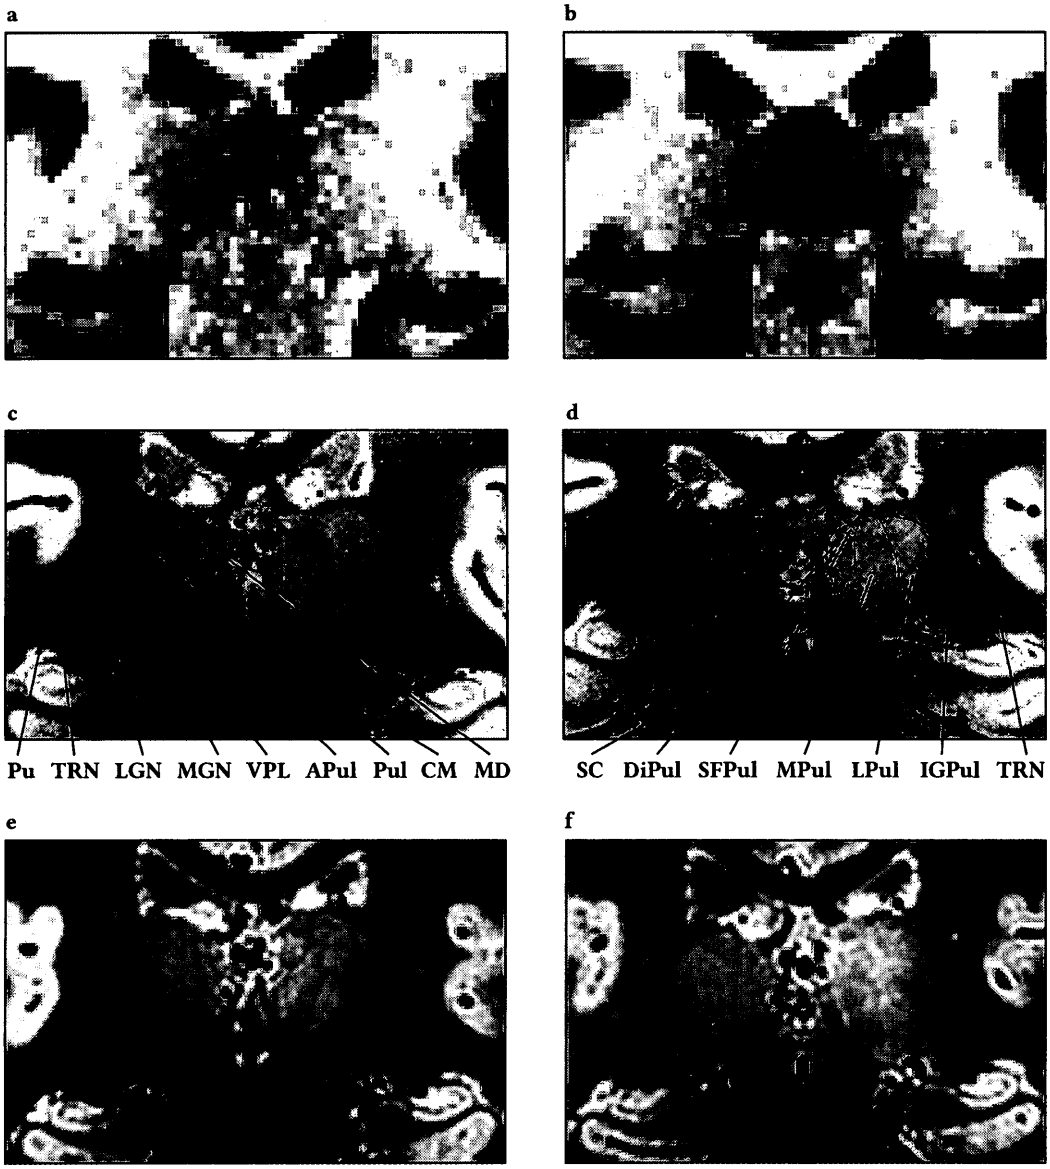

Figure1. Proton density weighted imaging of the human thalamus. All images in column 1 and 2 represent the same structures in the brain, respectively. a & b) T1-weighted images. c & d) Proton density (PD) imaging of the thalamus with a 512 matrix in S1. d & e) PD imaging of the thalamus with a 256 matrix in S6. Pu = Putamen, TRN = thalamic reticular nucleus, LGN = lateral geniculate nucleus, MGN = medial geniculate nucleus, VPL = ventral posteriolateral thalamic nucleus, APul = anterior pulvinar, Pul = pulvinar, CM = centromedian nucleus, MD = medial dorsal nucleus, SC = superior colliculus, DiPul = diffuse pulvinar nucleus, SFPul = superficial pulvinar nucleus, MPul = medial pulvinar nucleus, LPul = lateral pulvinar nucleus, IGPul = intergeniculate pulvinar.

I defined ROIs for the LGN, TRN, and pulvinar in both hemispheres using these anatomical images as reference; examples of both the 512 and 256 matrix acquisitions are shown in **Figure 2**. While the 512 matrix scans afford greater resolution the 256 matrix scans show the boundaries of the TRN well enough to reliably identify the boundary between the LGN and TRN. It is also apparent from these images that the exact boundaries of the structures are not always obvious, for example, where the TRN meets the lateral pulvinar in these images, or the boundary between the LGN and MGN. Furthermore, the utility of these ROIs is directly dependent on the accuracy of functional and anatomical image co-registration. Therefore these ROIs serve as a useful guide, and way to independently verify the performance of data driven methods, but should not be considered to be voxel-accurate representations of the anatomy.

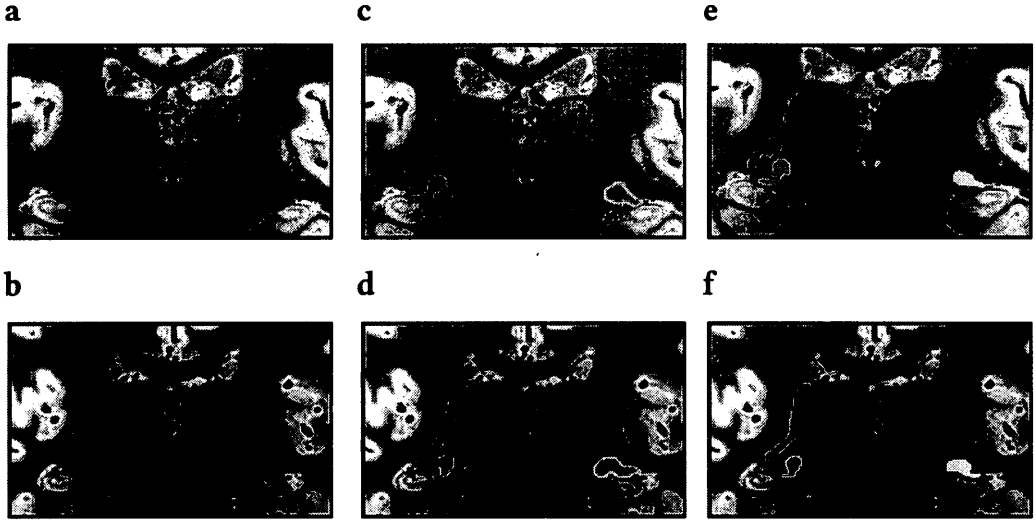

Figure 2. Anatomical masks of the LGN, TRN, and Pulvinar ROIs. a & b) Proton density (PD) weighted image showing the lateral geniculate nucleus (LGN), thalamic reticular nucleus (TRN), just dorsal-lateral to the LGN in S1 (512 matrix, top) and S4 (256 matrix, bottom). The pulvinar is shown dorsal to the LGN. c & d) Estimates of the structure's outlines in S1 and S4. Yellow = right LGN; Green = left LGN; Red = right TRN; Orange = left TRN; Purple = right pulvinar; Blue = left pulvinar. e & f) example masks derived from the estimated anatomy in S1 and S4.

2.2 Principals of Tremotopic Mapping and Stimulus Design

Electrophysiologists have long been interested in the resonance properties of isolated neurons, one of the reasons being they lend insight into the oscillatory

properties of neural networks (for a review, see Hutcheon and Yarom, 2000). The intrinsic resonance of a neuron at rest predicts its oscillation properties near threshold when driven, suggesting that the resonance properties of a neuron can bias it towards communicating with other neurons with similar oscillatory properties (Erchova et al., 2004). These experiments defined the impedance of a neuron using the so-called "ZAP" input — a current with a defined oscillatory waveform that would sweep through multiple frequencies (Puil et al., 1986). The output voltage recorded from the neuron driven at multiple frequencies over time will often show a resonant peak: a particular input frequency that produces the largest output voltage for a given neuron. In this way, neurons can be conceived of as band-pass filters that are most likely to respond to a driver with given oscillatory properties, as predicted by the communication through coherence hypothesis (Fries, 2005). I should clarify here that I am addressing a distinct issue from that regarding the neural code, which is believed to depend on more information than the phase-of-firing (Shadlen and Newsome, 1994).

EEG of the human suggests a similar phenomenon in the visual cortex in response to flicker (Herrmann, 2001). Specifically, they found resonance phenomena around 10, 20, 40, and 80 Hz. While the relationship between population-level oscillations as recorded by EEG and the resonance phenomena recorded at the single-neuron level is not direct, the choice of flicker frequency is known to modulate the BOLD response on a region-by-region basis. In the visual system, the LGN and V1 were found to respond preferentially to 7.5 and 20 Hz flicker, while the higher-order visual cortex (V2, V3, V3a, V4, and MT) responded preferentially to 20 Hz flicker alone (Kastner et al., 2004). One major limitation of this study was that only three flicker frequencies were investigated (0.5, 7.5, and 20 Hz), which did not allow for the experimenters to determine the fine flicker frequency tuning of these brain regions. Other studies have looked at the effect of flicker frequency more closely (e.g., D'Souza et al., 2011), but to the best of the author's knowledge, no study has investigated a wide range of flicker frequencies above 20 Hz. Most visual studies, especially in the thalamus, have made use of a single flicker frequency that best drives the primary visual regions

(e.g., Engel et al., 1997; Schneider and Kastner, 2005; Schneider et al., 2004). If there is indeed a relationship between visual flicker presentation and the resonance properties of neurons, one should be able to record this at the population-level non-invasively using MRI. This should allow us to better define structures based on their known resonance properties as established through invasive methods, and potentially functionally drive structures not previously studied in the human, such as the TRN.

To address this, I developed a stimulus that shares some properties with the aforementioned ZAP function, which I call a tremotopic stimulus. Conceptually, a tremotopic stimulus should periodically travel through retinotopic space while periodically sampling multiple flicker frequencies. These two periods should be different, and both should complete an integral number of cycles per scan. For technical reasons explained later, these periods should be relatively far apart, and should be numbers approximately larger than 6 to prevent the signal of interest from being

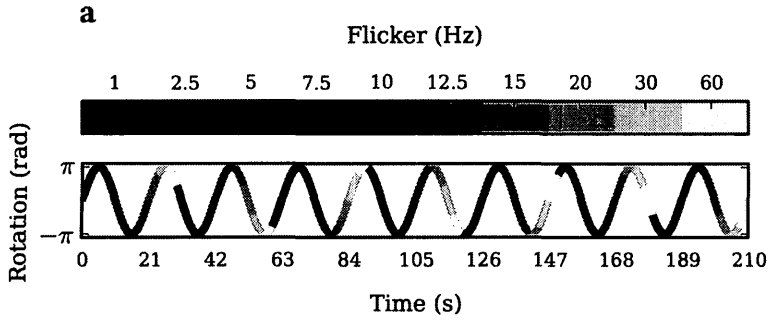

Figure 3. The tremotopic stimulus. a) Stimulus properties, demonstrating the superposition of flicker modulation and stimulus rotation throughout each run. b) A rendering of the stimulus as presented.

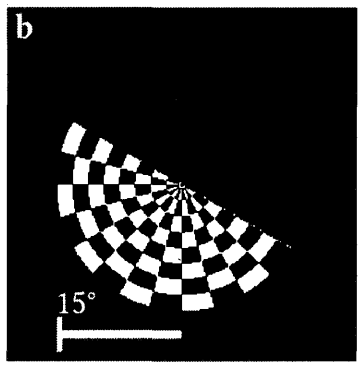

corrupted unduly by low-frequency drifts caused by MR scanner instabilities (Smith et al., 1999). Visually driven voxels should show sinusoidal BOLD modulation at the retinotopic frequency, and if they also express a flicker frequency preference, this

sinusoidal modulation should express an envelope reflecting the magnitude of a voxel's response to a given flicker frequency. Therefore, this should produce amplitude modulated BOLD signals in visually-driven voxels that also possess a flicker frequency tuning.

Specifically, the stimulus (**Figure 3**) consisted of a rotating hemifield containing a radial checkerboard spanning 15° of visual angle and an effective resolution of 100×100 pixels. This checkerboard consisted of pure white and black checkers presented on top of a mid-grey background, with a stationary fixation point presented in the center of the display for the duration of each run. The checkerboard rotated 360° every 21 s, a total of 10 times per scan. The checkerboard also stair-cased through 10 flicker frequencies (defined as a full contrast-reversal cycle) every 30 s, 7 times a scan [Hz = 1, 2.5, 5, 7.5, 10, 12, 15, 20, 30, 60]. The stimulus included a 10 s rewind to allow for T1 stabilization, and the first 10 s were therefore identical to the last 10 s of each run, to allow us to safely discard the first 10 s of data.

2.3 Collection and Pre Processing of Functional Data

For all functional scans, participants were instructed to maintain fixation and not move throughout each run, which lasted 220 s. The stimulus was presented on an InFocus IN112 3D DLP projector allowing for rapid pixel offset times and finer control of flicker frequencies due to the 120 Hz refresh rate. The stimulus was projected onto a plastic screen hung from inside the head of the magnet bore, which participants viewed through a head-mounted mirror.

Functional imaging in the thalamus was accomplished using echo-planar imaging (EPI) sequence consisting of 12 coronal slices at a 128×128 resolution matrix with 2 mm thickness and no gap between slices leading to an in-plane resolution of 1.5 mm [TR = 1250 ms, TE = 41 ms, flip angle = 90°; partial Fourier factor = 7/8]. These slabs were placed such that they covered the posterior thalamus including the LGN, TRN, posterior pulvinar, MGN, and SC. A whole-head EPI image with the exact same imaging parameters save an extended TR was also acquired for registration purposes.

The utility of this image was that is preserved the exact EPI distortions present on the particular scanning day, facilitating co-registration with the high-resolution T1-weighted anatomical image. Each run consisted of 176 TRs spanning the full 220 s of stimulus presentation.

I chose this relatively low TR length in order to balance the conflicting demands of appropriate brain coverage and reasonable sampling of cardiac and respiratory-induced signals in the BOLD signal, which start around 0.15 Hz and extend well beyond the Nyquist frequency of these scans (Ash et al., 2011; Liu et al., 2006; Triantafyllou et al., 2005). Since both stimulus frequencies fell well below 0.15 Hz (~0.048 and ~0.33 Hz, for retinotopy and flicker modulation, respectively), this allows for the data to be low-passed without interfering with our signals of interest, with the intention of removing the major contributions of both physiological and high-frequency MR sequence noise.

Functional data was pre-processed in AFNI. The first 10 s of data (8 TRs) were then deleted from each run (3dCalc), leaving 168 time points for each run. Next, I slice time corrected each run to align the data in time, as slices in EPI data are collected sequentially throughout each TR (3dTshift). These slabs were then co-registered with the T1 anatomical and resliced to 0.75 mm^3 in one step (3dWarp, 3dVolreg, and align_epi_anat.py). I recorded the following six rigid-body realignment parameters from 3dVolreg to later analyze head motion on a run-by-run basis: shifts in the *x*, *y*, *z* planes, and rotations roll (A), pitch (B), and yaw (Γ). Large, non-physiologically relevant spikes in the functional data were removed by fitting a curve to each time series using L1 regression, defined by

$$f(t) = a + bt + ct^2 + \left\{\sum_{k=1}^{k=L} d_k \sin\left(\frac{2\pi kt}{T}\right) + e_k \cos\left(\frac{2\pi kt}{T}\right)\right\} \quad [1]$$

where T is the duration of the time series, a, b, c, d, and e are parameters chosen to minimize sum of the error over t, and L is the number of time points divided by 30, which is 168 in this case. Next, the mean absolute deviation (MAD) was calculated of the difference between this curve and the data, and I estimate the standard deviation

(SD) of the residuals as $\sigma = \sqrt{\frac{\pi}{2}MAD}$, and replace values in the time series that are greater than 2.5 SDs from the fitted curve with values from the time series

$$s(t) = 4\tanh\frac{SD - 2.5}{1.5} \quad [2]$$

individually for each voxel (3dDespike). I next removed the linear and quadratic trend from each time series and recorded the resulting time series mean (3dDetrend) before subtracting the mean from each time series and dividing by it, converting each time series to units of percent signal change (3dCalc). Next, I fit each time series with the six head motion parameters estimated by 3dVolreg using linear least squares and retain the residuals to account for BOLD signals arising from head motion (3dDetrend; Van Dijk et al., 2010). Finally, I spatially smoothed each run with a Gaussian kernel with a 2 mm full-width half-max (FWHM).

I further pre-processed each dataset using custom software written in the MATLAB programming language (version 7.12). First, I band-passed each time series before producing two sets of means: one high-passed run retaining frequencies $0.009 < f < 0.4$ Hz, to remove only the lowest frequencies attributable to MR scanner drift, and the other band-passed run retaining $0.009 < f < 0.08$ Hz, as is done in resting-state analysis (Van Dijk et al., 2010). Both were accomplished with a dual-pass bi-directional digital FIR Kaiser filter with an order of 55 and beta value of 2.5. The use of this bi-directional filter ensures that no phase lag is introduced into the signal by running the filter once in the ascending direction (from $t = 1$ to $t = n$), and again in the negative direction (from $t = n$ to $t = 1$). In both the high-passed and band-passed case, I also folded the time series over on itself to produce a mean flicker cycle (168 TRs / 7 cycles = 24 TR/cycle). More details on their utility will be described later.

As local time series correlations can be generated by small head motions even after the regression of the estimated head motion parameters, I next censored TRs in

each run in the following way: I estimated frame-wise displacement (FD) from the six motion parameters as

$$FD_i = |\Delta d_{ix}| + |\Delta d_{iy}| + |\Delta d_{iz}| + |\Delta A_i| + |\Delta B_i| + |\Delta \Gamma_i| \quad [3]$$

where $\Delta d_{ix} = d_{(i-1)x} - d_{ix}$, and identically for $[d_{iy}\ d_{iz}\ A_i\ B_i\ \Gamma_i]$ (Power et al., 2012). I converted from degrees rotation to millimetres displacement assuming the radius of each participant's head was 50 mm. I then masked every TR associated with an FD over 0.25 mm, as well as one frame back and two frames forward from these TRs to allow for the re-establishment of steady state spins. This procedure is normally used in resting-state experiments, which allow for the complete deletion of individual TRs. In the present stimulus-driven paradigm, however, I have the option to average time-locked signals. I therefore chose to take a frame-wise average across runs, using only those TRs that survive the threshold procedure. This allows one to generate an averaged time

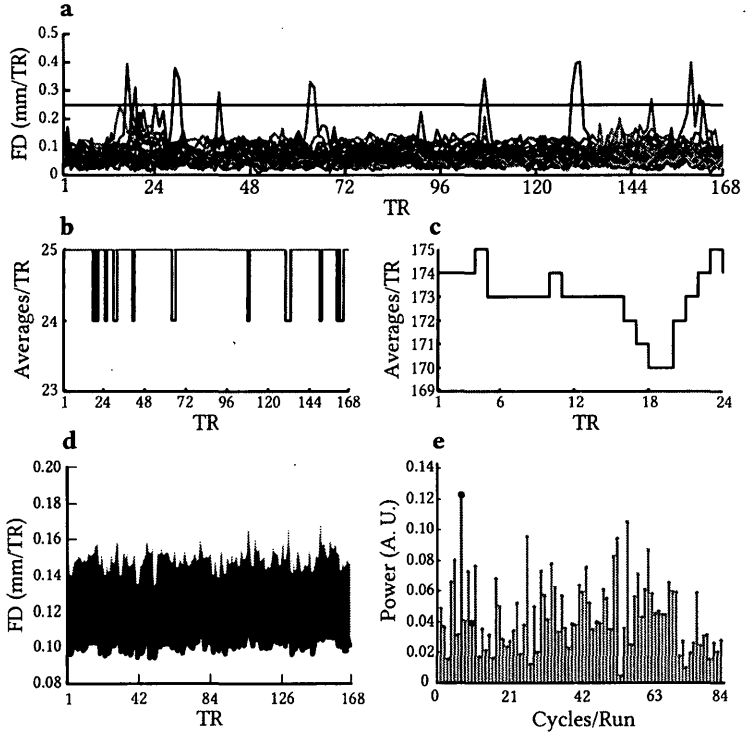

Figure 4. Analysis of head motion. a) A plot showing the frame-wise displacement (FD) for each run collected from S1 and the threshold of 0.25 mm/TR. b) A plot showing the number of averages per TR for S1 over the course of the experiment. c) A plot showing the number of averages per TR for S1 over the course of a flicker cycle. d) A plot showing the mean and SD of the group-averaged FD and absolute SD. e) A periodogram of the group-mean FD, with 7 cycles and 10 cycles per scan in red, corresponding to the stimulus flicker and rotation cycles per run, respectively.

series with no head displacements over 0.25 mm at any given TR. This procedure generally removed a relatively small number of TRs from each participant's total, as is

demonstrated by a representative subject (S1) in **Figure 4a. Figure 4b,c** show the number of averages per TR for the full run and mean cycle, respectively.

The removal of the head motion parameters in this way is important, as it removes any signal arising from the participant moving their head in synchrony with the stimulus. To test whether there was any group-level consistency in head motion, I calculated the FD for all runs across subjects and calculated the resulting mean and SD. I then calculated the periodogram of this average FD time series; both are shown in **Figure 3d,e.** At the group level, there were significant modulations of the FD at seven cycles per scan, which corresponds to the rate of flicker modulation ($F_{2,83} = 5.17$, $p = .0078$), but not at the rate of stimulus rotation ($F_{2,83} = 0.155$, $p = .86$).

Finally, I created two sets of masked averages for each set of band-passed inputs: one of the full run (168 TRs), and one of the mean flicker cycle (24 TRs). Because both the retinotopic and flicker signals are integral within each run, any retinotopically-driven sinusoidal modulation of the BOLD signal should cancel when I average over the mean flicker cycle. Therefore, the mean flicker cycle is a measure of the frequency preferences for a given voxel. I will refer to this mean cycle as the frequency profile. Finally, I masked these data sets with a thalamus mask constructed from the Freesurfer segmentation process, removing voxels belonging to the surrounding white matter and ventricle masks. An example thalamus mask is shown in **Figure 5.**

Figure 5. An example Freesurfer-derived thalamus mask (red overlay) from S1. This mask was mainly used to limit computation times.

Chapter 3

The Attempts Made to Identify the Human Thalamic Reticular Nucleus

Here I describe the various analysis preformed on these data to identify the functional properties of the human thalamic reticular nucleus. Each subsection is presented as a self-contained experiment with associated method, result, and discussion. In Section **3.1**, I describe my attempt to apply standard traveling-wave techniques and a Fourier analysis to define retinotopic regions and flicker preference in the thalamus. In Section **3.2**, I introduce the utility of graph theory in tackling the present problem and define subregions of the thalamus with particular flicker-frequency preferences using a community detection algorithm. I then further apply graph-theoretic techniques in an attempt to find a network property unique to the TRN in the human thalamus. In Section **3.3**, I utilize the seed-based correlation approach to look for evidence of TRN-mediated lateral-inhibition in the human thalamus.

3.1 Tremotopic Mapping of the Human Thalamus

I defined visually-driven voxels as those that strongly modulate at the same frequency as either the rotation or flicker-modulation frequency of my stimulus. Retinotopy studies utilizing the travelling-wave paradigm have traditionally utilized the analysis proposed by (Engel et al., 1997) which finds the maximum correlation value of the stimulus harmonic as the amplitude of the Fourier component at frequency f divided by the square root of the time series power. In this case, this analysis is not appropriate for the data: only the largest response would be detectable in a voxel that expresses a large-amplitude retinotopic and moderate flicker response, or vice versa. To avoid this, I implemented a modified version of the local spectra-based F-test for periodicities proposed by (Wei and Craigmile, 2010). I used the adaptive multi-taper method as implemented in the MATLAB 7.13 to estimate the power spectral density for each time series with three orthogonal windows (Thomson, 1982). This method provides a more accurate estimation of the noise variance than the estimation of the power spectral

density multiplied by a rectangular window, which introduces a large amount of variance into the data due to Gibbs phenomenon. An accurate estimation of the noise variance is required to ensure accurate *F*-values, in order to maximize delectability of visual regions.

Briefly, the modified local spectra-based *F*-test assumes that each signal is composed of Gaussian white noise with zero mean and a variance greater than zero, some known nuisance signals at particular frequencies, and a test signal at a frequency of interest. I conducted this test on the high-passed run, as this view of the data affords more degrees of freedom for the significance test. I defined the two stimulus frequencies and their first two harmonics as nuisance frequencies. Specifically, this procedure tests whether the amplitude of the test frequency is not zero using a regression-based *F*-test against a model of the noise process. I evaluated the estimated spectral power $\hat{S}_X^{(p)}$ separately at the retinotopic and flicker modulation frequencies (7 and 10 cycles per scan, respectively) f_ζ divided by the average power of the frequencies $[f_{k,\mathrm{N}};\ k = 1, \ldots, m]$, where m is the Nyquist frequency and $f_{k,\mathrm{N}} \neq f_{k^*,\mathrm{N}};\ k^* \in \{7, 10, 14, 20, 21, 30\}$, as

$$F(\xi) = \hat{S}_x^{(p)}(f_{k_\xi,N}) / \left\{ \sum_{k \neq k^u} \hat{S}_x^{(p)}(f_{k,N}) / m - n \right\} \qquad [4]$$

with n being the number of nuisance frequencies, and N the number of TRs.

This yields an *F*-statistic that can be tested against the non-central *F*-distribution on 2 and $2(m - n)$ degrees of freedom to produce a series of *p*-values. In this case, the number of tests was large per participant (~ 125,000 voxels), therefore, the number of expected false-positives for each participant was expected to be ~6250 voxels with a critical $p = .05$. Given the goal of identifying a structure with an ambiguous physical location in the brain, this is unacceptable. To correct for multiple comparisons, I adjusted the critical *p*-value using the false discovery rate (FDR; Benjamini and Hochberg, 1995) and define the maximum number of false positives among all discoveries q to be 5% (~ 350 voxels per participant), yielding a corrected *p*-value

threshold for each participant. This procedure is more appropriate in this case than stringent family-wise error rate corrections such as the Bonferroni method, as the small number of false positives should not be consistent among participants, while a large number of false negatives may not allow for the identification of the structure of interest (Lieberman and Cunningham, 2009). I further refined these maps by only retaining significant voxels participating in a contiguous cluster of at least 9 voxels, under the assumption that clusters smaller than this are insignificant or arise due to noise. Considering that the data is resampled to produce ~10× the number of voxels (original voxel resolution = 1.5 mm^2 × 2 mm, or 4.5 ml; resliced voxel resolution = 0.75 mm^3, or ~0.42 ml), it seems appropriate to remove the influence of signal arising from a volume smaller than 1 voxel in the original data. The results of the multiple-comparisons corrected tests for the retinotopy and flicker analysis can be found in **Figure 6** and 7, respectively. These figures are organized such that panel **a** falls on the anterior LGN, and panel **d** falls on the posterior SC, covering most of the visual thalamus.

The retinotopy analysis revealed extensive visual responses in the thalamus. Statistical map threshold of q[FDR] = 0.05 resulted in critical *p*-values ranging .0014–.0049 ($\bar{x}$ = 0.0034, SEM = 0.00053). The data shows the expected retinotopic organization in the LGN, intergeniculate pulvinar, lateral pulvinar, and SC in most, but not all, subjects (Schneider, 2011; Schneider and Kastner, 2005; Schneider et al., 2004). In the interest of brevity, I will only list activated ROIs outside of the LGN-TRN-pulvinar complex, and reference the pulvinar as a contiguous whole. Of note, one can observe reliable drive of all primary visual nuclei (LGN, SC, intergeniculate pulvinar) by the contralateral visual field, as expected. Surprisingly, the MGN also shows strong visual drive in multiple participants, which has not, to the best of the author's knowledge, yet been reported in the human thalamus. It is not completely unexpected, however, as multiple cortical auditory regions are known to respond to visual stimulation within 110 ms, and the BOLD signal here may be driven by feedback from the cortex (Beauchamp

et al., 2004a, 2004b; Logothetis and Wandell, 2004; Proverbio et al., 2011). Furthermore, multiple regions of the dorsal pulvinar respond to the stimulus, are area known to have

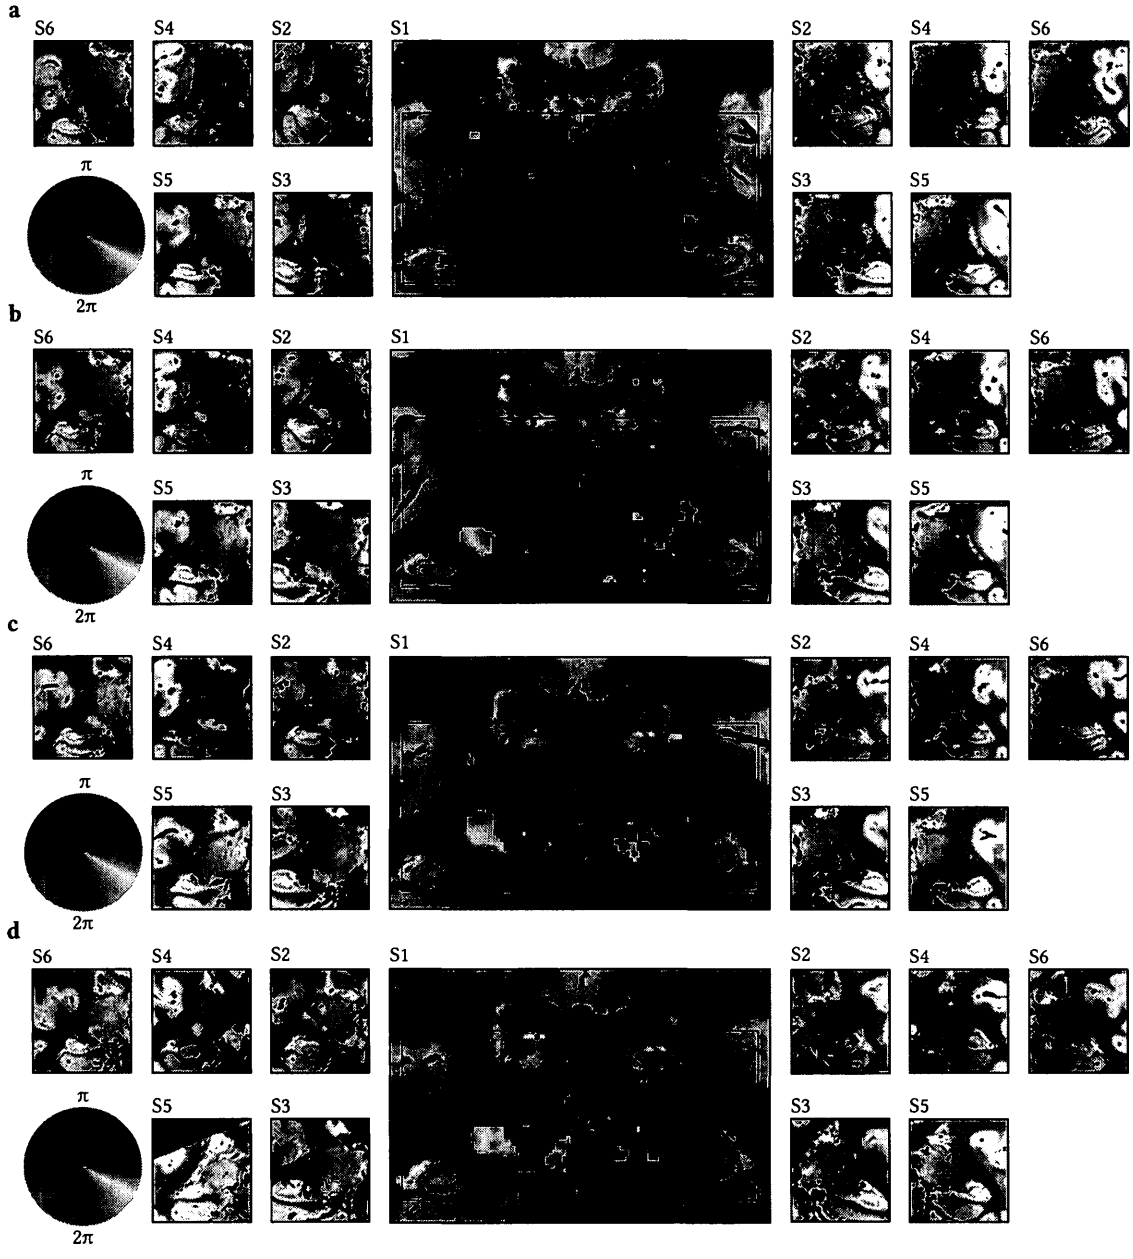

Figure 6. Retinotopy of the thalamus. All maps thresholded at q[FDR] = 0.05. Slices move anterior to posterior from panel a to d. The data is presented with natural left and right. a,b) Panels demonstrating retinotopy of the LGN, intergeniculate pulvinar, lateral pulvinar, and MGN. Of note, primary sensory nuclei represent the contralateral visual field. c,d) Panels demonstrating retinotopy of the above, medial pulvinar, dorsal pulvinar, SC, and TRN. TRN retinotopy can be found in the following panels: a) S2 R, S3 R, S4 R, S6 R; b) S2 R&L, S3 R, S4 R, S6 R; c) S1 L, S2 L, S3 R, S4 R, S5 R, S5 R&L; d) S2 R&L, S3 R, S4 R&L, S5 R, S6 L.

widespread audio-visual connectivity in the brain, as well as dense connectivity with the pre-frontal cortex in the monkey (Gutierrez et al., 2000).

Multiple regions of the TRN respond in multiple participants just dorsal-lateral to the medial LGN, and in multiple locations along the lateral pulvinar. Upon visual inspection, it isn't clear from this analysis that the TRN is reliably driven by the present stimulus, because these significant ROIs are generally lateralized to the right and appear to fall in inconsistent locations along the TRN. Furthermore, due to the variability in the relative locations of these thalamic structures, moving these data to a standard space is not feasible. The spatial normalization algorithms required to move the data into standard space do not respect the spatial fidelity required to delineate between the LGN and TRN, rendering them suboptimal for high resolution functional studies of this nature.

The data demonstrating flicker-frequency preference in each structure also shows consistency across participants, but this analysis was less sensitive. Statistical map threshold of q[FDR] = 0.05 resulted in critical p-values ranging .0014–.0056 ($\bar{x}$ = 0.0024, SEM = 0.00061). This analysis shows that the tremotopic method does indeed allow one to probe the resonance properties of various neural populations. The LGN is resonant at, and responds to, flicker frequencies surrounding 10 Hz stimulation, mirroring electrophysiology experiments (Hughes et al., 2004). This result also replicates the fMRI finding of (Kastner et al., 2004), demonstrating preferential LGN activity when presented with 3.75 or 10 Hz flicker, when compared with the responses elicited by 0.25 Hz. The MGN and dorsal pulvinar are seen to respond primarily to 20–30 Hz flicker. The SC is only detectable in a few participants, but in those that it is, it expresses a 5–10 Hz flicker tuning. Very little of the TRN expresses flicker preference when analyzed in this way. However, considering that this analysis is significantly less sensitive than the retinotopy analysis (identifying only 71.9% of the number of voxels identified by the retinotopic analysis, or 38755 of 53837), it is possible that any flicker modulation of the TRN is not robustly detectable by this analysis. To overcome this

limitation, I produced a third set of maps demonstrating the flicker preference of each significantly retinotopically-modulated voxel, shown in **Figure 8**.

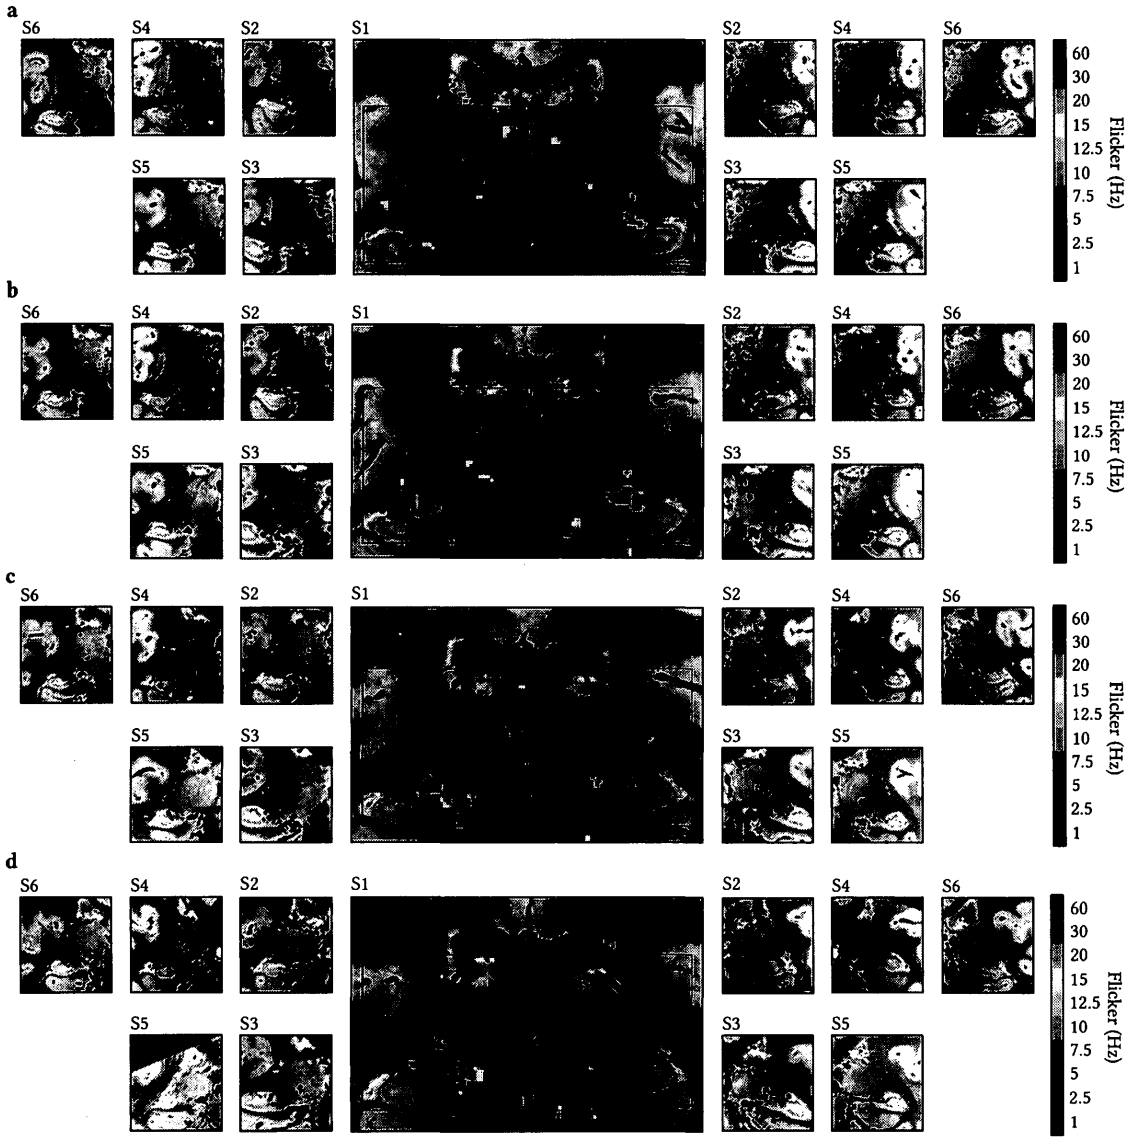

Figure 7. Flicker responses of the thalamus. All maps thresholded at q[FDR] = 0.05. Slices move anterior to posterior from panel a to d. The data is presented with natural left and right. a,b) Panels demonstrating flicker responses to the LGN, intergeniculate pulvinar, and dorsal pulvinar. c,d) Panels demonstrating flicker preferences in the MGN, lateral pulvinar, and SC. Primary visual nuclei prefer flicker frequencies in the α-frequency range, while most pulvinar regions and the MGN prefer β-frequency flicker.

Qualitatively, these data suggest a delineation of the thalamic nuclei into two groups: low-frequency responses in first order thalamic nuclei, and high-frequency responses in second order thalamic nuclei, as predicted (Hughes et al., 2004; Wang et al., 2012; Wróbel et al., 2007). In order to better quantify the functional properties of the

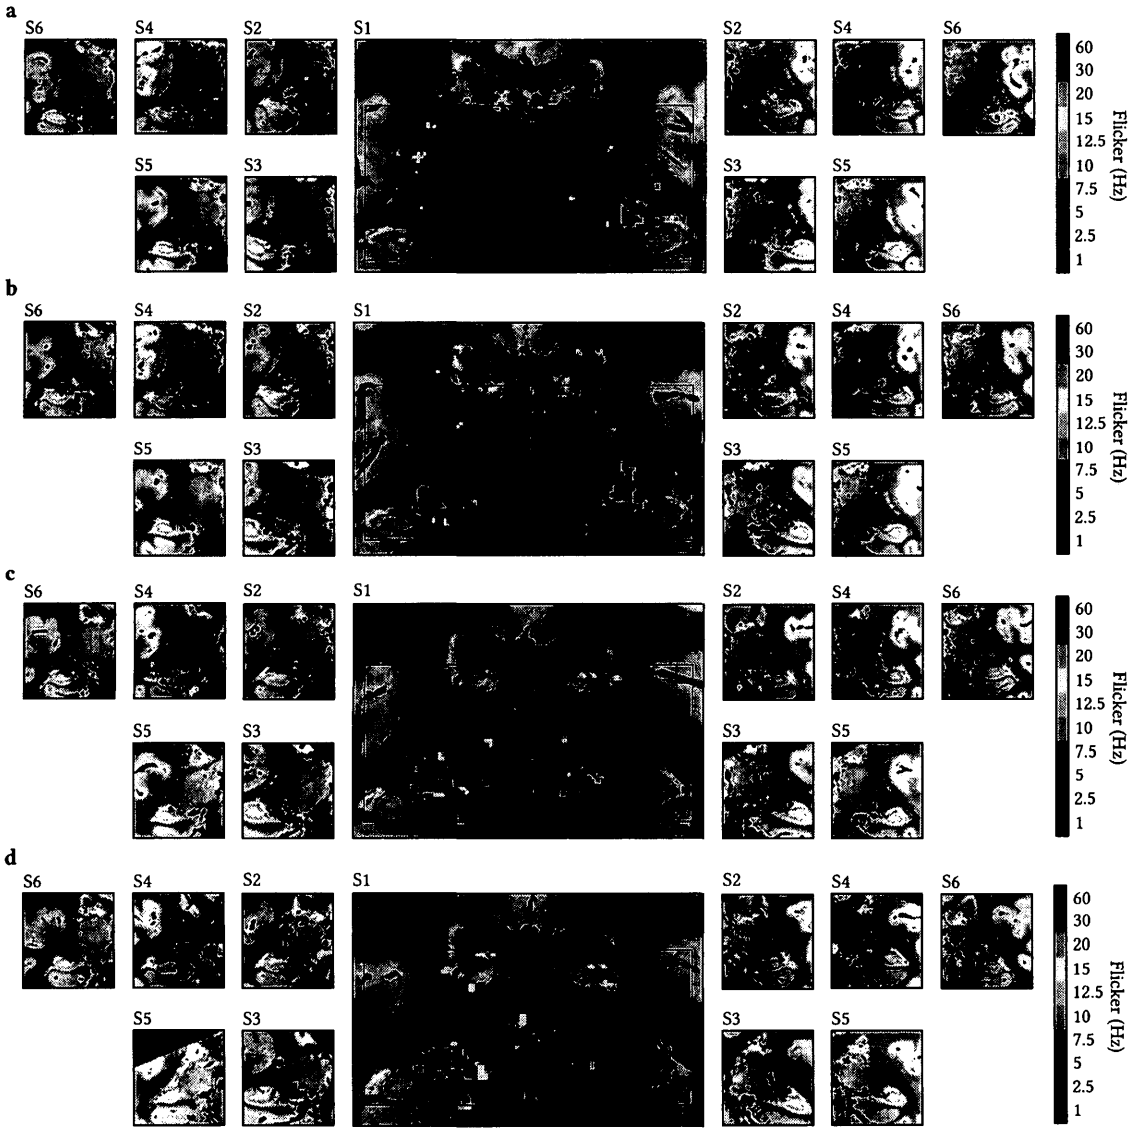

Figure 8. Flicker responses of the thalamus, with the retinotopy threshold. Data thresholded at q[FDR] = 0.05 by rotation frequency response. Slices move anterior to posterior from panel a to d. The data is presented with natural left and right. This analysis shows multiple regions of both high and low flicker frequency responses in the TRN. Both high (12.5–60 Hz) flicker frequency preferences can be seen in the dorsal pulvinar. TRN flicker preferences can be found in the following panels: a) S2 R, S3 R, S4 R, S6 R; b) S2 R&L, S3 R, S4 R, S6 R; c) S1 L, S2 L, S3 R, S4 R, S5 R, S5 R&L; d) S2 R&L, S3 R, S4 R&L, S5 R, S6 L.

visual thalamus, and perform a group-level analysis, I pooled both the retinotopy and flicker phase statistics from the Fourier analysis across subjects within the defined proton-density ROI masks: the left and right LGN, TRN, and pulvinar, respectively (**Figure 9**). I divided the visual field into 36 even segments, and bin the flicker phase according to the 10 flicker frequencies presented. I conducted a series of Kolmogorov-Smirnov (KS) tests to compare the response properties of these ROIs. The distribution of polar angle responses are significantly different between ROIs (LGN vs. TRN: D_n = 0.126, $p = 1.2\times10^{-8}$; TRN vs. pulvinar: D_n = 0.167, $p = 1.07\times10^{-16}$; LGN vs. pulvinar: D_n = 0.170, $p = 2.9\times10^{-61}$), and similarly for the flicker frequency preference (LGN vs. TRN: D_n = 0.450, $p = 3.6\times10^{-62}$; TRN vs. pulvinar: D_n = 0.275, $p = 3.61\times10^{-28}$; LGN vs. pulvinar: D_n = 0.506, $p = 3.73\times10^{-294}$).

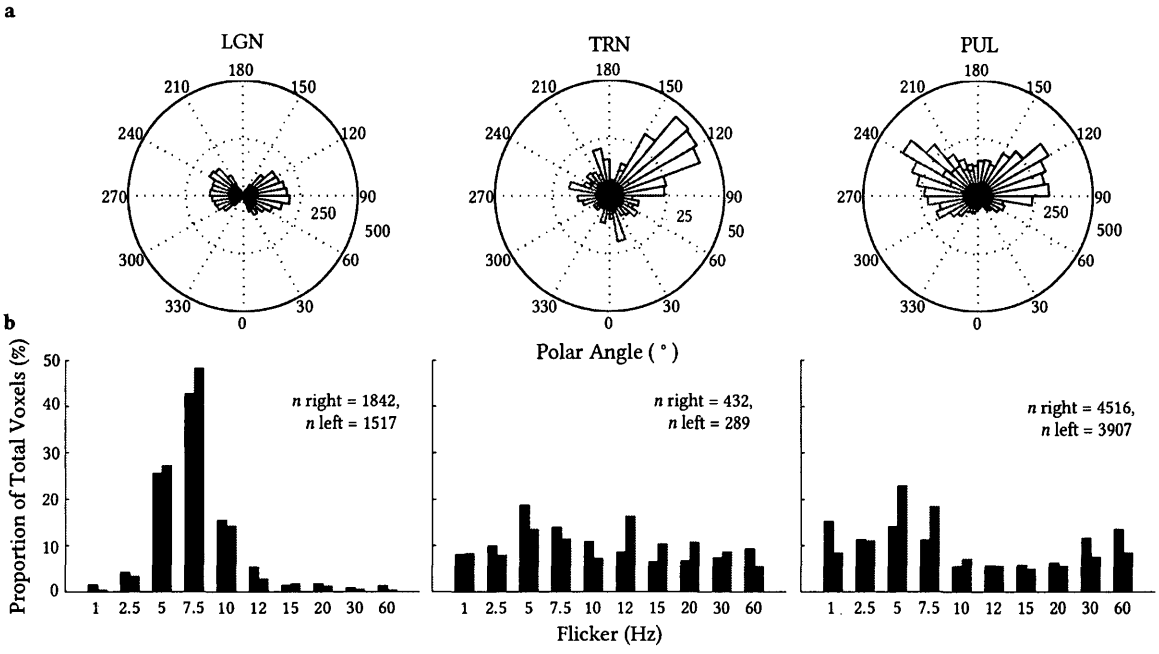

Figure 9. Group-level retinotopy and flicker responses of the LGN, TRN, and pulvinar for the thalamus-masked analysis. Black and red represent the right and left ROIs, respectively. a) Polar angle distributions from the LGN, TRN, and pulvinar ROIs. The visual field was segmented into 36 bins for visualization. b) Flicker frequency distributions from the LGN, TRN, and pulvinar ROIs.

This analysis clearly shows the expected flicker frequency tuning of the LGN. Furthermore, the distribution of polar angle representations replicates the finding of (Schneider et al., 2004), including the under-representation of the vertical meridian. This

is of significance, as this analysis was restricted to a well-defined anatomical mask. In previous work, the LGN ROI was defined using a functional localizer, and was found to be approximately 2× the volume of the human LGN as measured during post-mortem studies. These data suggest that the functional under-representation of the horizontal meridian is not an artefact of the reliance on functional ROIs, and generally that these two methods for defining ROIs produce similar results regardless of the extent of BOLD activation. The pulvinar ROI also shows an overrepresentation of the upper visual field, and minor contributions of the pulvinar on each side of the brain to both the contralateral and ipsilateral space, replicating previous animal studies (Bender, 1981). This overrepresentation of the upper visual field may suggest that at least some portions of the pulvinar are preferentially involved in the linear search of extra personal space: there is evidence that the upper and lower visual fields are specialized for processing distant and near objects, respectively (Previc, 1990). Recent evidence lends showing cross-frequency coupling between α-and β-frequency LFPs in the pulvinar and temporal-occipital cortex during an attention task lends credence to this theory, suggesting the pulvinar is required for the integration of information across the senses (Saalmann et al., 2012). Interestingly, this data shows the pulvinar has populations with flicker-resonances in both the α-and high-β ranges.

The TRN ROIs shows both ipsi and contra-lateral representations of the full visual field in both ROIs, although the number of driven voxels is very small. To the best of my knowledge, this has not been shown before, but this result is also not wholly unexpected. The TRN participates in a thalamo-TRN-thalamo circuit which is hypothesized to be a thalamic lateral inhibitory circuit (Deschênes et al., 2005; Pinault and Deschênes, 1998). In this case, the BOLD representation of the contralateral visual field may be driven by inhibitory inputs during visual drive of the contralateral LGN. This data also shows a broad distribution of flicker-frequency preferences in the TRN, with slight over-representations in the α-and high-β ranges, as seen in the pulvinar. The data also suggests a large preferential drive of the right TRN. It is not clear whether this finding is an artefact of smoothing between the pulvinar and right TRN, which lie side-

by-side. This explanation seems plausible as it has the same visual field representation properties as the right pulvinar, preferring the upper visual field. An alternative explanation would be that, like all structures studied here, one hemisphere of the brain

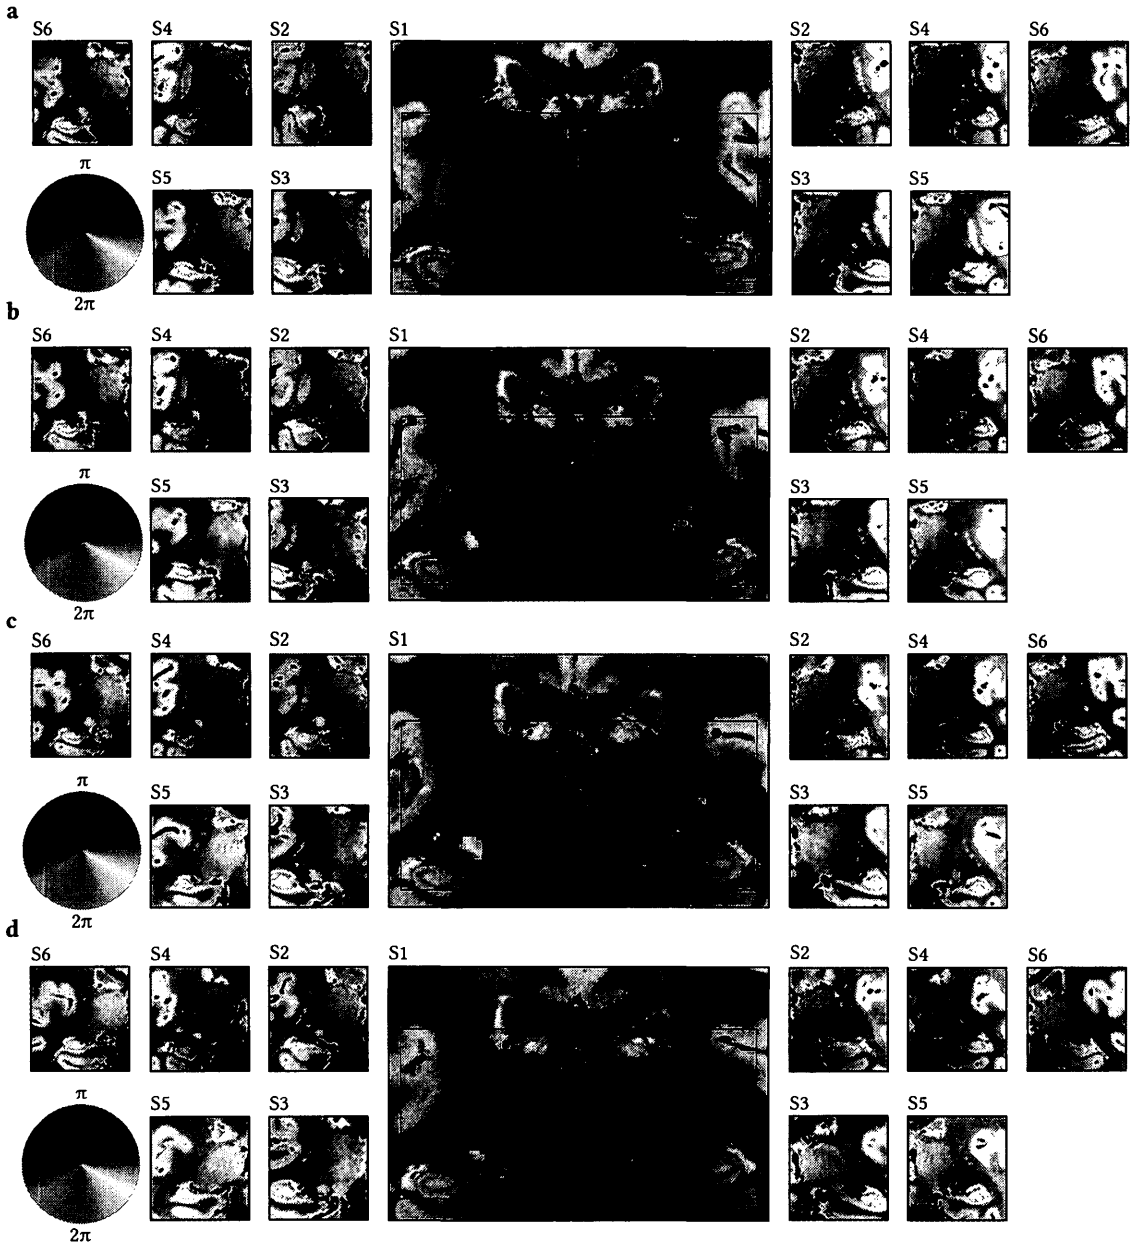

Figure 10. Retinotopy of the LGN-TRN complex. All maps thresholded at q[FDR] = 0.05. Slices move anterior to posterior from panel a to d. The data is presented with natural left and right. These data demonstrate both contra and ipsilateral retinotopic activation in the TRN in both hemispheres.

tends to respond more strongly to the stimulus, and given the extreme difficulty in detecting TRN activity with BOLD, I may not have sufficient power to elicit a response from the (on average) weaker left TRN.

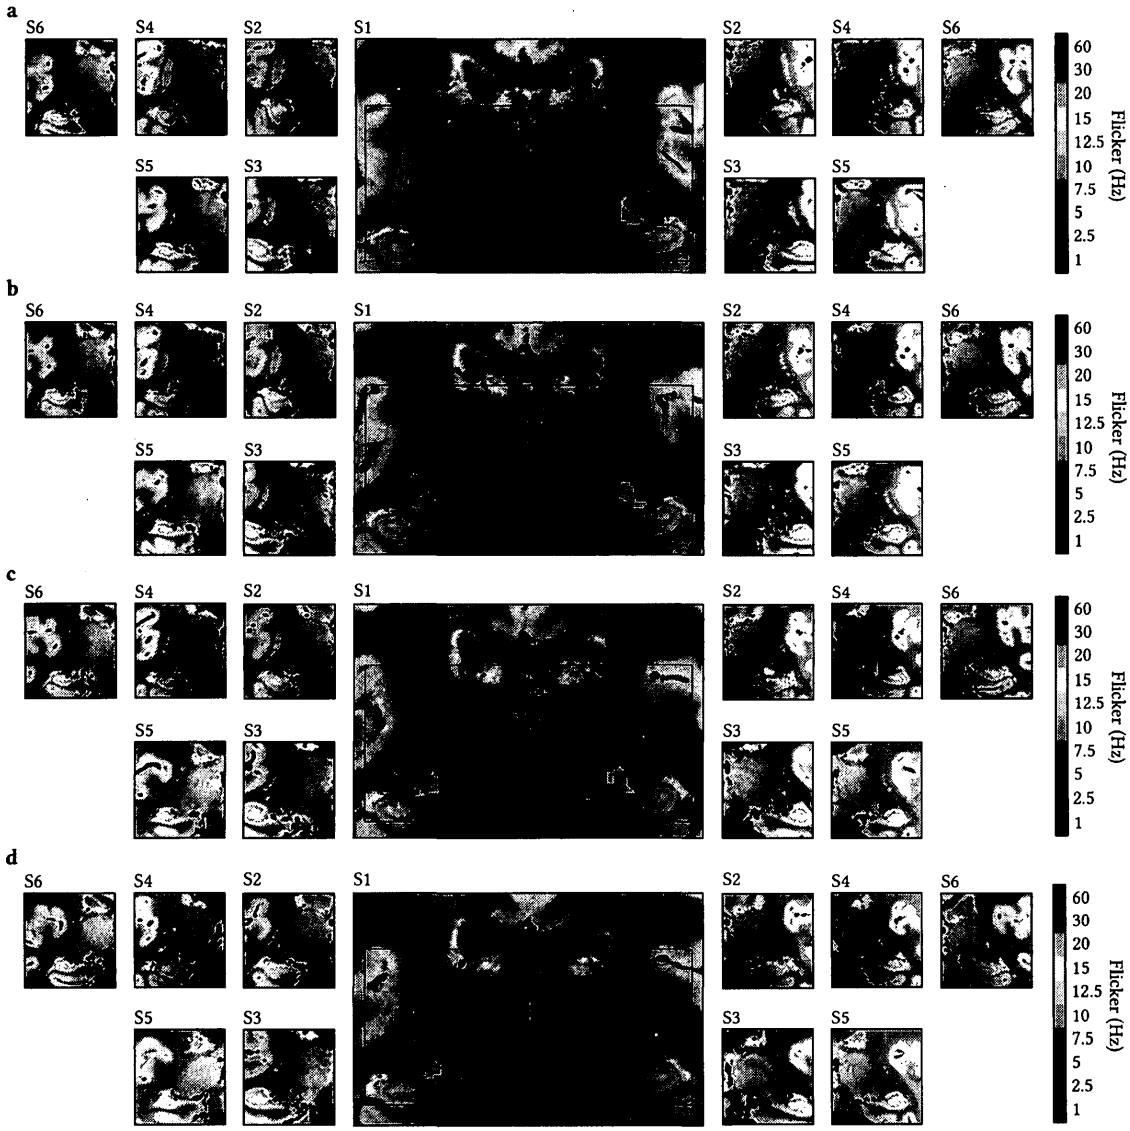

Figure 11. Retinotopy of the LGN-TRN complex. All maps thresholded at q[FDR] = 0.05. Slices move anterior to posterior from panel a to d. The data is presented with natural left and right. These data demonstrate the broad flicker frequency tuning of the TRN.

One way to test this would be to re-run the analysis using only the LGN and TRN ROIs, substantially reducing the number of comparisons made and therefore increasing the statistical power of the analysis. Furthermore, no cluster-size threshold

was imposed on these analysis. The results of the retinotopy, flicker preference masked by retinotopy, and group-level analysis are shown in **Figures 10**, **11**, and **12**, respectively. Statistical map threshold of q[FDR] = 0.05 resulted in critical p-values ranging .0057–.015.

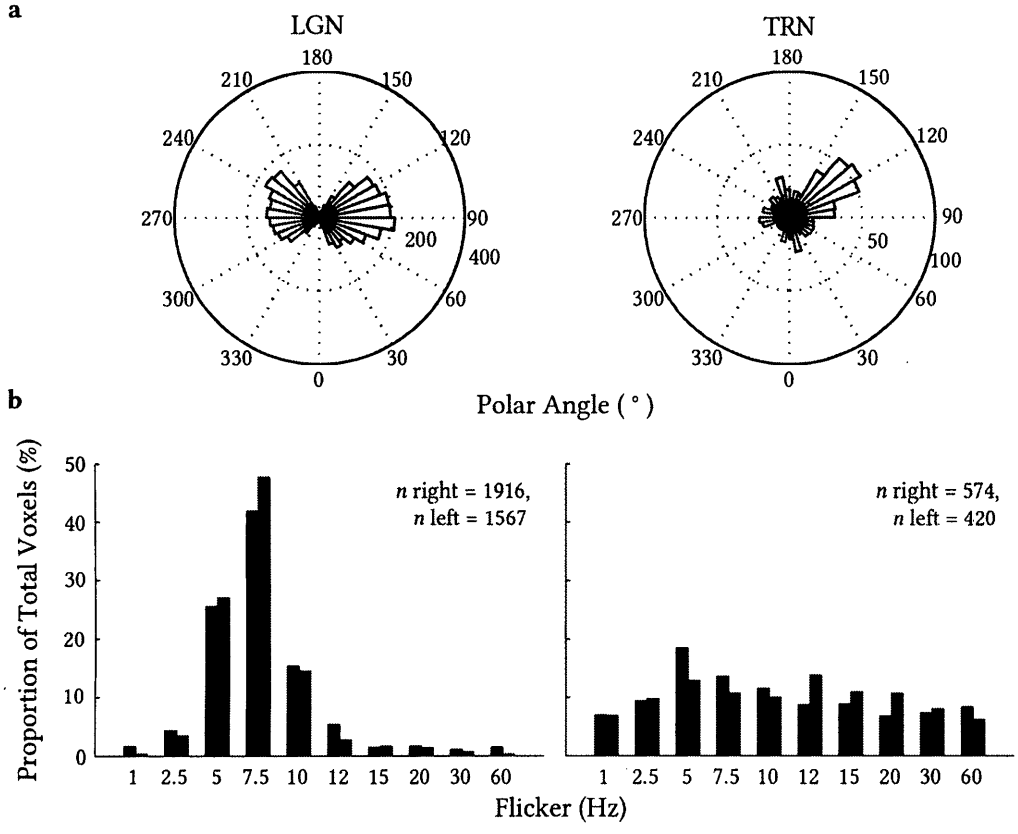

Figure 12. Group-level retinotopy and flicker responses of the LGN and TRN for the TRN-LGN masked analysis. Black and red represent the right and left ROIs, respectively. a) Polar angle distributions from the LGN and TRN ROIs. The visual field was segmented into 36 bins for visualization. b) Flicker frequency distributions from the LGN and TRN ROIs.

Figure 10 and 11 demonstrate two properties of the TRN: a full representation of the visual field in each hemisphere, and a broad flicker frequency tuning. Figure 12 shows that these findings are consistent at the group level. The distribution of polar angle responses are significantly different between the LGN and TRN (KS test: $D_n = 0.113$, $p = 4.8\times10^{-9}$), and similarly for the flicker frequency preference ($D_n = 0.356$, $p = 5.1\times10^{-86}$). This analysis indentified substantially more voxels in the TRN, especially in the weaker

left hemisphere (number of voxels identified: right 1.33×, left 1.45×). Interestingly, the over-representation of the upper-left visual field remained relatively constant. To determine whether this is a simple artefact of smoothing, I re-ran this analysis on unsmoothed data. The results of this analysis can be found in **Figure 13**.

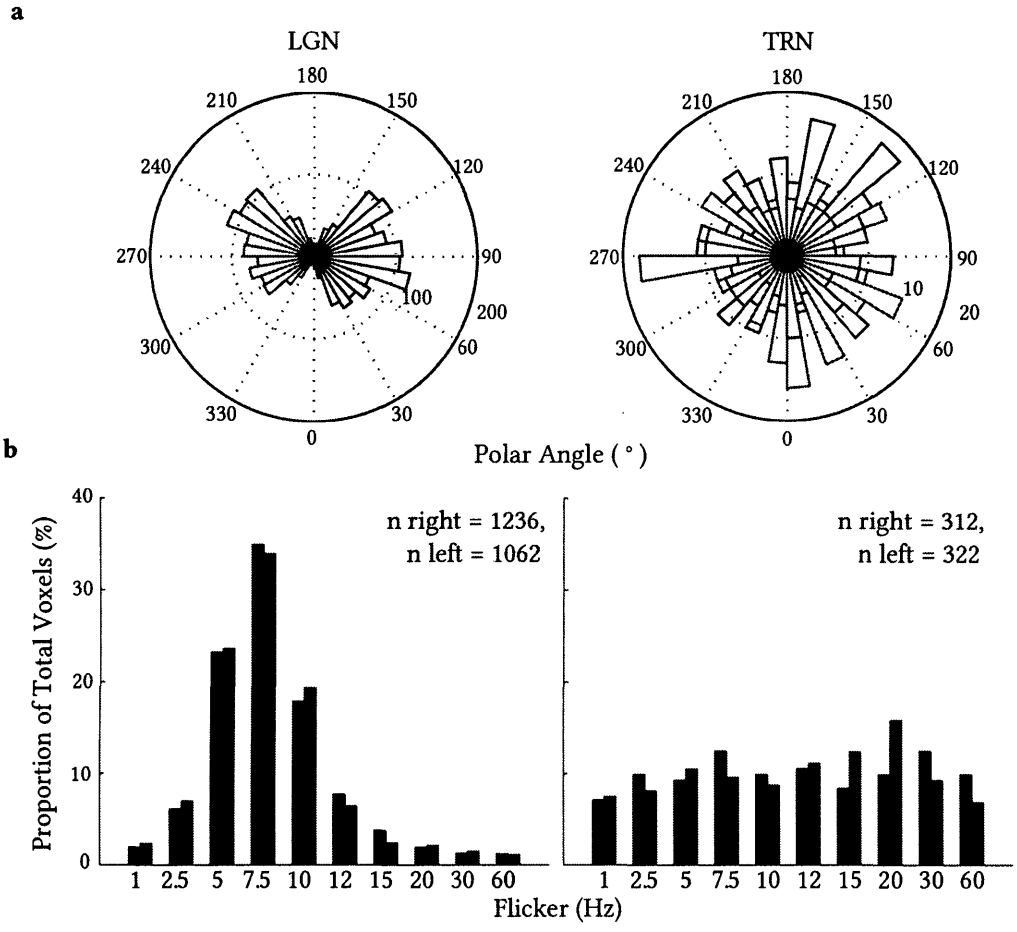

Figure 13. Group-level retinotopy and flicker responses of the LGN and TRN for the TRN-LGN masked analysis with unsmoothed data. Black and red represent the right and left ROIs, respectively. a) Polar angle distributions from the LGN and TRN ROIs. The visual field was segmented into 36 bins for visualization. b) Flicker frequency distributions from the LGN and TRN ROIs.

Statistical map threshold of q[FDR] = 0.05 resulted in critical *p*-values ranging .0030–.0082. The distribution of polar angle responses are again significantly different between the LGN and TRN (KS test: $D_n = 0.0863$, $p = .0011$, and similarly for the flicker frequency preference ($D_n = 0.393$, $p = 7.8\times10^{-68}$). This analysis indentified substantially

less voxels in the TRN and LGN in both hemispheres than the when using the smoothed data (number of TRN voxels identified: right 0.54×, left 0.76×; number of LGN voxels identified: right 0.64×, left 0.67×). From this analysis, a few things are apparent. First, it seems plausible that the over-representation of the upper left visual field in the TRN was an artefact of smoothing, and therefore it seems safe to conclude that the TRN represents the full visual field without preference in each hemisphere. Second, it appears that if the TRN does not prefer a single flicker frequency, and instead has regions which resonate at across those frequencies tested here. Further analysis will be required to determine whether there is a meaningful structure to the distribution of flicker preferences of the TRN. Third, even when using a small 2mm FWHM kernel, the procedure helps greatly for detecting visual structures with high resolution fMRI, which is crucial for this paper. Even the LGN was substantially less detectable when using unsmoothed data within the anatomically defined ROI. Since this set of experiments is focused on the LGN and TRN, it seems preferable to use the smoothed data for the detection and analysis of the TRN. However, work investigating the TRN adjacent to the pulvinar should take special precautions to avoid blurring these signals between these ROIs.

In general, this work shows that the TRN is not easily detectable in a principled and purely data-driven way. These participants were scanned for multiple hours on the same 210 s stimulus in order to obtain a very small number of TRN voxels. However, the inclusion of anatomical ROI masks increases the sensitivity of the analysis greatly. Because of participant drop-out, I was unable to obtain 30 functional runs from each participant. This affords us the opportunity to investigate the relationship between the number of runs collected and the number of visual voxels identified. I conducted a linear regression with the number of runs collected as my set of predictors and the percentage of active voxels out of the total number of voxels tested as my dependent variable. I preformed this test on the smoothed data within the full thalamus mask, and again within only the LGN-TRN complex (**Figures 14** and **15**, respectively).

While the 95% confidence interval is very large for both analysis due to the small *n* of this experiment, I nonetheless observe a significant relationship between the number of averaged runs and the voxels identified in the purely data-driven analysis

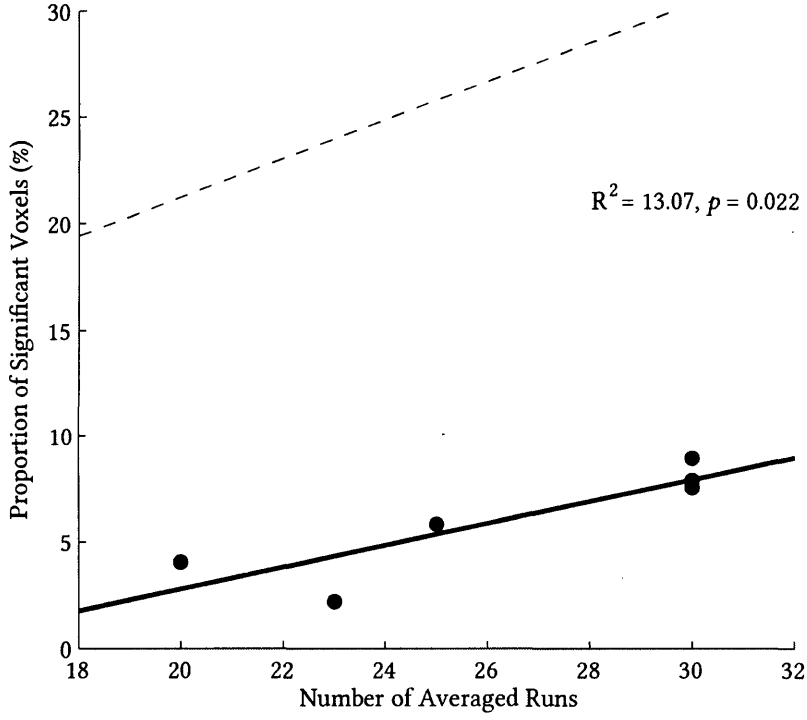

Figure 14. Relationship between number of averages and proportion of significant voxels in the data-driven analysis. Dashed line represents the 95% confidence interval.

(R^2 = 13.07, p =.022). This analysis suggests that more data would be beneficial for the development of a completely data-driven method for TRN localization, and may also suggest a plausible explanation for some of the variance observed between subjects (apparent in **Figures 6–8**). On the other hand, no relationship between the number of averages and detectable voxels was found using the ROI analysis (R^2 = 0.0012, p =.97), suggesting that using a conjunction of the smoothed data and anatomically-defined ROIs will allow one to detect these very small structures with significantly less data. Therefore, the time spent collecting the high-resolution proton density scans may be

offset by the reduced required number of functional runs. This bears further investigation, as our sample size is very small.

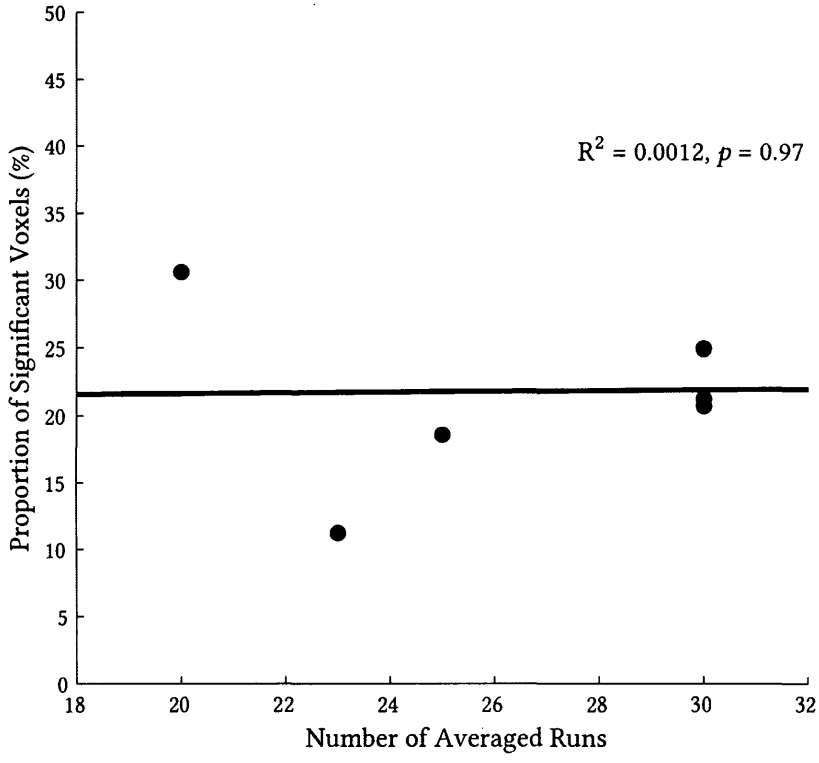

Figure 15. Relationship between number of averages and proportion of significant voxels in the ROI analysis. The 95% confidence interval is not shown due to scale.

3.2 Network Analysis of the Human Visual Thalamus

The previous analysis is not sufficient to capture the full dynamics of flicker modulation, as voxels may express multiple flicker frequency preferences. A Fourier analysis will only allow one to identify the flicker frequency that drives the maximal response. According to my hypothesis, certain structures in the thalamus, such as the TRN, should show multiple flicker-frequency preferences arising from its connectivity with first and second order thalamic structures. Therefore I further quantified the

similarities of flicker modulation between these visually driven voxels using the tools of graph theory. Graph theoretic analysis has recently be used to study the functional and structural organization in the brain with great success (for a review, see Bullmore and Sporns, 2009), providing holistic insights into the functional dynamics of neural networks. Here, I sought to identify the TRN via the similarity of its BOLD responses with those of the remaining thalamic nuclei. Another approach to this problem would have been to use a clustering algorithm, which would hopefully partition the data into a few meaningful regions of interest, including the TRN. Most clustering algorithms, however, are not well suited to this particular problem. A small minority of them are able to estimate the number of clusters in a data set, and it isn't clear exactly how many partitions are meaningful in this case. Those that do estimate the number of clusters require one to define a set of tuning parameters. These solutions are reasonable if one has a test data set with a known ground truth to tune the algorithm on, but in this case, the results of clustering algorithms become nearly impossible to interpret without bias.

Here, we expect the visual sector of the TRN to be a densely-connected structure, expressing functional relationships with both first and second order thalamic nuclei. Furthermore, we expect that the thalamo-cortical and cortico-thalamo-cortical pathways to express a different flicker frequency preference. For a clustering algorithm to detect this, it would need to allow for "fuzzy" partitioning of the data, with each voxel expressing multiple potential cluster memberships depending on its relationship with the rest of the data. Another approach is to construct an undirected, weighted graph of the visually-responsive voxels via a correlation matrix. This matrix records the similarities and differences between each voxel with every other voxel, and allows one to make use of sophisticated tools for the detection of densely connected voxels, as we would expect in the TRN.

I first partitioned the data using a method described in Rubinov and Sporns, 2011. Network analysis of the brain is a rapidly evolving and complex field, with multiple accepted best practices. Here, I opted to use a data-driven implementation that requires no threshold on the graph. I constructed a fully-connected and undirected

graph with only positive weights by finding the correlation matrix of the band-passed mean flicker profile data that survived the FDR-corrected retinotopy mask described in **Section 3.1**. I used the band-passed flicker profile data here as it was temporally filtered in a similar manner to resting-state data sets, which these network techniques are commonly used on (Van Dijk et al., 2010). Furthermore, I did not include the voxels identified by only the Fourier analysis of flicker-modulation to avoid introducing any bias into the analysis by way of circular reasoning. Each graph consisted of n nodes and $\frac{1}{2}n(n-1)$ connections, or edges. I define the weight of an edge as the correlation value between two voxels. I discarded all negative weights, as it isn't clear that the proposed method has a similar effect on heamodynamics that also give rise to negatively correlated networks found at rest (Chai et al., 2012; Smith et al., 2012). Therefore, for the purposes of proof-of-concept, I only considered the positive weights of the graph, and defined the strength of a node's connections as the sum of all positive weights, and did not directly consider the number of connections (or node degree).

Graphs can be partitioned into a set of communities to cluster a data set by defining a set of edge cuts that maximize the strength of the edges within a community and minimize the strength of the edges between non-overlapping communities (Fortunato, 2010; Schaeffer, 2007). One advantage of this approach is one can then find the modularity coefficient Q to measure the quality of a given partitioning as

$$Q = \frac{1}{v}\sum_{ij}\left(w_{ij} - e_{ij}\right)\sigma_{M_iM_j} \qquad [5]$$

where w_{ij} denotes within module connection weight, and e_{ij} denotes between module connection weights, between node i and j. $\sigma_{M_iM_j}$ = 1 when i and j are in the same module, and is 0 otherwise. v scales the result by the number of edges in the graph. Q falls in the range [0,1] with Q = 0 when there is no modular structure in the graph, and Q = 1 when the modular structure of the graph is pure, i.e., there are no edges between communities (Rubinov and Sporns, 2011). Crucially, I have not regressed any global

signals from the data, so there is no chance that I have significantly altered the correlation structure of the graph before this analysis (Murphy et al., 2009).

The graph partitioning proceeded as follows. I first generated a partition with the greedy modularity-maximization algorithm described by Blondel et al., 2008. This algorithm was an attractive choice as it scales well to very large networks, finds a high-level modular structure with few communities, and automatically selects the number of communities within the graph. It should be stressed here that we have no *a priori* reason to expect a particular number of clusters to exist in the data, but also that it seems reasonable that multiple partition resolutions would serve as sensible solutions to the problem of defining the flicker response properties of the thalamus. Therefore, as a initial step, it seems logical to define a reliable and broad partition of the data with few clusters. The algorithm assigns each node to its own community, and then iteratively considers the increase of Q with the joining the communities of adjacent nodes. Each join that produces a larger Q is retained, and this procedure repeats until no further improvements result from further iterations. This partition is then fine-tuned using an algorithm that considers, once for each node, any possible improvement to Q that might result from moving a node into a neighbouring community (Sun et al., 2009). Both procedures are heuristic, and find marginally different partitions with each run. I therefore computed Q over 100 iterations with this procedure to generate 100 candidate partitions. I retained the partition with the largest Q, under the assumption that the partition with the largest modularity was the best partitioning of the graph.

I quantified the similarity of the individual partitions using the variation of information, a measure of distance in partition space (Meilă, 2007). To accomplish this, one first defines the entropy associated with a partition M as

$$H(M) = \sum_{u \in M} P(u) \log P(u) \qquad [6]$$

where $P(u) = \frac{n_u}{n}$, and n_u is the number of nodes in module u. Next, one finds the mutual information between the partitions M and M' as

$$I(M,M') = \sum_{u \in M} \sum_{u \in M'} \log \frac{P(u,u')}{P(u)P(u')} \quad [7]$$

where $P(u,u') = \frac{n_{uu'}}{n}$, and $n_{uu'}$ is the number of nodes that simultaneously members of partition M and M'. With this, one can define the variation of information VI as

$$VI = \frac{1}{\log n}[H(M) + H(M') - 2I(M,M')] \quad [8]$$

where a VI falls in the range [0,1] with $VI = 0$ denoting equivalent partitions, and $VI = 1$ denoting maximally distant partitions. This measurement is useful for checking the stability of the algorithm over multiple runs, and for comparing the content of runs with identical modularity coefficients. In the case that the algorithm returned multiple partitions with identical Qs, I retained a single partition from the largest set of Qs with $VI = 0$, under the assumption that the most stable partition was the best partition of the graph. These algorithms are available as MATLAB code in the Brain Connectivity Toolbox (http://www.brain-connectivity-toolbox.net; Rubinov and Sporns, 2010). An overview of this strategy is presented in **Figure 16**.

I next aligned the partitions across participants in the following way. I first ranked each community by the sum of the associated F statistics determined as described in **Section 3.1**. I next scaled each voxel time series to the range [-1 1] by dividing it by its maximum absolute value. I did this because the amplitude of the BOLD signal between various neural populations may not convey meaningful information in this context, as BOLD amplitude is dependent on the density of vascularisation, and this could distort the computed centroid (Logothetis and Wandell, 2004). Next, I computed the centroid of each partition as the mean of the normalized time series within each partition, as well as the SD of each partition. I next sequentially aligned each the

partitions of S2–6 to the registration target centroids of S1 in descending ranked order by computing the squared Euclidean distance of each centroid from the registration target, and then re-assigned partition labels such that the labels for each partition in each group matched that of the partition in S1 with the smallest distance. I finally computed a group centroid for each partition as the mean of the scaled data with each community across subjects, and the associated SD. The results of this registration can be found in **Figure 17.** In **Figure 18** I show the distribution of clusters within each ROI, in **Figure 19** I show the registered cluster maps projected back onto each participant's brain.

Here, I used voxels found using by analyzing the retinotopy of the smoothed data, and input flicker profiles also produced from the smoothed data. Exactly three clusters were found algorithmically in each subject. These clusters were relatively similar in size (*n* Cluster 1 = 16030, *n* Cluster 2 = 16103, *n* Cluster 3 = 14479). Cluster 1, which represents voxels with a flicker preference spanning 2.5–15 Hz, represents a large proportion of

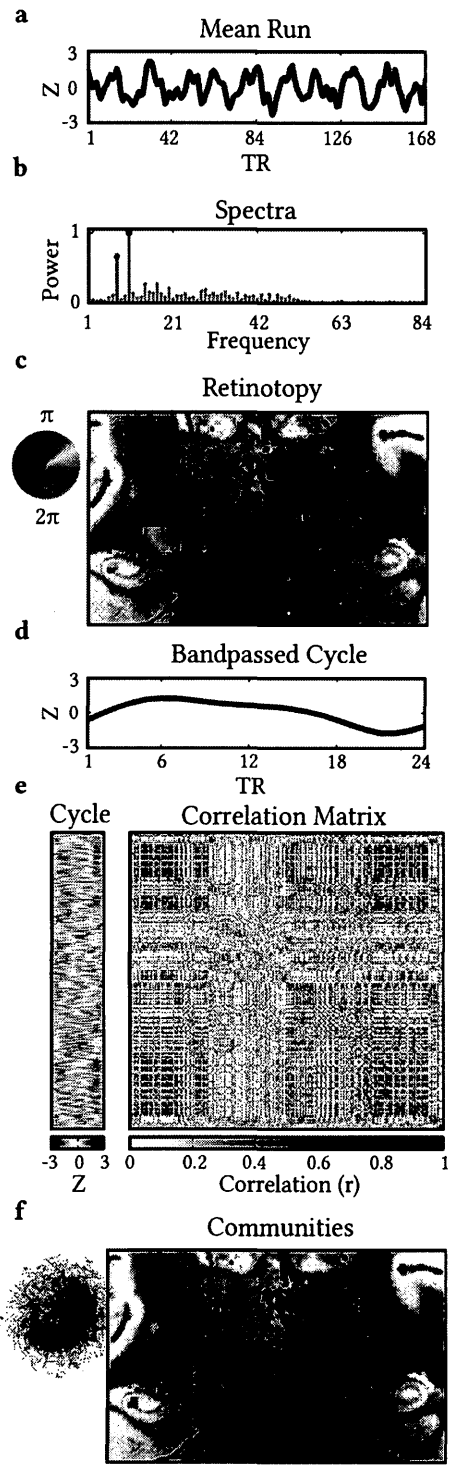

Figure 16. Overview of the community detection pipeline. a) Sample time series from the LGN. b) Power spectral densitiy estimate of a, stimulus frequencies in red. c) FDR-corrected retinotopy from S1. These serve as the ROIs for further analysis. d) Bandpassed cycle from the same voxel as a. e) Matrix of cycles for all retinotopic voxels and the associated correlation matrix. f) Community memberships projected back into the thalamus, showing three major partitions of the graph.

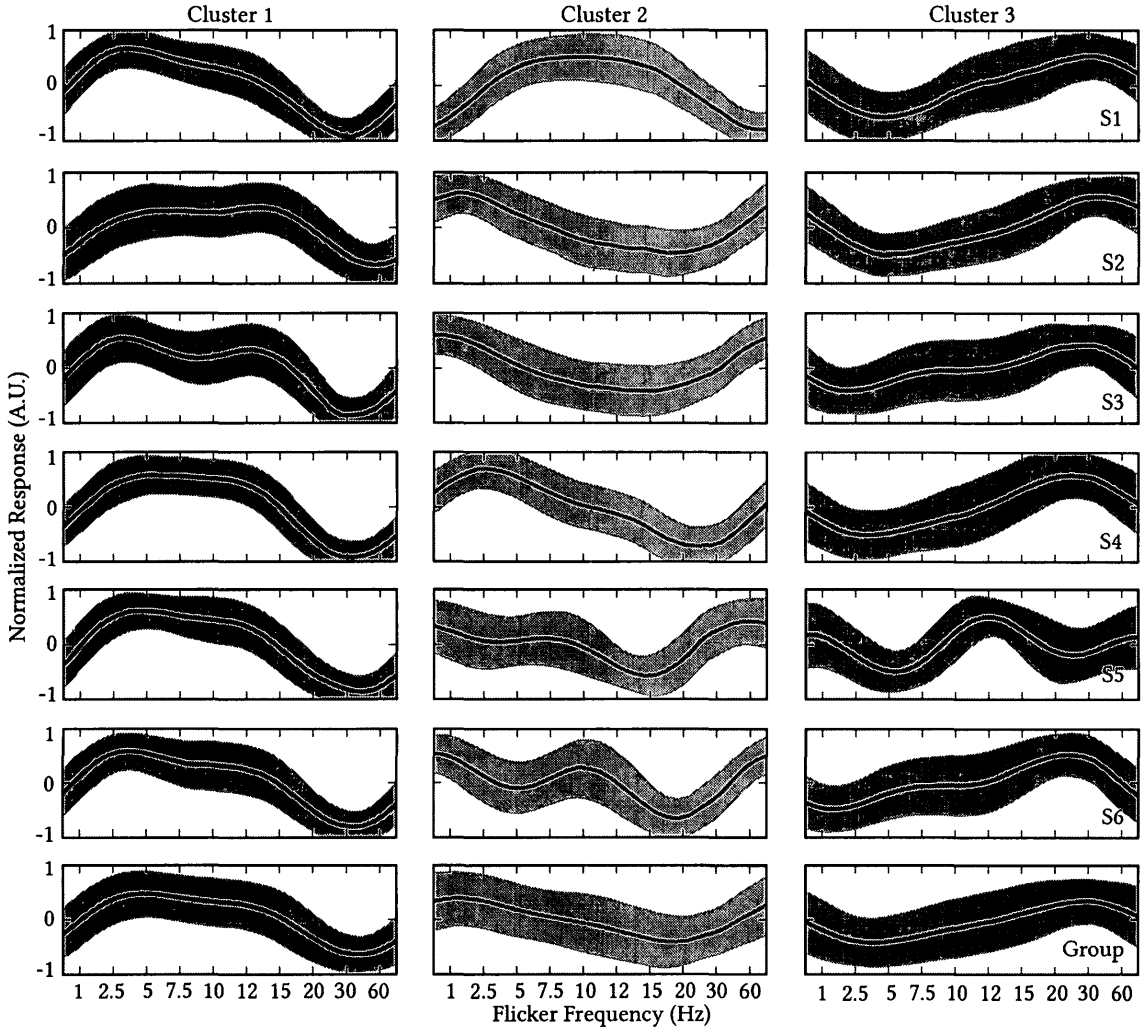

Figure 17. Group-level community flicker profiles obtained from smoothed data. The results of the group analysis are found in the final row.

LGN ROI voxels in all participants (**Figure 18**). Visual inspection of **Figure 19** also shows that this cluster represents the SC in some, but not all, participants. This clustering of the primary visual nuclei is the most reliable effect: patterns in the other clusters are more difficult to interpret. Cluster 2 is non-uniformly represented among all ROIs across subjects, and the group average masks rather different underlying flicker profiles. The profile itself represents preference for very high flicker (60 Hz), and flicker from 1–7.5 Hz . In some cases, Cluster 2 appears to represent some primary visual nuclei, or the TRN along the lateral edge of the pulvinar, but the inconsistency amongst participants make this observation hard to interpret. Cluster 3, which is the smallest

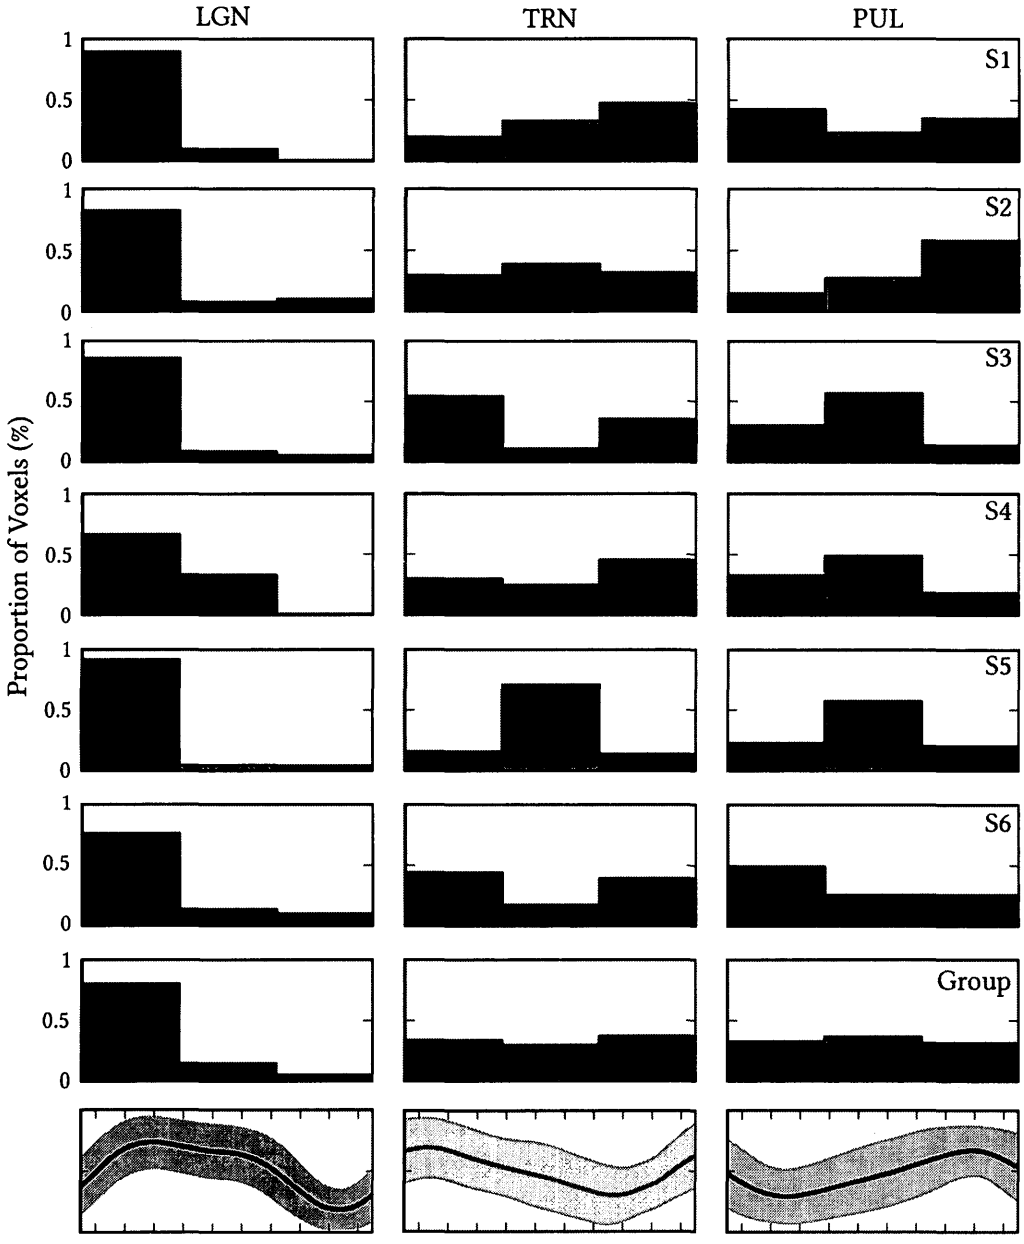

Figure 18. Distribution of smoothed cluster memberships across ROIs in percentages.

cluster by number of voxels, appears to be more consistent among participants. It generally represents voxels which are most active when presented with a flicker between 15 and 60 Hz. Large, contiguous regions of Cluster 3 appear to represent both the MGN and pulvinar region in multiple participants. However, the group level analysis does not show any consistent pattern of preference between Cluster 1, 2 or 3 for

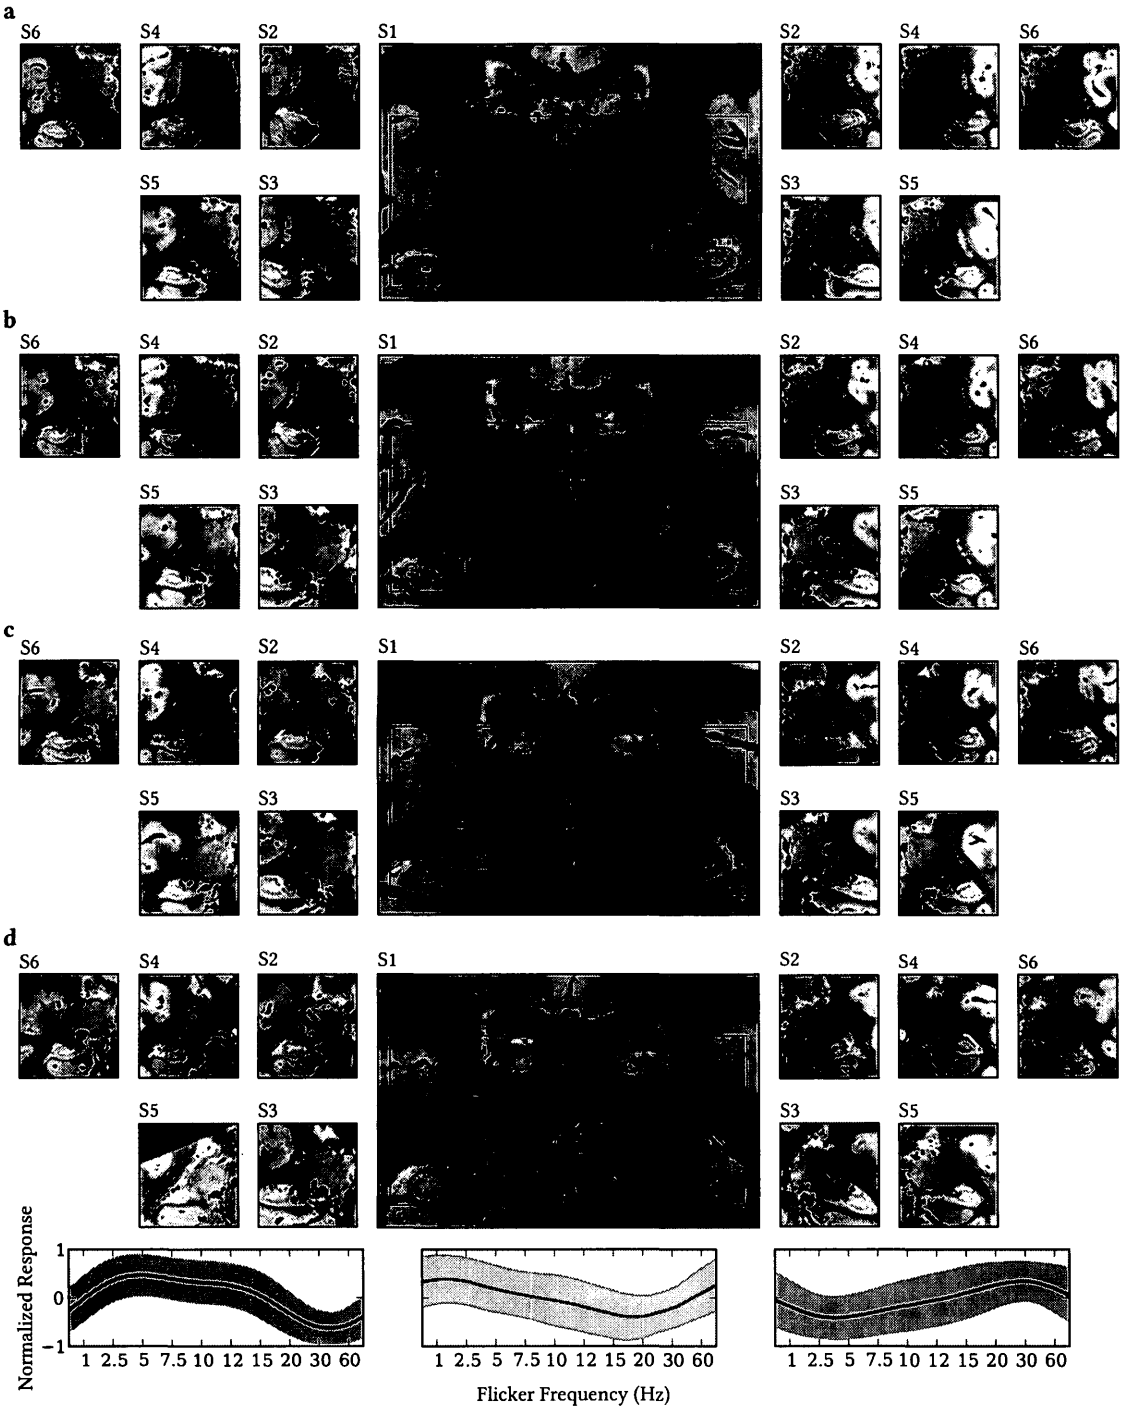

Figure 19. Communities of flicker profiles in the human thalamus. Order of the clusters are arbitrary.

either the pulvinar or TRN. This in essence replicates the results obtained using the Fourier analysis (**Figure 9**), and also demonstrates that the use of the flicker profile is a

viable alternative to the Fourier analysis for the purpose of classifying neural structured by their flicker frequency preference.

These partitions were remarkably stable across iterations of the algorithm. The modularity coefficient was modest across all participants (*Q*: $\bar{x}$ = 0.283, SEM = 0.0056), but each iteration of the algorithm returned small variations of the resultant partition (*VI*: $\bar{x}$ = 0.0305, SEM = 0.00031). If these voxels were, in fact, simple noise, it would seem unlikely that each iteration of the algorithm would produce such similar partitions of the graph. Instead, we would expect to see a very high VI, even in cases that produce similar overall distributions within ROIs. One explanation for the similar distribution of cluster memberships between the pulvinar and TRN ROIs was that the smoothed data blurred pulvinar or TRN-specific flicker profiles into the adjacent structure, confusing the results. This would particularly explain the large variability seen in cluster membership between participants. I tested this by re-running the analysis on the unsmoothed data, while re-using the retinotopy masks obtained using the smoothed data (therefore including the exact same voxels in both analysis). The results for this analysis are shown in **Figures 19** and **20**.

Two things are evident from this alternative analysis. First, the registration procedure, which performed well for the smoothed data, does not work as well with noisier centroids. Of note is the mis-registration of Cluster 3 and 1 in S5. Although the algorithm still identified exactly three clusters for each participant, there is clearly more variance in the cluster centroids between participants. Nonetheless, the group-level mean cluster centroids are remarkably similar between the smoothed and unsmoothed analysis, and these partitions were very stable across iterations (*Q*: $\bar{x}$ = 0.243, SEM = 0.0067; *VI*: $\bar{x}$ = 0.0038, SEM = 0.00065).

Second, the LGN still preferentially clusters under a low-frequency flicker profile. The TRN seem to have a slight preference for Cluster 2 here (which is remarkably similar to Cluster 3 from the smoothed analysis), preferentially responding to flickers between 15–60 Hz. However, it might be safer to conclude that neither the

pulvinar nor TRN have any preferential flicker tuning that is observable using this method. Therefore, this analysis shows that the obtained cluster-centroid flicker profiles

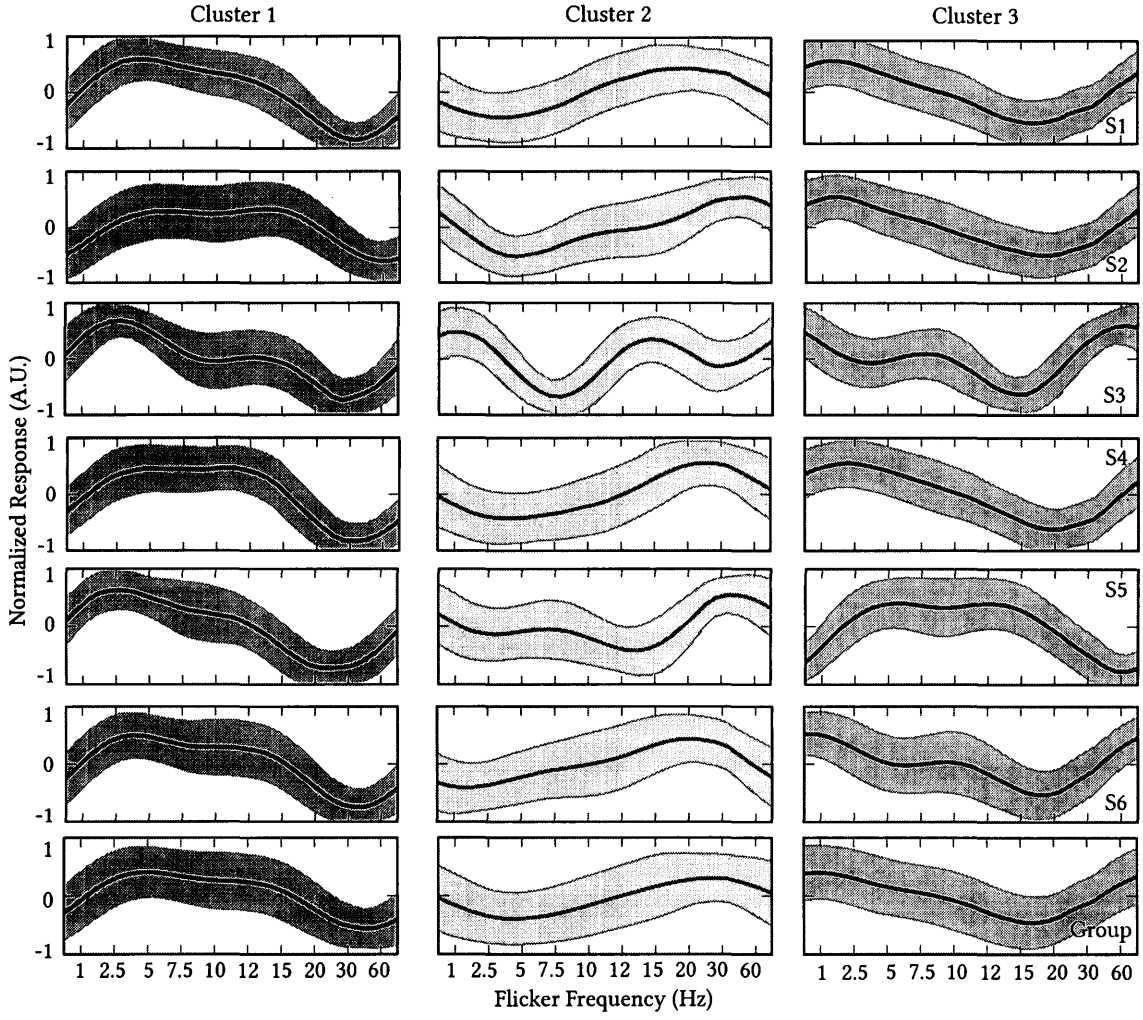

Figure 20. Group-level community flicker profiles obtained from unsmoothed data. The results of the group analysis are found in the final row.

are not an artefact of smoothing, and this analysis can reliably segment first order visual nuclei in the visual thalamus, but it remains unclear whether the findings relating to the TRN or pulvinar are simply due to noise, or are, in fact, due to these structures having a wide range of resonant frequencies.

Another explanation for why this analysis found no interesting cluster patterns would be that this data-driven approach would is highly sensitive to inter-subject differences in flicker profiles. In order to visualize the variability inherent in the data, I

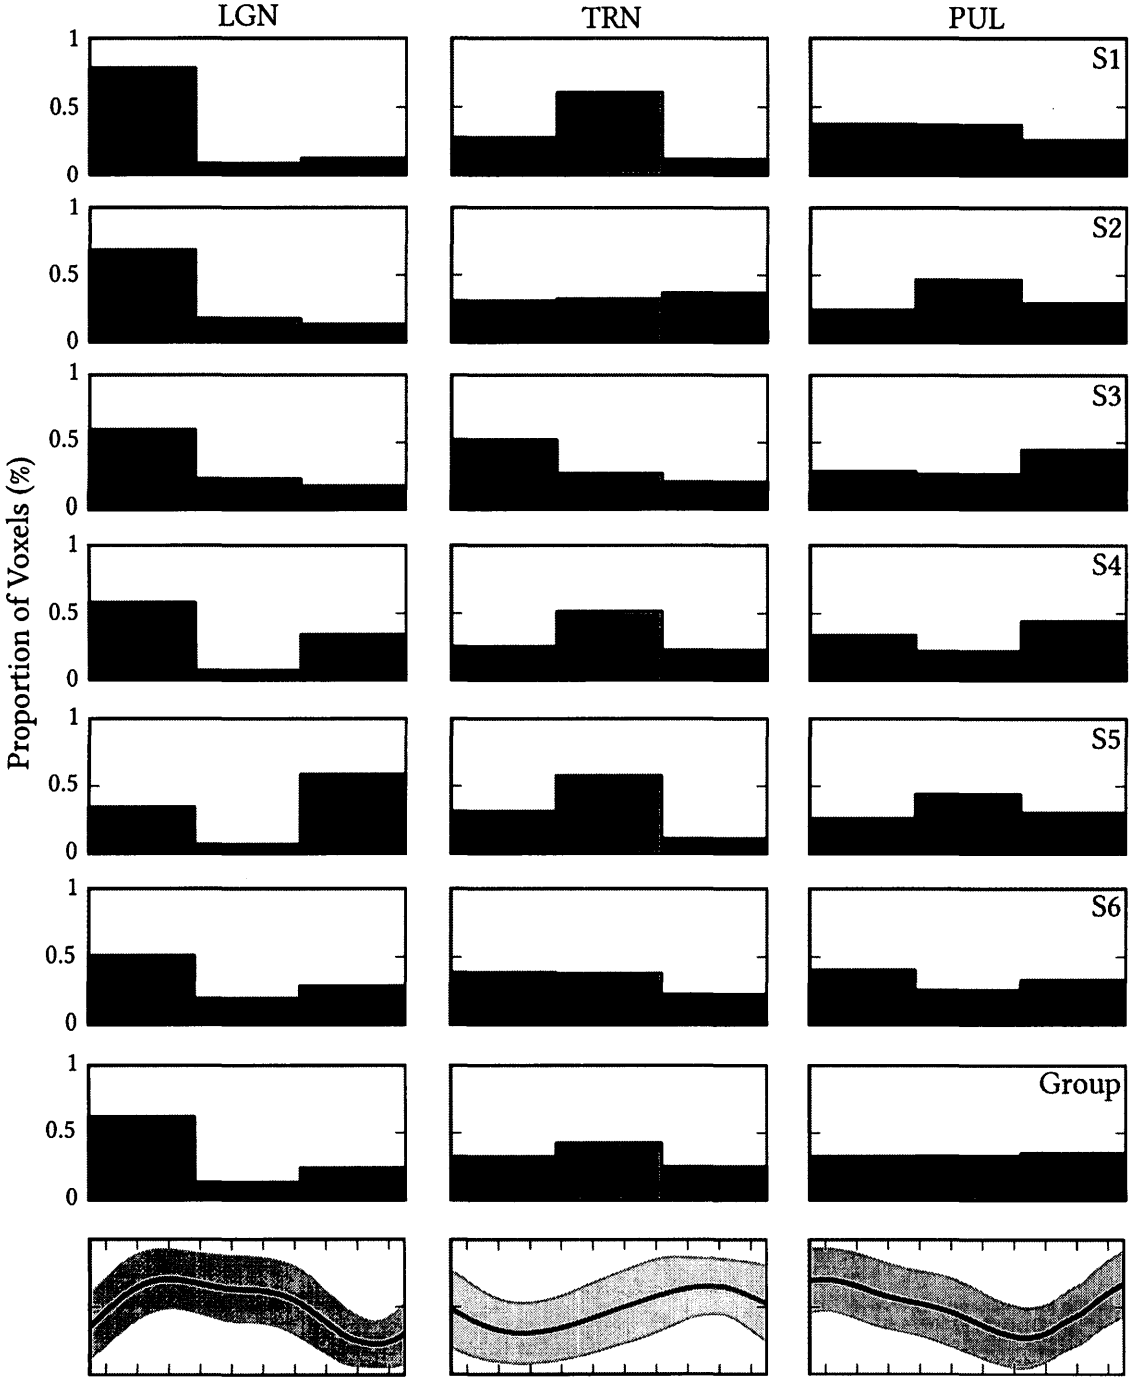

Figure 21. Distribution of unsmoothed cluster memberships across ROIs in percentages.

computed the mean and SD of the LGN flicker profile in each hemisphere from each participant and plot them in **Figure 22**. Since we are not comparing between structures, I present these data in percent signal change form. It is clear that there are differences in both BOLD amplitude and exact flicker-frequency tuning between participants in this structure, despite the fact that it reliably clusters in all participants. This suggests inter-subject differences in either BOLD lag, neural-population resonance, or a combination of the two, and this poses a significant problem for the group-level analysis of this sort of data. Previous work shows variability of up to 4 Hz in the relatively well-defined spindle frequency between participants during sleep (11–15 Hz), as well as considerable variability of BOLD-α coupling between participants, which are normally positively correlated in the thalamus, but in one participant was negatively correlated (Gonçalves et al., 2006; Werth et al., 1997). It seems like a reasonable next step to better quantify the relationship between these flicker profiles and EEG oscillations in the cortex. This information would be useful both for better understanding the relationship between flicker and neural resonance, and for understanding if there is an electrophysiological basis behind this inter-participant variation. If not, it may be the case that this variation is due to much less interesting things, such as inter-participant differences in physiological noise, head motion, or vigilance (Van Dijk et al., 2012; Olbrich et al., 2009; Shmueli et al., 2007).

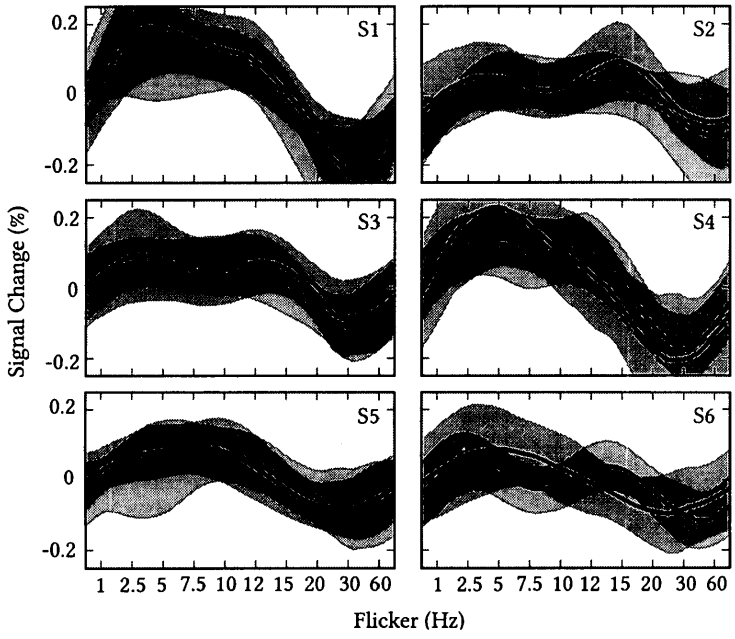

Figure 22. Inter-participant variability of the LGN flicker profile. Red = left LGN, Blue = right LGN.

As a follow-up analysis, I set out to determine whether the TRN serves as hub between the LGN and pulvinar. Specifically, I expected the LGN to prefer low frequencies, as shown in **Figures 12–17**, and at least some regions of the pulvinar to prefer higher β frequencies, as shown by (Saalmann et al., 2012). If these two regions are driven by the stimulus, and the visual sector of the TRN modulates communication between these regions, I would expect to see a bimodal flicker profile in the TRN. In the context of a graph, the TRN would act as a bridge between the group of voxels responding most to low frequencies and the group of voxels responding most to high frequencies (Cluster 1 and 3, respectively, from **Figure 17**).

To accomplish this, I preformed the analysis proposed by (Guimera and Nunes Amaral, 2005). In essence, this procedure seeks to find nodes in a network that have a large number of inter-community connections, as defined in the previous analysis. First, one finds the degree of each node in the network, that is, the number of nodal connections per node. Crucially, I preformed this analysis on a modified form of the data. Any analysis of node degree, or similar measurements such as the clustering coefficient, are biased towards identifying larger structures by virtue of the homogenous BOLD response found within contiguous brain regions (Power et al., 2011). Therefore, I constructed a sub graph of the thalamus from a single time series found within each contiguous region of activation, as identified by the retinotopy analysis of the smoothed data. I selected each time series as the voxel with the maximum F statistic from the retinotopy analysis, under the assumption that this voxel would provide the cleanest signal and best representation of the surrounding region's flicker profile. In cases where a single contiguous region of significant retinotopy activation spanned two or more anatomical ROIs, I subdivided along the anatomical boundaries. I next constructed each graph as the positive correlations derived from a correlation matrix of the band passed, smoothed flicker profiles. In this case, the degree of node i is simply the number of nodes with any positive correlation to node i.

I identified communities in this graph using the same algorithm defined previously. Next, I found the normalized within-module degree for each node as

$$Z_i = \frac{k_i - k_{s_i}}{\sigma_{k_{s_i}}} \quad [9]$$

where k_i is the number of links of node i to other nodes within its community s_i, k_{s_i} is the average of k over all nodes in s_i, and σ_{k_i} is the standard deviation of k in s_i. This serves as a measure of how well a particular node is connected to other nodes within its community. I then found the participation coefficient P_i for each node i as:

$$P_i = 1 - \sum_{S=1}^{N_M} \left(\frac{k_{is}}{k_i} \right)^2 \quad [10]$$

where N_M is the number of detected communities, k_{is} is the number of links of node i to the nodes in community s and k_i is the total degree of node i. Therefore, the participation coefficient of a node is close to 1 if its edges are uniformly distributed among all of the communities, and 0 if all its links are contained within its own community.

Since this analysis makes use of data-driven ROIs, and not all of the voxels identified by the retinotopy analysis, I re-computed the communities for each participant. While the results were not drastically different from the previous analysis, the overall modularity was slightly higher and remained remarkably stable across iterations (Q: $\bar{x}$ = 0.314, SEM = 0.0053; VI: $\bar{x}$ = 0.0382, SEM = 0.0082). In S4 and S5, 4 communities were detected. No attempt was made to register these communities across participants, so this discrepancy did not pose a problem.

Briefly, I set out to define each ROI as being either peripheral (small P) or well connected (large P), as well as being either a non-hub (small z) or hub (large z) node. Based on previous work, I expected to find the TRN to be a well-connected hub, mediating communication between those communities that represent both low and high flicker frequency responses. I preformed this analysis on data considering all positive

weights, and again only retaining the top 30% strongest weights. The results of these analysis are shown in **Figure 23**.

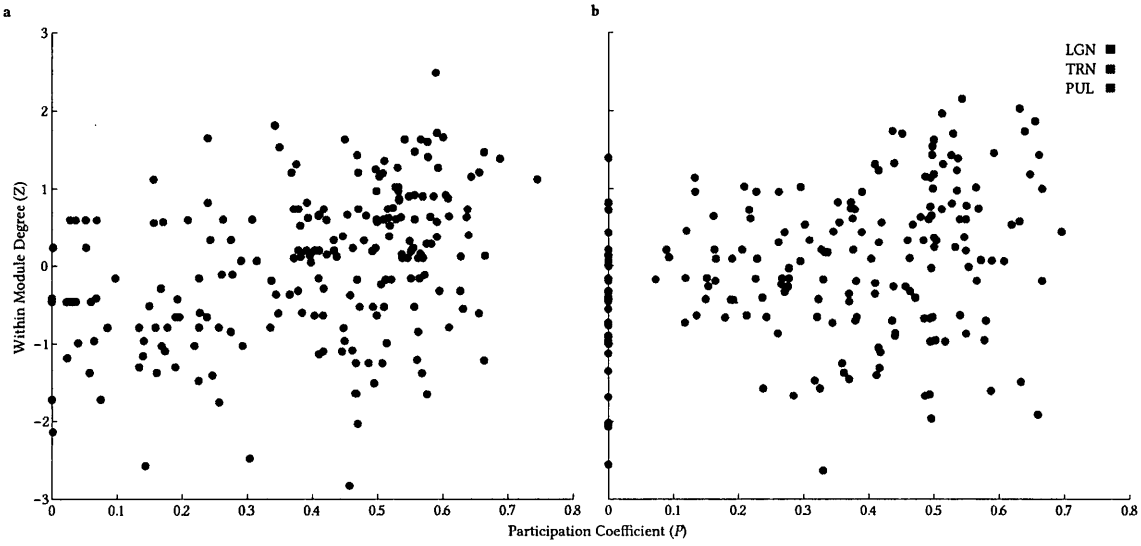

Figure 23. Modularity of the thalamic nuclei. Scatterplot depicting the participation coefficient vs. normalized within-module degree across all participants of the LGN, TRN, and pulvinar. a) Node degree calculated as the number of connected nodes with any positive correlation. b) Node degree calculated as the number of connected nodes with correlations in the top 30%.

The paper this method was derived from defined hub nodes as those with a normalized within-module degree of at least $z = 2.3$ (Guimera and Nunes Amaral, 2005). It appears that none of the ROIs defined in this analysis serve as particular hub within their own community. No class of ROI is obviously associated with a relatively large *P*, indicating a large number of inter-community connections relative to the other ROIs. In **Figure 24** I plot the proportion of voxels within each ROI with a particular participation coefficient (binned in increments of 0.1) to better visualize the differences between these brain regions. The data here shows a trend towards the expected effect with graph thresholds: approximately 15% of the identified TRN voxels have the largest participation coefficient ($P = 0.7$) when only the top 30% of positive weights are considered, but not when all positive weights are considered. This is in contrast to the pulvinar, which only has approximately 5% of its ROIs with an equivalent *P*.

The original paper also distinguishes between peripheral nodes, with most links within their own community ($P < 0.62$), non-hub connector nodes, with many links to other modules ($0.62 < P < 0.80$), and finally non-hub 'kinless' nodes, with links

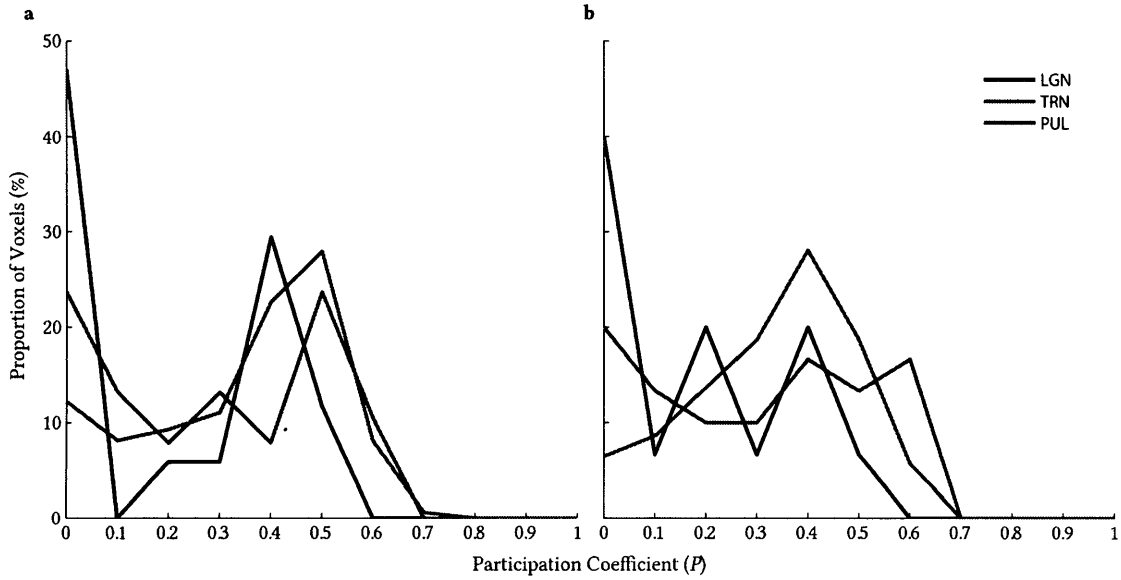

Figure 24. Distribution of participation coefficients within ROIs. Results pooled across participants. Data binned in intervals of 0.1. a) Distribution of participation coefficients across the three identified ROIs calculated using an unthresholded graph ($r > 0$). b) Distribution of participation coefficients across the three identified ROIs calculated using the thresholded graph ($r = 30^{th}$ percentile).

homogenously distributed among all communities. This analysis shows that the TRN and pulvinar have remarkably similar network characteristics, being composed of significantly more non-hub connector nodes than the LGN. That is, their respective flicker profiles cover a wide range of frequencies. These results fail to reject the null hypothesis that the TRN is not detectable as a hub in the thalamus using tremotopy. According to these results, the TRN cannot be isolated from the surrounding structures by its flicker profile connectivity with first and second order thalamic nuclei. Specifically, the resonance properties, and the network centrality, of the TRN are not unique, and the pulvinar may overshadow attempts to localize the TRN in a data-driven way using these measures. On the other hand, these results do support the findings of the Fourier analysis where both the pulvinar and TRN were found to have a broad

flicker frequency tuning (**Figure 9**). Considering that the community-detection algorithm detects three community centroids with flicker profiles preferring the α to high β range, it follows that a high participation coefficient in the pulvinar and the TRN would be due to broad range of flicker frequency profiles in these structures. Considering the vast difference in sizes between the TRN and pulvinar, it seems unlikely that this result was driven completely by noise: if noise was the sole contributor to each nodes participation coefficient, then we would expect to see a higher density of pulvinar nodes with large *P*s than in the TRN.

One issue that needs to be addressed in future network studies of the thalamus is appropriate brain coverage. The TRN is indeed an important hub in the thalamo-cortical loop, with widespread connections across every sensory modality and extensive connections with the pre-frontal cortex, but this information is not available in these data due to the very small region of the brain acquired. A possible solution, given the stimulus-driven nature of this experiment, would be to collect data sets at multiple resolutions and construct a whole-brain graph from parts of each image. Care would have to be taken to ensure similar signal-to-noise ratio across these data sets, but in theory this could be done. Graph analysis of the brain is a new technique and has almost always been used on whole-brain data (Wang Jinhui et al., 2010), so there is currently little evidence that this sort of analysis is useful when looking at such a restricted sub graph of the human brain.

3.3 Seed-Based Correlation Analysis of the Thalamic Reticular Nucleus

Following these failed attempts to develop a data-driven localizer for the human thalamic reticular nucleus, I set out to define functional properties of the TRN-LGN complex in the human brain using a seed-based correlation analysis commonly used in resting-state experiments (Raichle et al., 2001). One major advantage of this sort of analysis is the results are easily interpretable: positively correlated regions express temporally correlated metabolic demands, and negatively correlations imply that the two regions in question perform unrelated or opposite functions. Importantly, it isn't

clear how negative correlations are related to inhibitory processes in the brain. Modelling studies suggest that the negative BOLD response (a movement of the BOLD signal below baseline levels during stimulus presentation or a task), is not directly related to inhibitory processes, interneurons themselves require approximately 20% of the glucose required by pyramidal cells (suggesting inhibitory neural populations should express a low-amplitude positive BOLD response when contrasted with an excitatory process, but not negative), and GABA concentrations in the anterior cingulate cortex as measured using magnetic resonance spectroscopy positively correlates with the amplitude of negative BOLD responses seen at rest or during emotional processing (Logothetis and Wandell, 2004; Northoff et al., 2007; Sotero and Trujillo-Barreto, 2007). While this experiment is not well suited to directly address this question, as conclusive evidence would require causal manipulation of the TRN while simultaneously recording BOLD signal from the LGN (perhaps with a combination of optogenetics and fMRI; Lee et al., 2010), it at least provides some insight into the relationship between inhibitory processes and the BOLD signal. To the best of the author's knowledge, the TRN is the only structure in the human brain that consists entirely of inhibitory interneurons and is spatially distinct from other neural populations at a resolution of 0.75 mm. Therefore, the LGN-TRN complex seems ideally suited for answering this question.

The question I sought to address here is whether the TRN participates in a contralateral-inhibition network within the brain. In the cortex, interhemispheric inhibition is well defined, but its function is poorly understood. Magnetic stimulation of points on cortical topographic maps produces a polarized effect on topographically-matched locations of the contralateral hemisphere, with inhibitory stimulation eliciting contralateral facilitation, and vice versa (Ferbert et al., 1992; Hilgetag et al., 2001). A recent experiment in somatosensory cortex implicates $GABA_B$-mediated layer V projections to contralateral layer I apical dendrites underlies the inhibition of topographically-matched, contralateral evoked responses (Palmer et al., 2012). This provides a candidate mechanistic explanation for the perceptual phenomenon of

binocular rivalry, which has previously been proposed to be a form of intraocular suppression in areas as early in the visual stream as the LGN, but is also influenced by top-down object-level attention and seems to respect the topographic map (Haynes et al., 2005; Mitchell et al., 2004; Sengpiel et al., 1995; Wunderlich et al., 2005; Zhang et al., 2011). Here I build the hypothesis that the TRN participates in a contralateral inhibitory circuit which respects the topographic map that performs a similar function to that of interhemispheric Layer V to Layer I inhibition in the cortex. Therefore, I predict that at least part of the visual TRN adjacent to the LGN contralateral to the LGN driven by the stimulated visual field would be driven by the flickering checkerboard. Recently, profound walking-driven modulation of the cat ventrolateral (somatosensory) thalamus and associated TRN was demonstrated: spindle frequency discharges were recorded from the TRN during the "stance" phase of the forepaw contralateral to the recording site (Marlinski et al., 2012). That is to say, while the forepaw contralateral to the TRN was receiving sensory input, the TRN switched to a fast-spiking mode, driving IPSPs in the ipsilateral thalamus (Steriade et al., 1993). This would prevent any descending cortical process from interrupting the motor plan of the "swing" forepaw. Here, I would expect to see positive correlations between the LGN contralateral to at least some TRN seed voxels, and negative correlations between the same TRN seed voxel and the ipsilateral LGN.

I limited my analysis to the voxels identified by the FDR-corrected retinotopy analysis within the anatomically-defined LGN and TRN masks (**Figure 12**). I computed the correlation between a seed voxel time series taken from the TRN and each voxel in the LGN on either side of the brain. In some cases, this analysis produced positive correlations between the LGN and ipsilateral seed TRN, and negative correlations between the seed TRN and contralateral LGN (data not shown). This result was not particularly interesting: due to the close spatial proximity of the LGN and TRN in the human brain and the resolution of the functional data, these results could either be driven by the stimulus, or local correlations due to spatial smoothing, head motion, or both (Power et al., 2012). However, in all participants, I was also able to identify

retinotopic TRN voxels with the opposite pattern of correlations (**Figure 25**). In order to determine whether this effect is as all modulated by flicker frequency, I also produced correlation maps using the flicker profiles for each voxel.

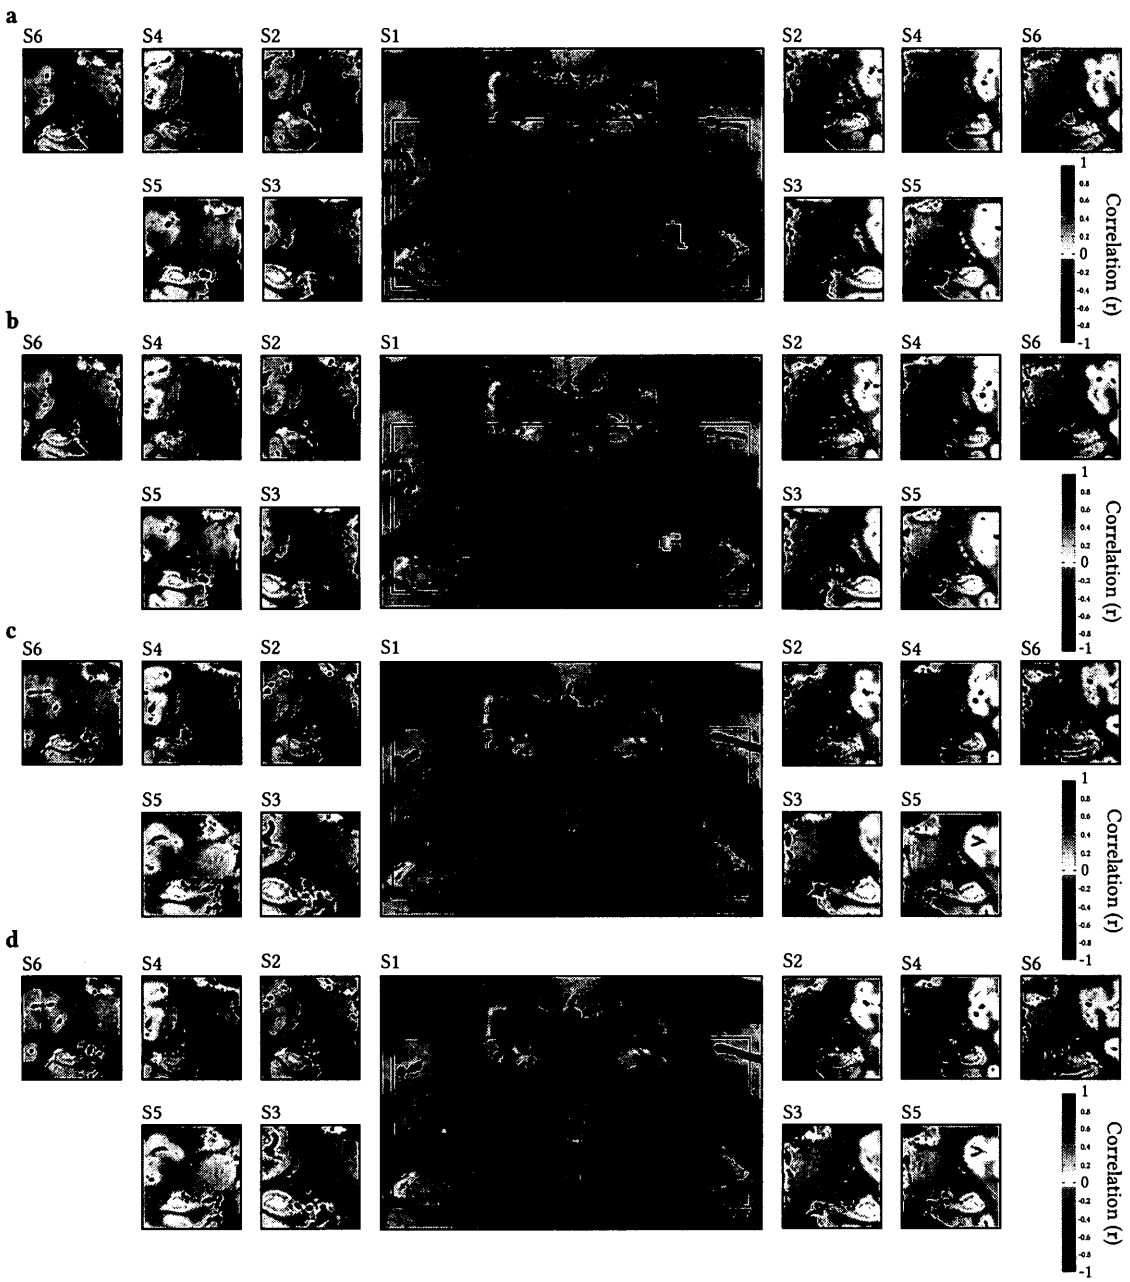

Figure 25. Contralateral functional connectivity of the LGN-TRN complex. Seeds for each participant circled in red. Panels a & b are more anterior than c & d. a,c) Correlation maps of the LGN-TRN complex produced by correlating the mean runs of each voxel. b,d) Correlation maps of the LGN-TRN complex produced by correlating the flicker profiles of each voxel.

This view of the data makes a few things clear. The location of the seeds required to produce this contralateral correlation effect is fairly consistent among participants. Generally, these seeds fall just dorsal lateral to the LGN and are within

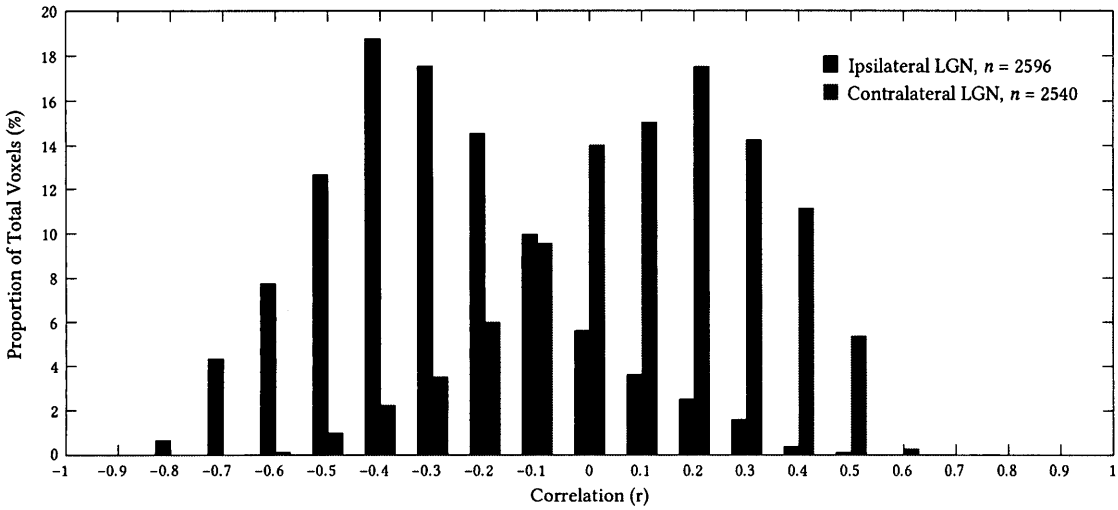

Figure 26. Distributions of mean run correlations with the TRN seed from the ipsilateral and contralateral LGN pooled across all participants.

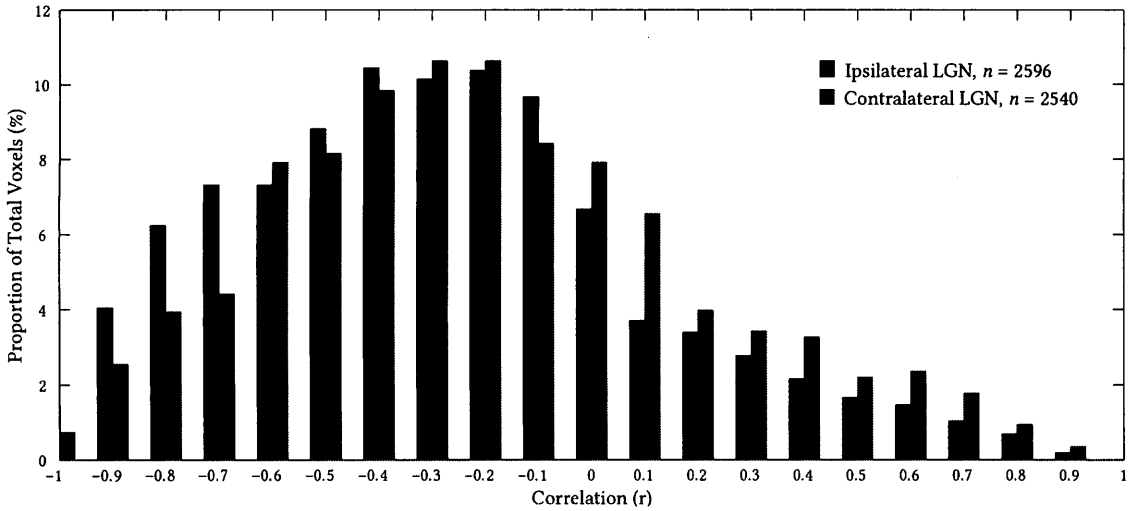

Figure 27. Distributions of flicker profile correlations with the TRN seed from the ipsilateral and contralateral LGN pooled across all participants.

1–3 mm of the dorsal edge of the anatomically-defined LGN. No seed produced this effect that was found to be further than the junction where the TRN meets the lateral pulvinar (the closest case is seen in Panel a, S4 R). In all cases, the mean run produced the contralateral positive correlation expected, while the correlation of flicker profiles

produced mostly negative correlations with both LGN (**Figure 26** and **27**, respectively). The distribution of correlations for the contralateral and ipsilateral LGN were found to be significantly different in both cases, with the contralateral LGN expressing more positive correlations in both tests (KS test run: $D_n = 0.644$, $p = 2.22 \times 10^{-308}$, KS test cycle: $D_n = 0.0936$, $p = 2.98 \times 10^{-10}$), although it should be stressed that the difference between distributions are order of magnitudes higher in the case of run correlations.

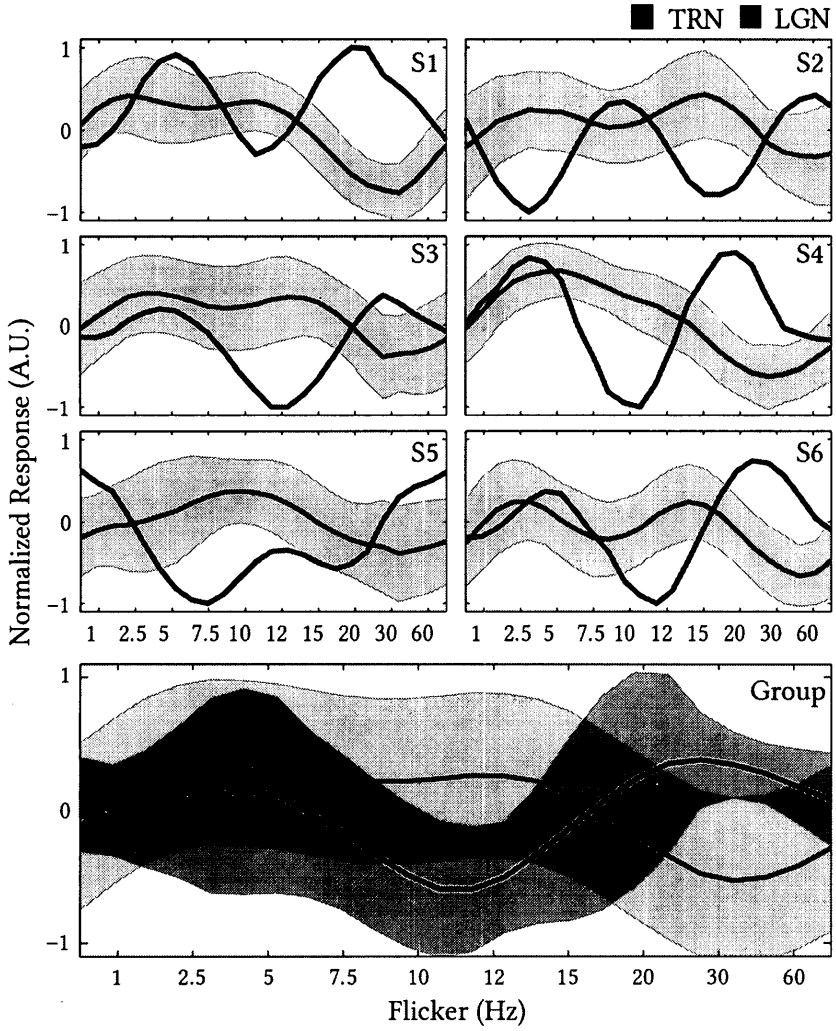

Figure 28. Flicker profiles of paired LGN and TRN. Flicker profiles were taken from the TRN seed from each participant, and averaged to produce the group-level analysis. Flicker profiles from the LGN represent the mean of the ipsilateral LGN to the seed.

The finding that the flicker profiles of the TRN and the LGN are largely negatively correlated is unanticipated. To further investigate this relationship, I calculated the mean flicker profile of the ipsilateral LGN in each participant and plot this mean flicker profile against the seed TRN flicker profile in **Figure 28**. It appears that in most participants, these regions are maximally negatively correlated when the retina is driven between 10–15 Hz, with the LGN responding more vigorously and the TRN responding less than average, and between 20–60 Hz, where the TRN responds maximally and the LGN responds less so than on average. It is interesting to note that the mean ipsilateral LGN flicker cycle largely resembles the flicker profile obtained by the community detection algorithm in **Section 3.2** (**Figures 17** and **19**), but the mean TRN cycle depicted here does not resemble any particular community. This may partially explain our negative results from **Section 3.2**: with a more fine-grained partitioning of the data, the participation coefficient might do a better job at locating voxels which serve as bridges between identified communities. This flicker profile would not correlate well with either Community 2 or 3 from **Figure 17** or **19**, depressing the *P* value of the TRN.

Second, these effects would be considerably stronger if I were to exclude S2 and S5. S2 appears to present the same shape of flicker profile, but with a considerably shifted phase. While some of this effect might be due to differences in BOLD lag between participants, which could not be corrected for due to the design of the stimulus, the difference between S2 and S1, 3, 4, & 6 is about 6 s: far greater than the maximum ~2.5 s BOLD lag variability previously observed between participants (Handwerker et al., 2004; although see Aguirre et al., 1998, which found lags in visual cortex up to 5 s). These results could be explained as a difference in flicker tuning between participants combined with differential BOLD lags. If flicker indeed gives us access to neural resonance, we would expect to see variability in the resonance of neural populations, which known to be variable in neural populations between participants. For instance, the optimal resonant frequency of 16 cat LGN were found to fall somewhere between 2 and 8 Hz (Mukherjee and Kaplan, 1995). The bias towards

higher flicker frequencies are clearly visible in S2's LGN in **Figure 8**. The results in S5 are less easy to explain away, however this participant's data was particularly atypical at all levels of analysis, including basic retinotopy (**Figures 8 and 17** clearly demonstrate the differences inherent to this participant). It may be that this data is different for an uninteresting reason such as fatigue at the time of scanning, a non-obvious mis-registration, or unusual blood vessel structure (see Handwerker et al., 2012 for a review).

Third, while the LGN and TRN appear to be co-active when the retina is presented with low-frequency flicker, the most consistent negative correlation between the seed TRN and mean LGN flicker cycle across participants was not in the spindle frequency range, but rather between 20–60 Hz. This is unexpected in the context of previous research, which has studied the LGN-TRN complex in isolation. In these studies, the TRN is most associated with spindle-frequency activity, and the 10 Hz signal is a well-understood thalamo-cortical phenomenon in the visual system (Fuentealba et al., 2004; Halassa et al., 2011; Pinault, 2004). There is, however, also some evidence of γ-frequency activity in the cat TRN during a motor task, and stimulation of the TRN with 500 Hz trains also evokes topographically-matched γ activity in the somatosensory and auditory cortex (Macdonald et al., 1998; Marczynski et al., 1984). Neither of these studies looked for such an effect in the visual TRN. Interestingly, in these data, we can see that the MGN (and putamen in some participants) is preferentially driven when the retina is presented with β to γ-frequency flicker (**Figures 7** and **8**).

A possible explanation for our findings is that driving the retina at these frequencies entrains non-visual cortical regions resonant at those frequencies–this would also explain the pulvinar responses observed in these data. The TRN is implicated in cross-modal filtering in response to ongoing attention demands (McAlonan et al., 2006, 2008). Here, the proposal is that the TRN prevents non-relevant incoming sensory information from interrupting an ongoing process. Cortico-thalamic β to γ-frequency oscillations are associated with attention, the transient integration of stereotypical

resting-state networks, and the execution of visual-motor behavior (de Pasquale et al., 2012; Saalmann et al., 2012; Womelsdorf et al., 2006). In this case, β to γ-frequency flicker may be driving a nonsense oscillation in the cortex, mimicking an ongoing attention process and entraining the TRN to inhibit visual input to the LGN. This feedback from the cortex would also explain the responses recorded in the MGN and pulvinar.

Furthermore, this effect acts in a lateralized fashion. That is, this flicker-frequency dependent TRN-mediated inhibition of the LGN is driven by flicker presented to the contralateral LGN (**Figures 25 & 26**). This finding has a few implications. First, the function of layer V to I contralateral inhibition observed in the cortex may extend to the thalamus, requiring the TRN (Palmer et al., 2012). While the findings of Marlinski et al., 2012 found TRN-mediated lateral inhibition in the walking cat, this serves as the first evidence of such a mechanism in the human brain, and also in the passive state. Curiously, while we found a spindle frequency effect in the TRN, like in the walking cat, but we also found a robust β to γ-frequency flicker response which drove the most dramatic suppression of the LGN.

I also plotted the full run of the seed TRN voxels alongside the ipsilateral LGN time series maximally negatively correlated with it. I am unable to average these time series due to the retinotopic signal being out of phase across the LGN, which would render the average meaningless. As these seeds represent different points in the retinotopic map, it isn't easy to compare between them. However, it is at least obvious that these regions are only strongly negatively correlated for a small fraction of the run, corroborating findings evident in the flicker profile response (**Figure 28**). Further, we can visually confirm that these TRN seeds respond at the retinotopic frequency, and are clearly amplitude modulated.

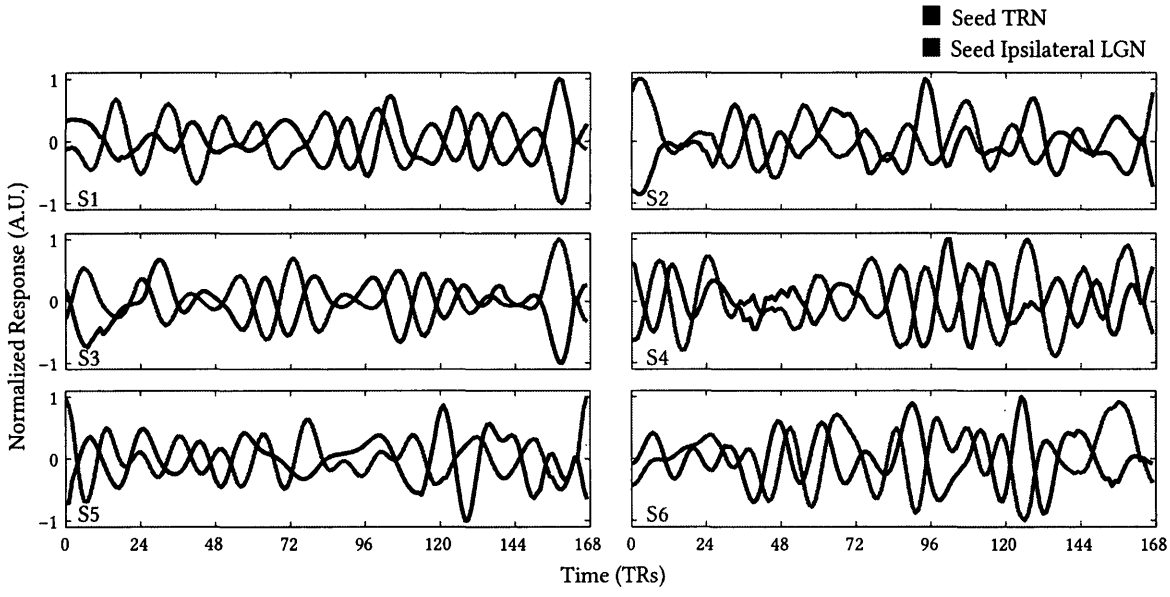

Figure 29. Seed time series from the TRN and maximally negatively correlated ipsilateral LGN.

To summarize, retinotopic analysis of the unsmoothed data (which prevents contamination from the adjacent pulvinar structures) revealed evidence for representations of each visual field in each TRN, but this finding could also have been due to noise. The attempt to triangulate the TRN as a structure uniquely expressing both low and high flicker resonances was unsuccessful, perhaps because both the TRN and pulvinar are heterogeneous structures. However, this approach did successfully identify the LGN as a structure with low-frequency flicker resonance, as well as identify the MGN as a structure with high-frequency flicker resonance, which suggests the proposed data-driven approach is valid. A seed-based correlation analysis in the TRN identified a sub-population of voxels exhibiting the expected bimodal flicker profile, suggesting these voxels mediate low and high-frequency resonance. These voxels also unexpectedly represent the visual field contralateral to that represented by the adjacent LGN, suggesting that the results of the initial retinotopic analysis were not due to noise, and may actually be the result of multiple neural populations in the TRN which express different resonance properties and visual receptive fields. Furthermore, it suggests the TRN is somehow involved in inter-hemispheric regulation of neural activity.

Chapter 4

An Overview of the Results and Future Experiments

Here I attempt to summarize and contextualize the results presented in Chapter 3. In **Section 4.1**, I give a brief overview of the conducted analysis and their potential implications. In **Section 4.2**, I briefly discuss the limitations of this experimental paradigm, and suggest improvements to the design. I also suggest a series of experiments that would be prudent given the results of the current experiment. In **Section 4.3**, I provide a brief summary of the major findings.

4.1 Properties of the Thalamic Reticular Nucleus

I attempted in general to localize the human TRN using fMRI. In order to verify that the method was successful, I compared the recorded response properties of the TRN with previous work using invasive electrical methods. One of the crucial assumptions made in this course of analysis is that the tremotopic stimulus allows one to access the resonance properties of a neural population, and is proposed as a non-invasive analogous method to the ZAP function used in single cell electrophysiology (Hutcheon and Yarom, 2000). In these analysis, we consistently identified LGN as being resonant at the expected α-frequency (**Figures 7–9, 11–13;** Hughes et al., 2004), and show that one is able to largely segment the first order visual thalamic nuclei in a data driven way using only the flicker profiles of the retinotopic voxels and a community-detection algorithm, providing independent verification that tremotopy and the flicker profile does indeed provide some insight into the resonance properties of neural populations, at least in early visual regions.

The retinotopic analysis suggests two general properties of the TRN. First, both the right and left TRN represent the entire visual field (**Figure 13**). This corroborates early labelling studies, which determined that this representation of the contralateral hemifield is due to cortical projections to the TRN, not projections from the LGN

(Montero et al., 1977). This seems a likely explanation for the present findings. Second, the TRN has a broad flicker frequency tuning, but has a slight preference for 7.5 Hz and 15-30 Hz, even when using unsmoothed data (so there is presumably little chance these flicker properties are due to the proximity with the LGN). The results of this analysis alone, however, are not convincing: we might see both of these results from a Fourier analysis by chance if the voxels were simply noisy. To the best of the author's knowledge, functional retinotopy of the TRN has not been described before. For example, two recent experiments implicating the TRN in attentional modulation of the LGN analyzed single unit activity in the TRN adjacent to the LGN, and also the activity of neurons with overlapping receptive fields in the ipsilateral LGN (McAlonan et al., 2006, 2008). The present results suggest concurrent recording from the TRN and LGN on each side of the brain may elucidate interesting interhemispheric properties of spatial or object-based attention.

I next sought to identify the TRN as a hub in the thalamic network. The rationale was that the visual sector of the TRN should mediate communication between first order and second order thalamic nuclei (the LGN and pulvinar, respectively (Lam and Sherman, 2007). While the visually responsive regions of the pulvinar and their associated resonance properties are not as well defined as in the LGN, one recent study found that pulvinar regions with strong connections to high-order visual cortex (V4 and the temporal-occipital area) express thalamo-cortical LFP coherence in the α to β frequency range (8–15 Hz; Saalmann et al., 2012). I therefore expected the TRN to respond to and mediate communication with both of these structures by expressing a flicker tuning to both of these frequencies. I attempted triangulate the TRN using a graph-theoretic statistic known as the participation coefficient. This analysis found no difference between the TRN and pulvinar. One possible explanation for this was that aforementioned community detection algorithm was not able to distinguish between these α to β frequencies: the LGN in some participants even expressed preference for 15 Hz (S2 and S6; **Figure 22**). Another explanation is that the pulvinar itself is a hub in the thalamo-cortical system, and due to the restricted field of view of the present

experiment (covering only a small region of the thalamus, and none of the cortex) the multi-modal network properties of the TRN were not accessible in this data set. Therefore, between the much larger number of pulvinar voxels sampled, and the improper coverage of the thalamic network, this analysis method was unable to identify any network properties unique to the TRN.

Finally, I sought to further probe the network properties of the TRN using a seed-based correlation analysis while ignoring the influence of the pulvinar completely. I found two interesting properties of the TRN. First, I found evidence that the TRN participates in a lateral inhibitory circuit: drive of the contralateral LGN was correlated with BOLD increases in the TRN, and BOLD decreases in the ipsilateral LGN. This corroborates previous work (Marlinski et al., 2012), and also strengthens the present retinotopy findings: if the equal representation of the entire visual field in the TRN was simply due to noise, it would seem very unlikely that these negatively correlated patterns would be apparent in all participants in such similar anatomical locations (**Figures 25 & 26**).

Second, I found a preference in the TRN to flicker in the 20–60 Hz range (**Figure 28**). During this time, the LGN was also substantially less driven than average. This relationship seems hard to reconcile with previous work. Specifically, driving one hemifield with β to γ-frequency flicker drives appears to drive the contralateral TRN, which possibly inhibits the associated LGN. I propose two possible explanations for this result.

One explanation follows from the unexpected observation that this stimulus drove the MGN at high flicker frequencies in most participants (20–60 Hz; **Figure 8**). At rest, the intrinsic oscillatory properties of the auditory cortex match the temporal properties of speech, with significant power in the δ, θ, & γ-frequency bands (Giraud et al., 2007; Morillon et al., 2010). It has also recently been shown that abnormalities in the low-gamma component underlie deficits in dyslexic verbal memory performance: those with dyslexia tend to entrain to acoustic modulations closer to 40 Hz, not 30 (Lehongre

et al., 2011). Therefore, low-gamma resonance is not only established in the auditory system, but it appears to be relevant to the active processing of the auditory world. Interestingly, the MGN, but not inferior colliculus, express flicker resonance in a similar frequency range in this data. This suggests that the origin of MGN modulation here is likely cortical in origin. Therefore, if low γ-frequency flicker is, in fact, driving the auditory cortex, this may result in cross-modal activity in the visual sector of the TRN. Previous work has shown the visual sector of the TRN is modulated in a selective-attention task involving audio-visual stimuli (with the distracter being either auditory or visual, and the target being the opposite; McAlonan et al., 2006). While any flicker-driven oscillation in the auditory system here would have no behavioural relevance, it remains a possibility that this nonsense signal is responsible for a cross-modal inhibition of visual inputs. If this is true, it suggests two things. First, it suggests that attention-driven cross-modal inhibition may operate in a lateralized fashion, as the inhibitory effect was to the hemifield contralateral to that stimulated. Second, it suggests that driving the visual system with different flicker frequencies entrains non-visual subsystems of the brain, and that this is determined by the resonance properties of a given neural population. Unexpectedly, this tremotopic stimulus may have the ability to map the resonance properties of a great deal more of the brain than simply the visual regions.

A second explanation follows from a line of research demonstrating that sub-perceptual γ-frequency flicker can bias the allocation of attention (Bauer et al., 2009; Cheadle et al., 2011; but see Bauer et al., 2012; van Diepen et al., 2010). The logic in this argument is that the visual cortex, which accurately entrains to oscillatory stimuli up to 120 Hz in the tree shrew, and 100 Hz in both the monkey and human (Herrmann, 2001; Veit et al., 2011; Williams et al., 2004), may be driven to produce phase-locked γ-frequency oscillations if driven by the correct flicker frequency. Phase-locked γ-band activity is known to be important for the propagation of task relevant signals between associated visual regions (Bosman et al., 2012), and therefore this sort of cortical activity may also entrain TRN neural populations, if the structure indeed plays a role in sensory

filtering and selective attention (Crick, 1984; McAlonan et al., 2008). In this case, γ-band flicker may have entrained attention to a particular point in visual space, driving suppression of the non-attended hemifield via the TRN.

Regardless of the explanation for this curious result, it should be taken in context with previous work showing some lateralized function of the TRN. One group found the TRN to be necessary for the 'validity effect' (the decrease in reaction time seen when a valid cue is presented in the same spatial location as an eventual target). Specifically, the validity effect was abolished in the rat when the TRN was completely lesioned contralateral to the visual field containing the cue/target pair (Weese et al., 1999). They also found no change in overall accuracy, suggesting this effect is not due to some sensory deficit. This result suggests the full retinotopic map observed in the present experiment may not be functionally equivalent across hemispheres, as they cannot compensate for one another. In light of the lateral-inhibitory BOLD connectivity observed here, it might be necessary for us to carefully interpret the negative BOLD response. In Weese et al., 1999, the TRN ipsilateral to the LGN receiving visual drive was responsible for the decreased observed reaction time. Therefore, filtering the irrelevant contralateral spatial location did not produce the decreased reaction time, but rather some ongoing process between the LGN and TRN representing the visual location of the target. In this experiment I have found some evidence that the two structures express negative BOLD correlations in some instances, but this should not be interpreted necessarily as overall decreased activity within the TRN. While the present experiment was not an attention task, and therefore it may well be that the present findings would be reversed in the case of an overt attention task, it may also imply that the TRN switches to a more specific, and less metabolically-demanding, firing mode during attentional filtering of the sensory world.

4.2 Problems With the Present Design and Future Experiments

Due to the task-free nature of this experiment, it remains impossible to know whether the presentation of γ-band flicker drove an attentional signal itself, whether

the participants directed attention towards the stimulus during these flicker frequencies intentionally, or we have simply mapped potential pathways for known attention phenomenon, but not observed any effects due to attention. A simple attention task outside of the MRI could be used to determine whether attention can indeed be entrained using particular flicker frequencies, by attempting to suppress the allocation of spatial attention with full-field flicker at different (but similar) frequencies. In the case that attention is associated with a small range of γ-band frequencies in early visual areas, one might be able to produce a tuning curve of distraction as the distracter flicker frequency deviates slightly from the resonant γ-frequency of the primary visual cortex. The logic of this experiment is that if coherent γ-frequency activity in disparate cortical regions is required for attention and flicker can drive meaningful oscillations in the brain, one might be able to inhibit spatial attention by entraining the entire visual field to the same frequency.

The lack of BOLD lag estimates in the present experiment posed significant problems for comparing the flicker profiles between participants (**Figures 17, 20, 22, & 28**). Due to the nature of the stimulus, BOLD lag would be best addressed with a second series of scans. I propose that each participant be presented with short presentations of flickering checkerboard stimuli in a blocked design as has been used before to estimate BOLD lag across participants (~ 200 ms; Handwerker et al., 2004). Next an HRF model would be fit to all voxels, in order to determine which time from stimulus onset gives the smallest error in the fit. Additionally, I would propose that this be done separately for all flicker frequencies present in the tremotopic stimulus, to measure the lags in regions not normally seen in visually driven experiments. Furthermore, one could observe whether BOLD lag is variable with presented flicker frequency.

While the EPI data collected here is of very high resolution, brain coverage was small. A number of the negative network-based results here are hard to interpret in light of this. In particular, we know the TRN receives inputs from every sensory system, the brainstem, the motor system, and limbic system in multiple mammalian species (Conley et al., 1991; Crabtree, 1992a, 1992b, 1996; Guillery and Harting, 2003; Künzle,

1976; Lozsádi, 1994; Montero et al., 1977; Pinault et al., 1995a, 1995b; Wang et al., 2001; Zikopoulos and Barbas, 2006, 2012). Our unexpected findings in the MGN suggest the tremotopic stimulus drives non-visual areas, and may entrain many areas of the cortex. If this is true, a whole-brain network constructed from high-resolution thalamic data and standard-resolution cortical data may be required for the expected network properties to be seen in the TRN. This sort of experiment would require the development of some new analytic techniques outside the scope of his paper in order to prevent biasing the results towards the higher-resolution data (as would be expected with more data points collected in closer spatial proximity; Power et al., 2011), and equalizing the SNR between these two data-sets. Physiological noise is a dominant source of noise in low-resolution data, but in high resolution data, thermal noise dominates (Triantafyllou et al., 2005). Therefore, sufficient temporal resolution and appropriate time series filtering will be required to equalize the properties of the noise distributions between these two data sets. Care should also be taken to use similar, if not identical, MRI sequences for both scans in order to minimize other differences between the two data sets which might be hard to detect without extensive tests. Fortunately, the simple stimulus-driven and task-free nature of the tremotopy design will make the combination of datasets across multiple resolutions otherwise trivial.

Finally, the BOLD signal does not carry much, if any, information about the LFP that gives rise to it (Goense and Logothetis, 2008; Logothetis, 2008; Logothetis et al., 2001; but see Magri et al., 2012). In these experiments we find some suggestion that visual flicker allows us to access the resonance properties of neural population, and some of the more interesting results depend on the concept of the flicker profile (**Figures 17** & **28**). These results demand further investigation using complementary methods. In conjunction with the whole-head cortical MRI data described above, it would be advantageous to collect complementary MEG data with the presentation of the same stimulus. This would allow one to verify that the actively presented flicker frequency is driving resonance in the brain at the same frequency. Unfortunately, the possibility remains that the BOLD signal seen here in downstream, particularly non-

visual areas such as the TRN and MGN, are actually expressing oscillations at a different frequency than that used to drive the retina. In the thalamus, this possibility might be best addressed using invasive electrical recordings. In both cases, it seems premature to conclude that we can access the exact resonant frequency of non-visual neural populations using visual flicker without on-line electrical data to support this claim. A second, MRI-only approach might be to look for correspondence of resonance in sensory areas using modulated stimuli in the different senses. For example, the present finding that 30 Hz flicker drives the MGN is in agreement with recent work demonstrating 30 Hz modulated auditory stimuli also preferentially entrains the auditory cortex in healthy controls (Lehongre et al., 2011). A satisfying proof-of-concept might be to see whether α-frequency modulated sound waves entrain early visual regions.

Finally, I did not collect eye-tracking or physiological data. This poses a potential confound in the interpretation of the results. Just as participants were found to move their heads in sync with flicker modulation (**Figure 4**), it seems a likely possibility that at least some of the participants were more likely to blink or saccade from fixation during particular flicker frequencies, both of which produce robust BOLD signals (Hupé et al., 2012; Kimmig et al., 2001). Any variability of these uninteresting events between participants could explain some or all of the flicker profile variability observed (**Figures 17, 20, 22, & 28**), and consistent behaviour between participants could explain some of these results, including flicker-specific drive of the TRN (**Figures 17, 20, & 28**). Heart rate and breathing variability are both known to modulate in participants performing concentration tasks (Negoescu et al., 1993), and this effect may extend to a task-free but stimulus driven case such as this one. Future work should ensure to collect these measurements to rule out any influence of physiological noise.

4.3 Summary of the Work

This work lays the foundation for future work on the TRN, and also provides compelling evidence that visual flicker can be used to segment neural populations

across their resonance properties using MRI, which is not information typically available in the BOLD signal. The TRN does show interesting resonance properties, exhibiting two different flicker resonances in the BOLD signal as predicted. However, its small size makes these features undetectable using data-driven segmentation approaches alone. For this reason, the collection of high-resolution proton density images is advantageous for the detection of TRN BOLD activity.

It is clear from the combined retinotopy and seed-based correlation analysis that some interesting retinotopic properties exist in the TRN, particularly the representation of the contralateral visual field. These functional properties have not been described before in the TRN, and are likely driven by descending cortical inputs. This seem particularly likely, as these contralateral TRN voxels are also responsive to 30 Hz flicker, and have negatively-correlated flicker profiles with the LGN in these high frequencies. These methods and findings should be of interest to those doing invasive recordings of the TRN, as they will allow one to triangulate ROIs for recording sessions, as well as clinical and cognitive psychologists who might like to define the role of the TRN in the manifestation of disorders or attention phenomenon. Furthermore, this work represents the first large-scale characterization of TRN activity in any animal, and highlights the need for future studies to investigate the network properties of this poorly-characterized structure.

References

Aguirre, G.K., Zarahn, E., and D'Esposito, M. (1998). The variability of human BOLD hemodynamic responses. Neuroimage *8*, 360–369.

Asanuma, C. (1994). GABAergic and pallidal terminals in the thalamic reticular nucleus of squirrel monkeys. Exp. Brain Res. *101*, 439–451.

Ash, T., Suckling, J., Walter, M., Ooi, C., Tempelmann, C., Carpenter, A., and Williams, G. (2011). Detection of physiological noise in resting state fMRI using machine learning. Hum. Brain Mapp. 985–998.

Bal, T., and McCormick, D.A. (1993). Mechanisms of oscillatory activity in guinea-pig nucleus reticularis thalami in vitro: a mammalian pacemaker. J. Physiol. *468*, 669–691.

Bauer, F., Cheadle, S.W., Parton, A., Müller, H.J., and Usher, M. (2009). Gamma flicker triggers attentional selection without awareness. Proc. Natl. Acad. Sci. *106*, 1666–1671.

Bauer, M., Akam, T., Joseph, S., Freeman, E., and Driver, J. (2012). Does visual flicker phase at gamma frequency modulate neural signal propagation and stimulus selection? J. Vis. *12*.

Bazhenov, M., Timofeev, I., Steriade, M., and Sejnowski, T. (1999). Self–sustained rhythmic activity in the thalamic reticular nucleus mediated by depolarizing GABAA receptor potentials. Nat. Neurosci. *2*, 168–174.

Bazhenov, M., Timofeev, I., Steriade, M., and Sejnowski, T. (2000). Spiking-bursting activity in the thalamic reticular nucleus initiates sequences of spindle oscillations in thalamic networks. J. Neurophysiol. *84*, 1076–1087.

Beauchamp, M.S., Argall, B.D., Bodurka, J., Duyn, J.H., and Martin, A. (2004a). Unraveling multisensory integration: patchy organization within human STS multisensory cortex. Nat. Neurosci. *7*, 1190–1192.

Beauchamp, M.S., Lee, K.E., Argall, B.D., and Martin, A. (2004b). Integration of auditory and visual information about objects in superior temporal sulcus. Neuron *41*, 809–823.

Bender, D.B. (1981). Retinotopic organization of macaque pulvinar. J. Neurophysiol. *46*, 672–693.

Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. J. R. Stat. Soc. Ser. B Methodol. 289–300.

Blondel, V.D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. J. Stat. Mech. Theory Exp. *2008*, P10008.

Bosman, C.A., Schoffelen, J.-M., Brunet, N., Oostenveld, R., Bastos, A.M., Womelsdorf, T., Rubehn, B., Stieglitz, T., De Weerd, P., and Fries, P. (2012). Attentional stimulus selection through selective synchronization between monkey visual areas. Neuron *75*, 875–888.

Breton, J.-D., and Stuart, G.J. (2012). Somatic and dendritic GABAB receptors regulate neuronal excitability via different mechanisms. J. Neurophysiol. *108*, 2810–2818.

Browne, S.H., Kang, J., Akk, G., Chiang, L.W., Schulman, H., Huguenard, J.R., and Prince, D.A. (2001). Kinetic and pharmacological properties of GABAA receptors in single thalamic neurons and GABAA subunit expression. J. Neurophysiol. *86*, 2312–2322.

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. *10*, 186–198.

Buzsáki, G., and Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science *304*, 1926–1929.

Buzsáki, G., Kaila, K., and Raichle, M. (2007). Inhibition and brain work. Neuron *56*, 771–783.

Canolty, R.T., Ganguly, K., Kennerley, S.W., Cadieu, C.F., Koepsell, K., Wallis, J.D., and Carmena, J.M. (2010). Oscillatory phase coupling coordinates anatomically dispersed functional cell assemblies. Proc. Natl. Acad. Sci. *107*, 17356–17361.

Castro-Alamancos, M.A. (2000). Origin of synchronized oscillations induced by neocortical disinhibition in vivo. J. Neurosci. *20*, 9195–9206.

Chacron, M.J., Doiron, B., Maler, L., Longtin, A., and Bastian, J. (2003). Non-classical receptive field mediates switch in a sensory neuron's frequency tuning. Nature *423*, 77–81.

Chai, X.J., Castañóna, A.N., Öngürb, D., and Whitfield-Gabrieli, S. (2012). Anticorrelations in resting state networks without global signal regression. Neuroimage *59*, 1420–1428.

Cheadle, S.W., Parton, A., Müller, H.J., and Usher, M. (2011). Subliminal gamma flicker draws attention even in the absence of transition-flash cues. J. Neurophysiol. *105*, 827–833.

Chen, W., Zhu, X.-H., Thulborn, K.R., and Ugurbil, K. (1999). Retinotopic mapping of lateral geniculate nucleus in humans using functional magnetic resonance imaging. Proc. Natl. Acad. Sci. *96*, 2430–2434.

Conley, M., Kupersmith, A.C., and Diamond, I.T. (1991). The organization of projections from subdivisions of the auditory cortex and thalamus to the auditory sector of the thalamic reticular nucleus in galago. Eur. J. Neurosci. *3*, 1089–1103.

Contreras, D., Destexhe, A., Sejnowski, T.J., and Steriade, M. (1997). Spatiotemporal patterns of spindle oscillations in cortex and thalamus. J. Neurosci. *17*, 1179–1196.

Cotillon-Williams, N., Huetz, C., Hennevin, E., and Edeline, J.-M. (2008). Tonotopic control of auditory thalamus frequency tuning by reticular thalamic neurons. J. Neurophysiol. *99*, 1137–1151.

Cox, R.W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. *29*, 162–173.

Crabtree, J.W. (1992a). The somatotopic organization within the rabbit's thalamic reticular nucleus. Eur. J. Neurosci. *4*, 1343–1351.

Crabtree, J.W. (1992b). The somatotopic organization within the cat's thalamic reticular nucleus. Eur. J. Neurosci. *4*, 1352–1361.

Crabtree, J.W. (1996). Organization in the somatosensory sector of the cat's thalamic reticular nucleus. J. Comp. Neurol. *366*, 207–222.

Crabtree, J.W., and Isaac, J.T.R. (2002). New intrathalamic pathways allowing modality-related and cross-modality switching in the dorsal thalamus. J. Neurosci. *22*, 8754–8761.

Crick, F. (1984). Function of the thalamic reticular complex: the searchlight hypothesis. Proc. Natl. Acad. Sci. *81*, 4586–4590.

Cueni, L., Canepari, M., Lujan, R., Emmenegger, Y., Watanabe, M., Bond, C.T., Franken, P., Adelman, J.P., and Luthi, A. (2008). T-type Ca2+ channels, SK2 channels and SERCAs gate sleep-related oscillations in thalamic dendrites. Nat. Neurosci. *11*, 683–692.

D'Souza, D.V., Auer, T., Strasburger, H., Frahm, J., and Lee, B.B. (2011). Temporal frequency and chromatic processing in humans: An fMRI study of the cortical visual areas. J. Vis. *11*.

Dale, A.M., and Sereno, M.I. (1993). Improved localizadon of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction: a linear approach. J. Cogn. Neurosci. *5*, 162–176.

Dale, A., Fischl, B., and Sereno, M.I. (1999). Cortical surface-based analysis: I. Segmentation and surface reconstruction. Neuroimage *9*, 179 – 194.

Deleuze, C., and Huguenard, J.R. (2006). Distinct electrical and chemical connectivity maps in the thalamic reticular nucleus: potential roles in synchronization and sensation. J. Neurosci. *26*, 8633–8645.

Deschênes, M., Timofeeva, E., Lavallée, P., and Dufresne, C. (2005). The vibrissal system as a model of thalamic operations. Prog. Brain Res. *149*, 31–40.

Devlin, J.T., Sillery, E.L., Hall, D.A., Hobden, P., Behrens, T.E.J., Nunes, R.G., Clare, S., Matthews, P.M., Moore, D.R., and Johansen-Berg, H. (2006). Reliable identification of the auditory thalamus using multi-modal structural analyses. Neuroimage *30*, 1112–1120.

Van Diepen, R.M., Born, S., Souto, D., Gauch, A., and Kerzel, D. (2010). Visual flicker in the gamma-band range does not draw attention. J. Neurophysiol. *103*, 1606–1613.

Van Dijk, K.R.A., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., and Buckner, R.L. (2010). Intrinsic functional connectivity as a tool for human connectomics: Theory, properties, and optimization. J. Neurophysiol. *103*, 297–321.

Van Dijk, K.R.A., Sabuncu, M.R., and Buckner, R.L. (2012). The influence of head motion on intrinsic functional connectivity MRI. Neuroimage *59*, 431–438.

Dubin, M.W., and Cleland, B.G. (1977). Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. J. Neurophysiol. *40*, 410–427.

Engel, S.A., Glover, G.H., and Wandell, B.A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb. Cortex *7*, 181–192.

Erchova, I., Kreck, G., Heinemann, U., and Herz, A.V.M. (2004). Dynamics of rat entorhinal cortex layer II and III cells: characteristics of membrane potential resonance at rest predict oscillation properties near threshold. J. Physiol. *560*, 89–110.

Estebanez, L., Boustani, S.E., Destexhe, A., and Shulz, D.E. (2012). Correlated input reveals coexisting coding schemes in a sensory cortex. Nat. Neurosci. *15*, 1691–1699.

Ferbert, A., Priori, A., Rothwell, J.C., Day, B.L., Colebatch, J.G., and Marsden, C.D. (1992). Interhemispheric inhibition of the human motor cortex. J. Physiol. *453*, 525–546.

Fischl, B., and Dale, A.M. (2000). Measuring the thickness of the human cerebral cortex from magnetic resonance images. Proc. Natl. Acad. Sci. U. S. A. *97*, 11050–11055.

Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al. (2002). Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. Neuron *33*, 341–355.

Fischl, B., Salat, D.H., Kouwe, A.J.W. van der, Makris, N., Ségonne, F., Quinn, B.T., and Dale, A.M. (2004). Sequence-independent segmentation of magnetic resonance images. Neuroimage *23*, S69–S84.

Fortunato, S. (2010). Community detection in graphs. Phys. Reports *486*, 75–174.

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cogn. Sci. *9*, 474–480.

Friston, K.J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., and Turner, R. (1998). Event-related fMRI: characterizing differential responses. Neuroimage *7*, 30–40.

Fuentealba, P., Timofeev, I., and Steriade, M. (2004). Prolonged hyperpolarizing potentials precede spindle oscillations in the thalamic reticular nucleus. Proc. Natl. Acad. Sci. U. S. A. *101*, 9816–9821.

Fujita, N., Tanaka, H., Takanashi, M., Hirabuki, N., Abe, K., Yoshimura, H., and Nakamura, H. (2001). Lateral geniculate nucleus: Anatomic and functional identification by use of MR imaging. Am. J. Neuroradiol. *22*, 1719–1726.

Funke, K., and Eysel, U.T. (1998). Inverse correlation of firing patterns of single topographically matched perigeniculate neurons and cat dorsal lateral geniculate relay cells. Vis. Neurosci. *15*, 711–729.

Giraud, A.-L., Kleinschmidt, A., Poeppel, D., Lund, T.E., Frackowiak, R.S.J., and Laufs, H. (2007). Endogenous cortical rhythms determine cerebral specialization for speech perception and production. Neuron *56*, 1127–1134.

Glover, G.H. (1999). Deconvolution of impulse response in event-related bold fMRI. Neuroimage *9*, 416–429.

Goense, J.B.M., and Logothetis, N.K. (2008). Neurophysiology of the BOLD fMRI signal in awake monkeys. Curr. Biol. *18*, 631–640.

Gonçalves, S.I., de Munck, J.C., Pouwels, P.J.W., Schoonhoven, R., Kuijer, J.P.A., Maurits, N.M., Hoogduin, J.M., Van Someren, E.J.W., Heethaar, R.M., and Lopes da Silva, F.H. (2006). Correlating the alpha rhythm to BOLD using simultaneous EEG/fMRI: Inter-subject variability. Neuroimage *30*, 203–213.

Gonzalez-Castillo, J., Saad, Z.S., Handwerker, D.A., Inati, S.J., Brenowitz, N., and Bandettini, P.A. (2012). Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. Proc. Natl. Acad. Sci. *109*, 5487–5492.

Guillery, R.W., and Harting, J.K. (2003). Structure and connections of the thalamic reticular nucleus: Advancing views over half a century. J. Comp. Neurol. *463*, 360–371.

Guillery, R.W., and Sherman, S.M. (2002). Thalamic relay functions and their role in corticocortical communication: Generalizations from the visual system. Neuron *33*, 163–175.

Guimera, R., and Nunes Amaral, L.A. (2005). Functional cartography of complex metabolic networks. Nature *433*, 895–900.

Gutierrez, C., Cola, M.G., Seltzer, B., and Cusick, C. (2000). Neurochemical and connectional organization of the dorsal pulvinar complex in monkeys. J. Comp. Neurol. *419*, 61–86.

Halassa, M.M., Siegle, J.H., Ritt, J.T., Ting, J.T., Feng, G., and Moore, C.I. (2011). Selective optical drive of thalamic reticular nucleus generates thalamic bursts and cortical spindles. Nat. Neurosci. *14*, 1118–1120.

Han, X., Jovicich, J., Salat, D., van der Kouwe, A., Quinn, B., Czanner, S., Busa, E., Pacheco, J., Albert, M., Killiany, R., et al. (2006). Reliability of MRI-derived measurements of human cerebral cortical thickness: The effects of field strength, scanner upgrade and manufacturer. Neuroimage *32*, 180–194.

Handwerker, D.A., Ollinger, J.M., and D'Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. Neuroimage *21*, 1639–1651.

Handwerker, D.A., Gonzalez-Castillo, J., D'Esposito, M., and Bandettini, P.A. (2012). The continuing challenge of understanding and modeling hemodynamic variation in fMRI. Neuroimage *62*, 1017–1023.

Haynes, J.-D., Deichmann, R., and Rees, G. (2005). Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus. Nature *438*, 496–499.

Hebb, D.O. (1949). The Organization of Behavior: A Neuropsychological Theory (Lawrence Erlbaum).

Herrmann, C.S. (2001). Human EEG responses to 1–100 Hz flicker: Resonance phenomena in visual cortex and their potential correlation to cognitive phenomena. Exp. Brain Res. *137*, 346–353.

Hilgetag, C.C., Theoret, H., and Pascual-Leone, A. (2001). Enhanced visual spatial attention ipsilateral to rTMS-induced "virtual lesions" of human parietal cortex. Nat. Neurosci. *4*, 953–957.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. J. Physiol. *148*, 574–591.

Hughes, S.W., Lörincz, M., Cope, D.W., Blethyn, K.L., Kékesi, K.A., Parri, H.R., Juhász, G., and Crunelli, V. (2004). Synchronized oscillations at α and θ frequencies in the lateral geniculate nucleus. Neuron *42*, 253–268.

Huntsman, M.M., Leggio, M.G., and Jones, E.G. (1996). Nucleus-specific expression of GABAA receptor subunit mrnas in monkey thalamus. J. Neurosci. *16*, 3571–3589.

Huntsman, M.M., Porcello, D.M., Homanics, G.E., DeLorey, T.M., and Huguenard, J.R. (1999). Reciprocal inhibitory connections and network synchrony in the mammalian thalamus. Science *283*, 541–543.

Hupé, J.-M., Bordier, C., and Dojat, M. (2012). A BOLD signature of eyeblinks in the visual cortex. Neuroimage *61*, 149–161.

Hutcheon, B., and Yarom, Y. (2000). Resonance, oscillation and the intrinsic frequency preferences of neurons. Trends Neurosci. *23*, 216–222.

Jones, E.G. (1975). Some aspects of the organization of the thalamic reticular complex. J. Comp. Neurol. *162*, 285–308.

Jones, E.G. (2007). The Thalamus (Cambridge University Press).

Jones, E.G. (2009). Synchrony in the Interconnected Circuitry of the Thalamus and Cerebral Cortex. Ann. N. Y. Acad. Sci. *1157*, 10–23.

Jourdain, A., Semba, K., and Fibiger, H.C. (1989). Basal forebrain and mesopontine tegmental projections to the reticular thalamic nucleus: an axonal collateralization and immunohistochemical study in the rat. Brain Res. *505*, 55–65.

Juhász, G., Emri, Z., Kékesi, K.A., Salfay, O., and Crunelli, V. (1994). Blockade of thalamic GABAB receptors decreases EEG synchronization. Neurosci. Lett. *172*, 155–158.

Kastner, S., O'Connor, D.H., Fukui, M.M., Fehd, H.M., Herwig, U., and Pinsk, M.A. (2004). Functional imaging of the human lateral geniculate nucleus and pulvinar. J. Neurophysiol. *91*, 438–448.

Kim, U., Sanchez-Vives, M.V., and McCormick, D.A. (1997). Functional dynamics of GABAergic inhibition in the thalamus. Science *278*, 130–134.

Kimmig, H., Greenlee, M., Gondan, M., Schira, M., Kassubek, J., and Mergner, T. (2001). Relationship between saccadic eye movements and cortical activity as measured by fMRI: quantitative and qualitative aspects. Exp. Brain Res. *141*, 184–194.

Künzle, H. (1976). Thalamic projections from the precentral motor cortex in Macaca fascicularis. Brain Res. *105*, 253–267.

Lam, Y.-W., and Sherman, S.M. (2007). Different topography of the reticulothalmic inputs to first- and higher-order somatosensory thalamic relays revealed using photostimulation. J. Neurophysiol. *98*, 2903–2909.

Lam, Y.-W., and Sherman, S.M. (2010). Functional organization of the somatosensory cortical layer 6 feedback to the thalamus. Cereb. Cortex *20*, 13–24.

Lam, Y.-W., Nelson, C.S., and Sherman, S.M. (2006). Mapping of the functional interconnections between thalamic reticular neurons using photostimulation. J. Neurophysiol. *96*, 2593–2600.

Landisman, C.E., Long, M.A., Beierlein, M., Deans, M.R., Paul, D.L., and Connors, B.W. (2002). Electrical synapses in the thalamic reticular nucleus. J. Neurosci. *22*, 1002–1009.

Lee, J.H., Durand, R., Gradinaru, V., Zhang, F., Goshen, I., Kim, D.-S., Fenno, L.E., Ramakrishnan, C., and Deisseroth, K. (2010). Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. Nature *465*, 788–792.

Lehongre, K., Ramus, F., Villiermet, N., Schwartz, D., and Giraud, A.-L. (2011). Altered low-gamma sampling in auditory cortex accounts for the three main facets of dyslexia. Neuron *72*, 1080–1090.

Lieberman, M.D., and Cunningham, W.A. (2009). Type I and Type II error concerns in fMRI research: re-balancing the scale. Soc. Cogn. Affect. Neurosci. *4*, 423–428.

Liu, X.-B., and Jones, E.G. (1999). Predominance of corticothalamic synaptic inputs to thalamic reticular nucleus neurons in the rat. J. Comp. Neurol. *414*, 67–79.

Liu, C.-S.J., Miki, A., Hulvershorn, J., Bloy, L., Gualtieri, E.E., Liu, G.T., Leigh, J.S., Haselgrove, J.C., and Elliott, M.A. (2006). Spatial and temporal characteristics of physiological noise in fMRI at 3T. Acad. Radiol. *13*, 313–323.

Logothetis, N.K. (2008). What we can do and what we cannot do with fMRI. Nature *453*, 869–878.

Logothetis, N.K., and Wandell, B.A. (2004). Interpreting the BOLD signal. Annu Rev Physiol *66*, 735–769.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature *412*, 150–157.

Long, M.A., Landisman, C.E., and Connors, B.W. (2004). Small clusters of electrically coupled neurons generate synchronous rhythms in the thalamic reticular nucleus. J. Neurosci. *24*, 341–349.

Lozsádi, D.A. (1994). Organization of cortical afferents to the rostral, limbic sector of the rat thalamic reticular nucleus. J. Comp. Neurol. *341*, 520–533.

Macdonald, K.D., Fifkova, E., Jones, M.S., and Barth, D.S. (1998). Focal stimulation of the thalamic reticular nucleus induces focal gamma waves in cortex. J. Neurophysiol. *79*, 474–477.

MacLean, J.N., Watson, B.O., Aaron, G.B., and Yuste, R. (2005). Internal dynamics determine the cortical response to thalamic stimulation. Neuron *48*, 811–823.

Magri, C., Schridde, U., Murayama, Y., Panzeri, S., and Logothetis, N.K. (2012). The amplitude and timing of the BOLD signal reflects the relationship between local field potential power at different frequencies. J. Neurosci. *32*, 1395–1407.

Mai, J.K., Paxinos, G., and Voss, T. (2007). Atlas of the Human Brain (Academic Press).

Marczynski, T.J., Burns, L.L., Livezey, G.T., Vimal, R.L.P., and Chen, E. (1984). Sleep and purposive behavior: Inverse deviations from randomness of neuronal firing patterns in the feline thalamus: A new form of homeostasis? Brain Res. *298*, 75–90.

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science *275*, 213–215.

Marlinski, V., Sirota, M.G., and Beloozerova, I.N. (2012). Differential gating of thalamocortical signals by reticular nucleus of thalamus during locomotion. J. Neurosci. *32*, 15823–15836.

McAlonan, K., Cavanaugh, J., and Wurtz, R.H. (2006). Attentional modulation of thalamic reticular neurons. J. Neurosci. *26*, 4444–4450.

McAlonan, K., Cavanaugh, J., and Wurtz, R.H. (2008). Guarding the gateway to cortex with attention in visual thalamus. Nature *456*, 391–394.

Meilă, M. (2007). Comparing clusterings—an information based distance. J. Multivar. Anal. *98*, 873–895.

Mitchell, J.F., Stoner, G.R., and Reynolds, J.H. (2004). Object-based attention determines dominance in binocular rivalry. Nature *429*, 410–413.

Montemurro, M.A., Rasch, M.J., Murayama, Y., Logothetis, N.K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. Curr. Biol. *18*, 375–380.

Montero, V.M., Guillery, R.W., and Woolsey, C.N. (1977). Retinotopic organization within the thalamic reticular nucleus demonstrated by a double label autoradiographic technique. Brain Res. *138*, 407–421.

Morillon, B., Lehongre, K., Frackowiak, R.S.J., Ducorps, A., Kleinschmidt, A., Poeppel, D., and Giraud, A.-L. (2010). Neurophysiological origin of human brain asymmetry for speech and language. Proc. Natl. Acad. Sci. *107*, 18688–18693.

Mukherjee, P., and Kaplan, E. (1995). Dynamics of neurons in the cat lateral geniculate nucleus: In vivo electrophysiology and computational modeling. J. Neurophysiol. *74*, 1222–1243.

Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., and Bandettini, P.A. (2009). The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? Neuroimage *44*, 893–905.

Negoescu, R.M., Csiki, I.E., Pafnote, M., and Wolf, S. (1993). Cortical control of sinus arrhythmia in man studied by spectral analysis. Integr. Physiol. Behav. Sci. *28*, 226–238.

Norrsell, U., Finger, S., and Lajonchere, C. (1999). Cutaneous sensory spots and the "law of specific nerve energies": history and development of ideas. Brain Res. Bull. *48*, 457–465.

Northoff, G., Walter, M., Schulte, R.F., Beck, J., Dydak, U., Henning, A., Boeker, H., Grimm, S., and Boesiger, P. (2007). GABA concentrations in the human anterior cingulate cortex predict negative BOLD responses in fMRI. Nat. Neurosci. *10*, 1515–1517.

Nuttall, A. (1981). Some windows with very good sidelobe behavior. Acoust. Speech Signal Process. IEEE Trans. *29*, 84–91.

Ogawa, S., Lee, T.M., Kay, A.R., and Tank, D.W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. *87*, 9868–9872.

Olbrich, S., Mulert, C., Karch, S., Trenner, M., Leicht, G., Pogarell, O., and Hegerl, U. (2009). EEG-vigilance and BOLD effect during simultaneous EEG/fMRI measurement. Neuroimage *45*, 319–332.

Osipova, D., Hermes, D., and Jensen, O. (2008). Gamma power is phase-locked to posterior alpha activity. PLoS One *3*, e3990.

Ozus, B., Liu, H.-L., Chen, L., Iyer, M.B., Fox, P.T., and Gao, J.-H. (2001). Rate dependence of human visual cortical response due to brief stimulation: An event-related fMRI study. Magn. Reson. Imaging *19*, 21–25.

Palmer, L.M., Schulz, J.M., Murphy, S.C., Ledergerber, D., Murayama, M., and Larkum, M.E. (2012). The cellular basis of GABAB-mediated interhemispheric inhibition. Science *335*, 989–993.

Paré, D., Hazrati, L.-N., Parent, A., and Steriade, M. (1990). Substantia nigra pars reticulata projects to the reticular thalamic nucleus of the cat: a morphological and electrophysiological study. Brain Res. *535*, 139–146.

De Pasquale, F., Della Penna, S., Snyder, A.Z., Marzetti, L., Pizzella, V., Romani, G.L., and Corbetta, M. (2012). A cortical core for dynamic integration of functional networks in the resting human brain. Neuron *74*, 753–764.

Paz, J.T., Bryant, A.S., Peng, K., Fenno, L., Yizhar, O., Frankel, W.N., Deisseroth, K., and Huguenard, J.R. (2011). A new mode of corticothalamic transmission revealed in the Gria4-/- model of absence epilepsy. Nat. Neurosci. *14*, 1167–1173.

Pinault, D. (2004). The thalamic reticular nucleus: structure, function and concept. Brain Res. Rev. *46*, 1–31.

Pinault, D., and Deschênes, M. (1998). Anatomical evidence for a mechanism of lateral inhibition in the rat thalamus. Eur. J. Neurosci. *10*, 3462–3469.

Pinault, D., Bourassa, J., and Deschênes, M. (1995a). The axonal arborization of single thalamic reticular neurons in the somatosensory thalamus of the rat. Eur. J. Neurosci. *7*, 31–40.

Pinault, D., Bourassa, J., and Deschênes, M. (1995b). Thalamic reticular input to the rat visual thalamus: a single fiber study using biocytin as an anterograde tracer. Brain Res. *670*, 147–152.

Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., et al. (2011). Functional network organization of the human brain. Neuron *72*, 665–678.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., and Petersen, S.E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage *59*, 2142–2154.

Previc, F.H. (1990). Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. Behav. Brain Sci. *13*, 519–542.

Proverbio, A.M., D'Aniello, G.E., Adorni, R., and Zani, A. (2011). When a photograph can be heard: Vision activates the auditory cortex within 110 ms. Sci Rep *1*.

Puil, E., Gimbarzevsky, B., and Miura, R.M. (1986). Quantification of membrane properties of trigeminal root ganglion neurons in guinea pigs. J. Neurophysiol. *55*, 995–1016.

Rager, G., and Singer, W. (1998). The response of cat visual cortex to flicker stimuli of variable frequency. Eur. J. Neurosci. *10*, 1856–1877.

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

Rubinov, M., and Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. Neuroimage *52*, 1059–1069.

Rubinov, M., and Sporns, O. (2011). Weight-conserving characterization of complex functional brain networks. Neuroimage *56*, 2068–2079.

Saad, Z.S., Glen, D.R., Chen, G., Beauchamp, M.S., Desai, R., and Cox, R.W. (2009). A new method for improving functional-to-structural MRI alignment using local Pearson correlation. Neuroimage *44*, 839–848.

Saalmann, Y.B., Pinsk, M.A., Wang, L., Li, X., and Kastner, S. (2012). The pulvinar regulates information transmission between cortical areas based on attention demands. Science *337*, 753–756.

Sanchez-Vives, M.V., and McCormick, D.A. (1997). Functional properties of perigeniculate inhibition of dorsal lateral geniculate nucleus thalamocortical neurons in vitro. J. Neurosci. *17*, 8880–8893.

Sanchez-Vives, M.V., Bal, T., and McCormick, D.A. (1997). Inhibitory interactions between perigeniculate GABAergic neurons. J. Neurosci. *17*, 8894–8908.

Schaeffer, S.E. (2007). Graph clustering. Comput. Sci. Rev. *1*, 27–64.

Schneider, K.A. (2011). Subcortical mechanisms of feature-based attention. J. Neurosci. *31*, 8643–8653.

Schneider, K.A., and Kastner, S. (2005). Visual responses of the human superior colliculus: A high-resolution functional magnetic resonance imaging study. J. Neurophysiol. *94*, 2491–2503.

Schneider, K.A., Richter, M.C., and Kastner, S. (2004). Retinotopic organization and functional subdivisions of the human lateral geniculate nucleus: A high-

resolution functional magnetic resonance imaging study. J. Neurosci. *24*, 8975–8985.

Schofield, C.M., Kleiman-Weiner, M., Rudolph, U., and Huguenard, J.R. (2009). A gain in GABAA receptor synaptic strength in thalamus reduces oscillatory activity and absence seizures. Proc. Natl. Acad. Sci. *106*, 7630–7635.

Segonne, F., Dale, A.M., Busa, E., Glessner, M., Salat, D., Hahn, H.K., and Fischl, B. (2004). A hybrid approach to the skull stripping problem in MRI. Neuroimage *22*, 1060–1075.

Sengpiel, F., Blakemore, C., and Harrad, R. (1995). Interocular suppression in the primary visual cortex: a possible neural basis of binocular rivalry. Vision Res. *35*, 179–195.

Shadlen, M.N., and Newsome, W.T. (1994). Noise, neural codes and cortical organization. Curr. Opin. Neurobiol. *4*, 569–579.

Sherman, S.M., and Guillery, R.W. (1996). Functional organization of thalamocortical relays. J. Neurophysiol. *76*, 1367–1395.

Sherrington, C. (1906). Observations on the scratch-reflex in the spinal dog. J. Physiol. *34*, 1.

Shmueli, K., van Gelderen, P., de Zwart, J.A., Horovitz, S.G., Fukunaga, M., Jansma, J.M., and Duyn, J.H. (2007). Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. Neuroimage *38*, 306–320.

Shu, Y., and McCormick, D.A. (2002). Inhibitory interactions between ferret thalamic reticular neurons. J. Neurophysiol. *87*, 2571–2576.

Shu, Y., Hasenstaub, A., Duque, A., Yu, Y., and McCormick, D.A. (2006). Modulation of intracortical synaptic potentials by presynaptic somatic membrane potential. Nature *441*, 761–765.

Sillito, A.M., and Jones, H.E. (2008). The role of the thalamic reticular nucleus in visual processing. Thalamus Relat. Syst. *4*, 1–12.

Singh, K.D., Smith, A.T., and Greenlee, M.W. (2000). Spatiotemporal frequency and direction sensitivities of human visual areas measured using fMRI. Neuroimage *12*, 550–564.

Sled, J.G., Zijdenbos, A.P., and Evans, A.C. (1998). A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Trans. Med. Imaging *17*, 87–97.

Smith, A.M., Lewis, B.K., Ruttimann, U.E., Ye, F.Q., Sinnwell, T.M., Yang, Y., Duyn, J.H., and Frank, J.A. (1999). Investigation of low frequency drift in fMRI signal. Neuroimage *9*, 526–533.

Smith, S.M., Miller, K.L., Moeller, S., Xu, J., Auerbach, E.J., Woolrich, M.W., Beckmann, C.F., Jenkinson, M., Andersson, J., Glasser, M.F., et al. (2012). Temporally-independent functional modes of spontaneous brain activity. Proc. Natl. Acad. Sci. *109*, 3131–3136.

Sohal, V.S., and Huguenard, J.R. (2003). Inhibitory interconnections control burst pattern and emergent network synchrony in reticular thalamus. J. Neurosci. *23*, 8978–8988.

Song, S., Miller, K.D., and Abbott, L.F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat. Neurosci. *3*, 919–926.

Sotero, R.C., and Trujillo-Barreto, N.J. (2007). Modelling the role of excitatory and inhibitory neuronal activity in the generation of the BOLD signal. Neuroimage *35*, 149–165.

Spaak, E., Bonnefond, M., Maier, A., Leopold, D.A., and Jensen, O. (2012). Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex. Curr. Biol. *22*, 2313–2318.

Steriade, M. (1997). Synchronized activities of coupled oscillators in the cerebral cortex and thalamus at different levels of vigilance. Cereb. Cortex *7*, 583–604.

Steriade, M., and Timofeev, I. (2003). Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. Neuron *37*, 563–576.

Steriade, M., Contreras, D., Curro Dossi, R., and Nunez, A. (1993). The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: Scenario of sleep rhythm generation in interacting thalamic and neocortical networks. J. Neurosci. *13*, 3284–3299.

Steriade, M., Timofeev, I., and Grenier, F. (2001). Natural waking and sleep states: A view from inside neocortical neurons. J. Neurophysiol. *85*, 1969–1985.

Sun, Y., Danila, B., Josić, K., and Bassler, K.E. (2009). Improved community structure detection using a modified fine-tuning strategy. EPL Eur. Lett. *86*, 28004.

Sun, Y.-G., Wu, C.-S., Renger, J.J., Uebele, V.N., Lu, H.-C., and Beierlein, M. (2012). GABAergic synaptic transmission triggers action potentials in thalamic reticular nucleus neurons. J. Neurosci. *32*, 7782–7790.

Thivierge, J.-P., and Cisek, P. (2011). Spiking neurons that keep the rhythm. J. Comput. Neurosci. *30*, 589–605.

Thomas, C.G., and Menon, R.S. (1998). Amplitude response and stimulus presentation frequency response of human primary visual cortex using BOLD EPI at 4 T. Magn. Reson. Med. *40*, 203–209.

Thomson, D.J. (1982). Spectrum estimation and harmonic analysis. Proc. IEEE *70*, 1055–1096.

Triantafyllou, C., Hoge, R.D., Krueger, G., Wiggins, C.J., Potthast, A., Wiggins, G.C., and Wald, L.L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. Neuroimage *26*, 243–250.

Ulrich, D., and Huguenard, J.R. (1996). GABAB receptor-mediated responses in GABAergic projection neurones of rat nucleus reticularis thalami in vitro. J. Physiol. *493*, 845–854.

Veit, J., Bhattacharyya, A., Kretz, R., and Rainer, G. (2011). Neural response dynamics of spiking and local field potential activity depend on CRT monitor refresh rate in the tree shrew primary visual cortex. J. Neurophysiol. *106*, 2303–2313.

Wandell, B.A., Dumoulin, S.O., and Brewer, A.A. (2007). Visual field maps in human cortex. Neuron *56*, 366–383.

Wang, H.-P., Spencer, D., Fellous, J.-M., and Sejnowski, T.J. (2010). Synchrony of thalamocortical inputs maximizes cortical reliability. Science *328*, 106–109.

Wang, L., Saalmann, Y.B., Pinsk, M.A., Arcaro, M.J., and Kastner, S. (2012). Electrophysiological low-frequency coherence and cross-frequency coupling contribute to BOLD connectivity. Neuron *76*, 1010–1020.

Wang, S., Bickford, M.E., van Horn, S.C., Erisir, A., Godwin, D.W., and Sherman, S.M. (2001). Synaptic targets of thalamic reticular nucleus terminals in the visual thalamus of the cat. J. Comp. Neurol. *440*, 321–341.

Wang Jinhui, Zuo Xinian, and He Yong (2010). Graph-based network analysis of resting-state functional MRI. Front. Syst. Neurosci. *4*.

Warren, R.A., Agmon, A., and Jones, E.G. (1994). Oscillatory synaptic interactions between ventroposterior and reticular neurons in mouse thalamus in vitro. J. Neurophysiol. *72*, 1993–2003.

Weese, G.D., Phillips, J.M., and Brown, V.J. (1999). Attentional orienting is impaired by unilateral lesions of the thalamic reticular nucleus in the rat. J. Neurosci. *19*, 10135–10139.

Wei, L., and Craigmile, P.F. (2010). Global and local spectral-based tests for periodicities. Biometrika *97*, 223–230.

Werth, E., Achermann, P., Dijk, D.-J., and Borbély, A.A. (1997). Spindle frequency activity in the sleep EEG: individual differences and topographical distribution. Electroencephalogr. Clin. Neurophysiol. *103*, 535–542.

Wespatat, V., Tennigkeit, F., and Singer, W. (2004). Phase sensitivity of synaptic modifications in oscillating cells of rat visual cortex. J. Neurosci. *24*, 9067–9075.

Williams, P.E., Mechler, F., Gordon, J., Shapley, R., and Hawken, M.J. (2004). Entrainment to video displays in primary visual cortex of macaque and humans. J. Neurosci. *24*, 8278–8288.

Williamson, A.M., Ohara, P.T., Ralston, D.D., Milroy, A.M., and Ralston, H.J. (1994). Analysis of gamma-aminobutyric acidergic synaptic contacts in the thalamic reticular nucleus of the monkey. J. Comp. Neurol. *349*, 182–192.

Womelsdorf, T., Fries, P., Mitra, P.P., and Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. Nature *439*, 733–736.

Womelsdorf, T., Schoffelen, J.-M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. Science *316*, 1609–1612.

Wood, M., Bronskill, M., Mulkern, R., and Santyr, G. (1993). Physical MR desktop data. J. Magn. Reson. Imaging *3*, 19.

Wróbel, A., Ghazaryan, A., Bekisz, M., Bogdan, W., and Kamiński, J. (2007). Two streams of attention-dependent β activity in the striate recipient zone of cat's lateral posterior–pulvinar complex. J. Neurosci. *27*, 2230–2240.

Wunderlich, K., Schneider, K.A., and Kastner, S. (2005). Neural correlates of binocular rivalry in the human lateral geniculate nucleus. Nat. Neurosci. *8*, 1595–1602.

Yu, X.-J., Xu, X.-X., He, S., and He, J. (2009). Change detection by thalamic reticular neurons. Nat. Neurosci. *12*, 1165–1170.

Zhang, L., and Jones, E.G. (2004). Corticothalamic inhibition in the thalamic reticular nucleus. J. Neurophysiol. *91*, 759–766.

Zhang, P., Jamison, K., Engel, S., He, B., and He, S. (2011). Binocular rivalry requires visual attention. Neuron *71*, 362–369.

Zikopoulos, B., and Barbas, H. (2006). Prefrontal projections to the thalamic reticular nucleus form a unique circuit for attentional mechanisms. J. Neurosci. *26*, 7348–7361.

Zikopoulos, B., and Barbas, H. (2012). Pathways for emotions and attention converge on the thalamic reticular nucleus in primates. J. Neurosci. *32*, 5338–5350.

List of Abbreviations

BOLD	Blood oxygen level dependent
Cx36	Connexin36
EEG	Electroencephalography
EPI	Echo-planar imaging
FD	Frame-wise displacement
FDR	False discovery rate
FWHM	Full-width half-max
GABA	γ-Aminobutyric acid
HRF	Hemodynamic response function
IPSP	Inhibitory post-synaptic potential
KS	Kolmogorov–Smirnov
LGN	Lateral geniculate nucleus
LPF	Local field potential
MAD	Mean absolute deviation
MEG	Magnetoencephalography
MGN	Medial geniculate nucleus
MRI	Magnetic resonance imaging
PD	Proton density
ROI	Region of interest
SC	Superior colliculus
TE	Echo time
TR	Repetition time
TRN	Thalamic reticular nucleus
V1	Primary visual cortex
VI	Variation of information

Table of Data Acquired for Each Participant

	S1	**S2**	**S3**	**S4**	**S5**	**S6**
Functional runs	25	30	20	30	23	30
Proton density voxel volume (ml)	0.14	0.56	0.56	0.56	0.56	0.56
Proton density runs	174	40	40	40	40	40
Number of TRs removed	16	31	4	41	11	6
TRs removed / run	0.64	1.03	0.2	1.36	0.48	0.2
Sex	M	M	F	M	M	M
Handed	R	L	R	R	R	R