NeuroImage 209 (2020) 116537

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

The role of transient spectral 'bursts' in functional connectivity: A magnetoencephalography study

Zelekha A. Seedat [a], Andrew J. Quinn [b], Diego Vidaurre [b,e], Lucrezia Liuzzi [g], Lauren E. Gascoyne [a], Benjamin A.E. Hunt [a,c], George C. O'Neill [f], Daisie O. Pakenham [a], Karen J. Mullinger [a,d], Peter G. Morris [a], Mark W. Woolrich [b], Matthew J. Brookes [a,*]

[a] Sir Peter Mansfield Imaging Centre, School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, NG7 2RD, UK
[b] Oxford Centre for Human Brain Activity, Wellcome Centre for Integrative Neuroimaging, Department of Psychiatry, University of Oxford, UK
[c] Diagnostic Imaging, Hospital for Sick Children, 555 University Avenue, Toronto, Ontario, M5G 1X8, Canada
[d] Centre for Human Brain Health, School of Psychology, University of Birmingham, Birmingham, B15 2TT, UK
[e] Department of Clinical Medicine, Palle Juul-Jensens Boulevard 82, Building 2, Incuba/Skejby, 8200 Aarhus N, Denmark
[f] Wellcome Centre for Human Neuroimaging, University College London, 12 Queen Square, London, WC1N 3AR, UK
[g] Mood Brain and Development Unit, Emotion and Development Branch, NIH/NIMH, Bethesda, MD 20892, USA

A B S T R A C T

Neural oscillations dominate electrophysiological measures of macroscopic brain activity and fluctuations in these rhythms offer an insightful window on cortical excitation, inhibition, and connectivity. However, in recent years the 'classical' picture of smoothly varying oscillations has been challenged by the idea that many 'oscillations' may actually be formed from the recurrence of punctate high-amplitude bursts in activity, whose spectral composition intersects the traditionally defined frequency ranges (e.g. alpha/beta band). This finding offers a new interpretation of measurable brain activity, however neither the methodological means to detect bursts, nor their link to other findings (e.g. connectivity) have been settled. Here, we use a new approach to detect bursts in magnetoencephalography (MEG) data. We show that a time-delay embedded Hidden Markov Model (HMM) can be used to delineate single-region bursts which are in agreement with existing techniques. However, unlike existing techniques, the HMM looks for specific spectral patterns in timecourse data. We characterise the distribution of burst duration, frequency of occurrence and amplitude across the cortex in resting state MEG data. During a motor task we show how the movement related beta decrease and post movement beta rebound are driven by changes in burst occurrence. Finally, we show that the beta band functional connectome can be derived using a simple measure of burst overlap, and that coincident bursts in separate regions correspond to a period of heightened coherence. In summary, this paper offers a new methodology for burst identification and connectivity analysis which will be important for future investigations of neural oscillations.

1. Introduction

Neural oscillations comprise rhythmic fluctuations in electrical potential observed across neuronal assemblies. These oscillations, which exist across a range of frequencies from one to several hundred Hertz, dominate electrophysiological measures of large-scale brain activity. A vast body of work suggests that an increase in the amplitude of 'low' frequency (e.g. alpha (8–13 Hz) and beta (13–30 Hz) band) oscillations is a marker of increased functional inhibition whereas increased high frequency (gamma (30+ Hz) band) amplitude is a marker of excitation (Pfurtscheller and Lopes da Silva, 1999). Modulation of neural oscillations during both simple sensory and cognitive tasks are some of the most robust measurements in brain imaging – for example, ballistic finger movements generate a drop in beta (13–30 Hz) amplitude during movement (the movement related beta decrease (MRBD)), followed by a transient increase in amplitude (above baseline) on movement cessation (the post movement beta rebound (PMBR)) (Jurkiewicz et al., 2006). Whilst the functional significance of these effects is unknown, their importance is underlined by several demonstrations of abnormalities across a number of disorders including developmental conditions (e.g. autism (Buard et al., 2018)), severe psychoses (e.g. schizophrenia (Robson et al., 2016)), and neurodegenerative disorders (e.g. Parkinson's disease (Gross et al., 2001)). Some light on the role of beta oscillations has been shed by pharmacological manipulation; for example several studies have shown that alteration of GABA levels results in changes in beta modulation (Muthukumaraswamy et al., 2013). Further studies have shown a direct link between GABA concentration and beta amplitude (Gaetz et al., 2011). Such studies support a hypothesis that beta oscillations are related to inhibition. However, other studies have suggested that beta oscillations are related to long range connectivity

* Corresponding author.
E-mail address: matthew.brookes@nottingham.ac.uk (M.J. Brookes).

https://doi.org/10.1016/j.neuroimage.2020.116537
Received 16 July 2019; Received in revised form 2 December 2019; Accepted 10 January 2020
Available online 11 January 2020

between distal brain regions (Engel et al., 2013), with demonstrations of large-scale spatio-temporal correlations in oscillatory envelopes measured across the cortex (Brookes et al., 2011; Hipp et al., 2012). Precisely how the inhibitory modulation of oscillatory amplitude helps to drive connectivity remains unknown.

Recent work has begun to change the way that the research community thinks about oscillations, particularly in the beta band (Ede et al., 2018). The 'classical' picture is that the brain generates an ongoing oscillation whose amplitude varies over time depending on the task being undertaken. The MRBD was thought of as a drop in oscillatory amplitude, whilst the PMBR represents a smooth increase in amplitude above the resting level. However, this picture was largely a result of time and/or trial averaging. More recent studies (Jones, 2016; Sherman et al., 2016; Shin et al., 2017) looking at unaveraged beta oscillations, both at rest and during tasks, have shown that rather than a smooth oscillation, the beta rhythm is actually formed from the recurrence of discrete and punctate events; each event can be thought of as a very short (e.g. a few hundred milliseconds) burst of activity. These 'beta bursts' occur with a characteristic probability, which is altered by a task. For example during movement execution, the probability of a beta bursts becomes lower; during the PMBR that probability becomes higher (Little et al., 2019). This means that, when summed over large numbers of trials, bursts combine to give the impression of a smooth decrease, followed by an increase in oscillatory amplitude (the MRBD and PMBR). Interestingly, Little et al., (2019) have shown behavioural relevance of bursts by demonstrating that the timing of the last burst, prior to movement, predicts movement onset time. The classical view of event related synchronisation and desynchronization is, therefore, likely inadequate. This, in turn, has significant implications across a range of nascent neuroscientific findings, including the application to a clinical setting (Tinkhauser et al., 2017a; Tinkhauser et al., 2017b), the inhibition hypothesis, and the interpretation of electrophysiological functional connectivity (Engel et al., 2013). This means that much work needs to be done to understand mechanisms and implications. For example, the methodological means to detect single-region bursts is not settled, with most studies choosing an empirically-derived thresholding of beta band limited data. However, it's likely that bursts are not limited to the beta band, and a model driven, broad band approach which, in addition to identifying bursts, can also characterise their spectral content would provide an important step forward.

The role of neural oscillations in functional connectivity is based primarily on the concept of 'communication by coherence' (Fries 2005, 2015). The premise is that if neural oscillations in two separate brain regions are in phase, then this provides periods of mutually high electrical potential, which offer optimal windows for the transfer of action potentials, and hence information. In resting state data, brain regions that are highly connected would therefore be hypothesised to exhibit high coherence, and this has proven to be the case with a number of studies showing phase-locking between regions (Engel et al., 2013; Vidaurre et al., 2018). However, other studies have suggested that temporal correlation between the amplitudes of neural oscillations also offers a means to measure connectivity. E.g. Brookes et al., (2011) and Hipp et al., (2012) independently showed that some resting state networks commonly observed in functional magnetic resonance imaging (fMRI), could also be observed in MEG by correlating the envelopes of beta band oscillations. More recent studies (Colclough et al., 2016; Liuzzi et al., 2017) have shown that whilst network measures made using coherence and amplitude envelope correlation can be similar, in individual subjects amplitude metrics are typically more repeatable. Moreover, there is now growing evidence that within-network functional connectivity is underpinned by coordinated neuronal dynamics that fluctuate on a very rapid (e.g. 100 ms) timescale (Baker et al., 2014). This implies that burst-like activity might be driving network coordination. However, no-one has specifically probed the extent to which bursts drive functional connectivity, and assessment of how these bursts temporally coincide across regions might offer a novel means to measure and interpret functional connectivity.

In this paper, in a subject cohort of resting state and task positive MEG data, we assess the role of bursts in mediating connectivity. In the first part of the paper, we introduce a method to detect single-region transient bursts in source localised MEG data. This method, which is based upon the premise of a Hidden Markov Model (Baker et al., 2014; Vidaurre et al., 2016; Woolrich et al., 2013; Quinn et al., 2018) offers an unsupervised and objective means to identify bursts in broadband (univariate) data. In order to ensure that we capture the full spectral profile of these bursts, we do not constrain our model to the beta band. Instead we look to detect transient spectral bursts within a broad frequency range (1–48Hz) and compare findings to a 'classical' beta burst framework. Following this, we assess the role of these bursts in functional connectivity. Specifically, we test a hypothesis that, and the extent to which, the beta band electrophysiological connectome can be derived based upon an analysis of coincident bursts.

2. Materials and methods

Data were acquired as part of the United Kingdom MEG partnership programme (see Acknowledgements) and have been published previously (Hunt et al., 2018). 75 subjects took part in the study, which was approved by the University of Nottingham Medical School Research Ethics Committee.

2.1. Paradigms and data acquisition

All MEG data were acquired using a 275-channel CTF MEG system, at a sampling rate of 1200 Hz. The system was operated inside a three layer magnetically shielded room (MSR) and in third order synthetic gradiometer formulation to reduce the effects of external interference. All subjects were seated. Data were acquired during two separate paradigms:

Resting state: Five minutes of resting state MEG data were acquired; the subject was asked to sit with their eyes-open, and 'think of nothing.' A fixation cross was displayed (by projection through a waveguide in the MSR) on a back projection screen which was placed approximately 40 cm in front of the subject in order to give them something to focus on during data acquisition.

Visuo-motor task: The task comprised presentation of a visual stimulus – a high contrast vertical square wave grating – for a jittered interval of 1.5 s – 2 s. Upon cessation of visual stimulation, the participant was asked to make a single right index finger abduction. Fifty trials employed a short inter trial interval (ITI) of 4 s, and a further 50 trials employed a longer ITI of 8 s. As the PMBR has been shown to last in excess of 6 s (Fry et al., 2016; Pfurtscheller and Lopes da Silva, 1999; Pakenham et al., 2019) only the long ITI trials are analysed in the current work.

All subjects also underwent an anatomical MRI scan, using a Philips 7T Achieva MRI scanner (Philips) running a phase sensitive inversion recovery (PSIR) sequence (field of view [FOV]: 240 × 216 × 160 mm^3, 0.8 mm isotropic resolution). In order to co-register MEG functional data to MRI structural data, we employed 3D head digitisation. Prior to MEG acquisition, three coils were placed at fiducial locations on the head (the nasion and pre-auricular points). The locations of these coils, relative to the subjects scalp and face surface were digitised (Polhemus Inc.). The equivalent head surface was extracted from the anatomical MRI scan and a surface matching algorithm was employed to compute the locations of the fiducial markers relative to the MRI (hence the brain anatomy). The coils were subsequently energised during MEG recording in order to localise their position inside the MEG helmet; this was done continuously throughout data acquisition. Knowledge of the location of the fiducial markers inside the MEG helmet enabled complete coregistration of the MEG sensor geometry to brain anatomy. It also allowed motion tracking of the subject's head.

Following acquisition, data were inspected visually; 9 subjects were removed from the resting state paradigm and 12 were removed from the visuo-motor paradigm either due to movement (>5 mm) or artifacts in the data. This left a total of 66 (age 38 ± 12; 35 female) and 63 (age 38 ±

12; 34 female) subjects for the resting and visuomotor paradigms respectively.

2.2. Source localisation

The cortex was parcellated into 78 regions according to the automated anatomical labelling (AAL) atlas (Tzourio-Mazoyer et al., 2002). Each region was divided into voxels (4 mm cubic grid) and a beamformer (Robinson and Vrba, 1998) was employed to generate a time-course representing electrophysiological activity at each voxel. A scalar linearly constrained minimum variance (LCMV) beamformer was used with covariance computed within a 1–150Hz frequency window and a time window spanning the whole experiment (Brookes et al., 2008). Regularisation was applied to the covariance matrix using the Tikhonov method with a regularisation parameter equal to 5% of the maximum eigenvalue of the unregularised covariance matrix. The forward model was based upon a dipole approximation (Sarvas, 1987) and a multiple local sphere head model (Huang et al., 1999). Dipole orientation was determined by rotating the dipole in the tangential plane and finding the orientation which gave the highest SNR. Following this, data from each voxel were weighted according to a Gaussian function of the distance of that voxel from the centre of mass of the AAL region (17 mm full width at half maximum) and summed. This process, which has been used previously (Brookes et al., 2018), resulted in a single regional timecourse of activity for each of the 78 parcels in the AAL atlas. These timecourses were then frequency filtered between 1 and 48 Hz, and temporally down-sampled to a sampling rate of 100 Hz. Symmetric orthogonalisation was used to reduce the effect of spatial leakage (Colclough et al., 2015).

2.3. Hidden Markov Model

In order to find single-region beta bursts, we used an HMM. Each regional timecourse was treated independently and a univariate 3-state Time-Delay Embedded (TDE) HMM (Vidaurre et al., 2018) was inferred. An HMM assumes that a series of mutually exclusive 'hidden' states governs the observed region's timecourse, such that each timepoint is associated with one of the hidden states. The sequence is assumed to be Markovian so that the state modelled as active at time point t is conditionally independent of previous time points, given time point t-1; that is, time point t only depends on time point t-1. An observation model links the HMM state to the observed value in the regional timecourse. In its simplest form, the model would describe each state by a different Gaussian distribution, from which the observed values can be extracted. The mean and standard deviation of each Gaussian would then define each state. This is shown schematically in Fig. 1, where the values y_t correspond to the measured data which are drawn from a distribution defining each state. The sequence of states is described by x_t. The number of states is defined a-priori (in the schematic in Fig. 1, there are 2 states) and model inference would learn the observation models, and sequence of states, from the observed data.

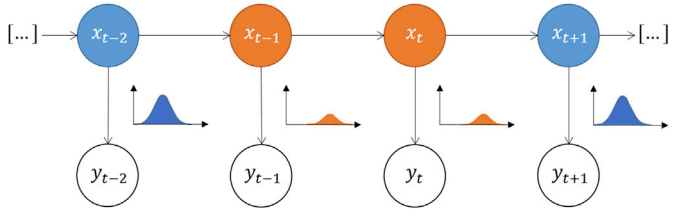

Fig. 1. Schematic diagram of a simple HMM. At each time point, t, there is an observation, y_t (the source localised MEG data) and an underlying (hidden) state, x_t. The model here assumes that each observation is drawn from a Gaussian distribution where each state corresponds to a different Gaussian (i.e. a different mean and standard deviation). The number of states (hence Gaussians) is set a-priori. In this schematic example we show just two states, in orange and blue.

Here we used the HMM with time-delay embedding (Vidaurre et al., 2018) where each state is characterised by a different autocovariance pattern defined over a specified time window (duration 230 ms). These state autocovariance patterns contain the spectral information of the signal when that state is active.

The model inference was undertaken using a variational Bayesian method which seeks to minimise the free energy of the system. For each separate regional timecourse, we assumed 3 states (though if insufficient evidence for a third state existed, the model was able to collapse to 2 states). Although the rest of this paper relates to our 3-state model, results from a 6-state and a 10-state model can be found in the supplementary information. The output of the model was therefore a set of 3 state timecourses, each representing the probability of a region's timecourse existing in a specified state at any time point. We then measured the correlation between these state probability timecourses and the amplitude envelope of beta oscillations in the same brain region (estimated by applying a (Morlet) wavelet transform to the regional timecourse and extracting those values corresponding to the 13–30Hz frequency band to derive the instantaneous envelope). The state whose probability timecourse correlated highest with the beta amplitude envelope was taken to represent the 'transient spectral burst state'. This burst probability timecourse (which reflects the inherent uncertainty in the model) was then binarised; where the *instantaneous* probability of being in the burst state was greater than two thirds, we assumed the state had been entered and that a burst had begun. (In post-hoc analyses we also tested a threshold probability of ½.) This is an improvement over a traditional, heuristic approach (thresholding the beta envelope) in two ways. First, the burst identification is not based purely on the amplitude timecourse of an a-priori specified frequency band, but on a generative model of bursting across a broader frequency range. This allows us to examine the full spectral profile of bursts without assuming that the only interesting part of their spectrum falls within the beta band. Second, the thresholding is more principled, being based on the actual probability of being in a burst state.

2.4. Burst parameters and comparison to established methods

Having applied the HMM, we sought to examine the characteristics of the state visits (i.e. the bursts), and how they vary across brain regions. To this end we used the binary timecourses associated with the burst state, alongside the regional timecourse data, to determine four burst 'features':

- **Burst duration:** The time that a region spent in the burst state, on each visit.
- **Burst amplitude:** The maximum value of the beta envelope during each visit to the burst state.
- **Frequency of occurrence:** The number of visits to the burst state normalised by time.
- **Burst interval:** The length of time between visits to the burst state.

Each of these values was measured independently for each region and subject, and the values were averaged over subjects and plotted as a function of cortical location (i.e. AAL region). For resting state data, these values were recorded as a time average over the entire 5 min resting state recording. For the visuo-motor task data, they were measured in three windows; during the movement related beta desynchronization, the post movement beta rebound, and a 'rest' period at the end of each trial. These windows were taken, approximately, to be 0–1 s (MRBD), 1–3 s (PMBR) and 4–7 s (rest) relative to the offset of the grating, however they were also allowed to modulate, independently for each trial, according to when the bursts actually occurred. I.e. when a burst fell at the edge of a time-window, the window was extended to include that burst. Similarly we extended the MRBD window to capture the time between the edge of the window and the occurrence of a subsequent (or previous) burst. These burst-modulated time windows are shown alongside our results for

completeness. For each subject, this procedure resulted in 16 measurements (4 burst parameters, each measured in the resting state and in 3 separate task windows). We used a Wilcoxon non-parametric rank sum test in order to statistically test for a difference in burst parameters between the resting state and the task. To correct for multiple comparisons, the Bonferroni correction was applied: because three tests were computed (one for each window - MRBD, PMBR and rest) for each of the four burst parameters, the significance threshold was divided by twelve and reduced from 0.05 to a corrected value of 0.0042.

We sought to compare our HMM derived broad-band (1–48 Hz) bursts, to a more conventional beta band thresholding approach, in order to test the extent to which the HMM was identifying the same bursts as the established method. To this end, using resting state data, we took the mean corrected beta (13–30 Hz) envelope of a single region (left sensory cortex) and applied three different thresholds (1.5, 2.5 and 3 times the standard deviation (measured independently for each subject)) in order to get a binary timecourse showing periods of high beta amplitude (the thresholded beta bursts). We then quantified the percentage of overlapping bursts picked up by both the HMM and thresholding techniques; the percentage picked up by the HMM only, and the percentage picked up by thresholding only. These values were measured independently for each subject and averaged over subjects. The broad-band nature of our HMM implementation also enabled us to characterise the spectral properties of these bursts in both the resting state and visuo-motor task data. This was done using a state-specific multitaper analysis (Vidaurre et al., 2016), resulting in a single spectrum for each state and each subject, in each location in the brain. Spectra were averaged and plotted alongside the standard error over subjects.

2.5. Measuring functional connectivity

In order to measure functional connectivity, we assessed relationships between all possible pairs of regional timecourses from the 78 AAL regions; this method results in 3003 measures of inter-regional functional connectivity. We quantified functional connectivity in two ways:

1. **Amplitude envelope connectivity (AEC):** AEC (Brookes et al., 2011; de Pasquale et al., 2010; Hipp et al., 2012) is a well characterised connectivity measure in which the Pearson correlation between oscillatory envelopes is measured. Here, we calculated AEC in the theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), low gamma (30–48 Hz) and high gamma (48–70 Hz) bands. Following frequency filtering, the envelope was calculated using a Hilbert transform and temporally down sampled to 1 s. AEC was measured between all pairs of brain regions to yield a (78 x 78 element) matrix. These were generated independently for each subject.
2. **Burst coincidence:** In accordance with our hypothesis we wanted to test whether AEC could be reproduced by simply considering the likelihood of transient spectral bursts co-occurring in spatially separate brain locations. In order to calculate this, for any one pair of AAL regions, we took the binary timecourses (i.e. the outputs of the HMM) for the burst state in each region and calculated the Jaccard index. This index measures the intersection over union (and hence the similarity) between two binary timecourses; larger values indicate greater burst coincidence. This was computed for all possible region pairs and averaged across subjects. Jaccard indices were then transformed to pseudo-z-statistics by subtracting the mean and dividing by the standard deviation (across the whole matrix).

It should be pointed out that AEC and burst coincidence are, to a degree, related. Specifically, if we assume bursts are periods of high amplitude beta oscillations then the burst timecourses should mirror, to an extent, the amplitude envelope. However, a significant amount of information from the envelope itself has been removed by generating the binary timecourses. Therefore, we might assume that if the burst coincidence and the AEC derived connectomes are in strong agreement, then AEC is mainly driven by coincident bursts as distinct from lower amplitude coherent fluctuations.

To assess the relationship between AEC and burst coincidence, we generated an element-by-element scatter plot and measured linear correlation between the burst coincidence matrix (which recall is measured in the broad band (1–48Hz)) and the AEC matrices, derived in the theta, alpha, beta, low and high gamma bands. To assess how these correlation values compared to a null hypothesis we used a pseudo-matrix approach (Tewarie et al., 2016; Hunt et al., 2016). We first performed an eigenvalue decomposition of the burst coincidence (Jaccard) matrix. Each eigenvector was then randomised using a phase based technique. Post-randomisation reconstruction yielded a pseudo-matrix, similar in structure to the genuine matrix, but not reflecting genuine functional connectivity. Comparison of correlation between real, and pseudo-matrices revealed the statistical significance of the relationship. This was done using a non-parametric Wilcoxon rank-sum test between real and pseudo-correlation values obtained for each subject.

In order to further investigate the contribution of bursts to AEC, we recalculated the AEC matrix (in the beta band) having removed the periods of time during which the coincident bursts occurred – we expected that if bursts were driving the AEC measure, then AEC values would be significantly diminished when the bursts were removed. We quantified the amount by which the AEC values had changed (average across the whole matrix) and also the correlation with the coincident burst matrix.

Finally, we sought to test whether coincident bursts were also representative of time windows of high coherence (i.e. phase-locking) between regions. To this end, we selected 6 exemplar (interhemispheric) connections across the left and right precentral, postcentral, occipital, parietal and frontal cortices (see Fig. 6 for precise locations). We denote these the 'seed' and 'test' regions. For these 6 connections, we found all of the coincident bursts in both regions and, within a time window centred on the midpoint of each of the overlapping bursts, we measured the phase difference derivative (PDD) between regions (Breakspear et al., 2004; Tewarie et al., 2019). PDD is a measure that captures the stability of phase relationships between timecourses: first the instantaneous phase of the seed and test signals is computed and the difference between them calculated. This phase difference is then differentiated. Values close to zero indicate a constant phase difference (hence high coherence) and high values indicate a rapidly changing phase difference (hence low coherence). These values were negatively transformed (using the equation $x_{new} = e^{-|x|}$) to ensure that a high number relates to high coherence. The result is a timecourse of coherence between regions. Here, we measured PDD, in the −0.5 s to 0.5 s time window (relative to burst centre) for all coincident bursts, and averaged the results (yielding an average timecourse reflecting changing coherence throughout a coincident burst). Note that we used PDD, as distinct from simpler (and more direct) measures of phase locking (e.g. coherence) because it yields a timecourse measure throughout the burst duration. In order to contrast this with a control condition, we also looked at 'non-coincident' bursts. We applied the same measure, in the same time window, surrounding an equivalent number of bursts that occurred in the seed region, but without a coincident burst in the test region (i.e. non-coincident bursts). We expected to see higher coherence for genuinely coincident bursts compared to non-coincident bursts. We generated these measurements for all subjects, and tested our hypothesis using a non-parametric Wilcoxon rank-sum test.

3. Results

3.1. Identification of bursts in resting state data using the Hidden Markov Model

Fig. 2 shows, qualitatively, a comparison between a simple thresholding technique (Fig. 2A), and the HMM (Fig. 2B) for identifying bursts. The probability timecourse (blue), and the binarised state timecourse (red), for the burst state, are shown (overlaid) in Fig. 2C. Here we show a

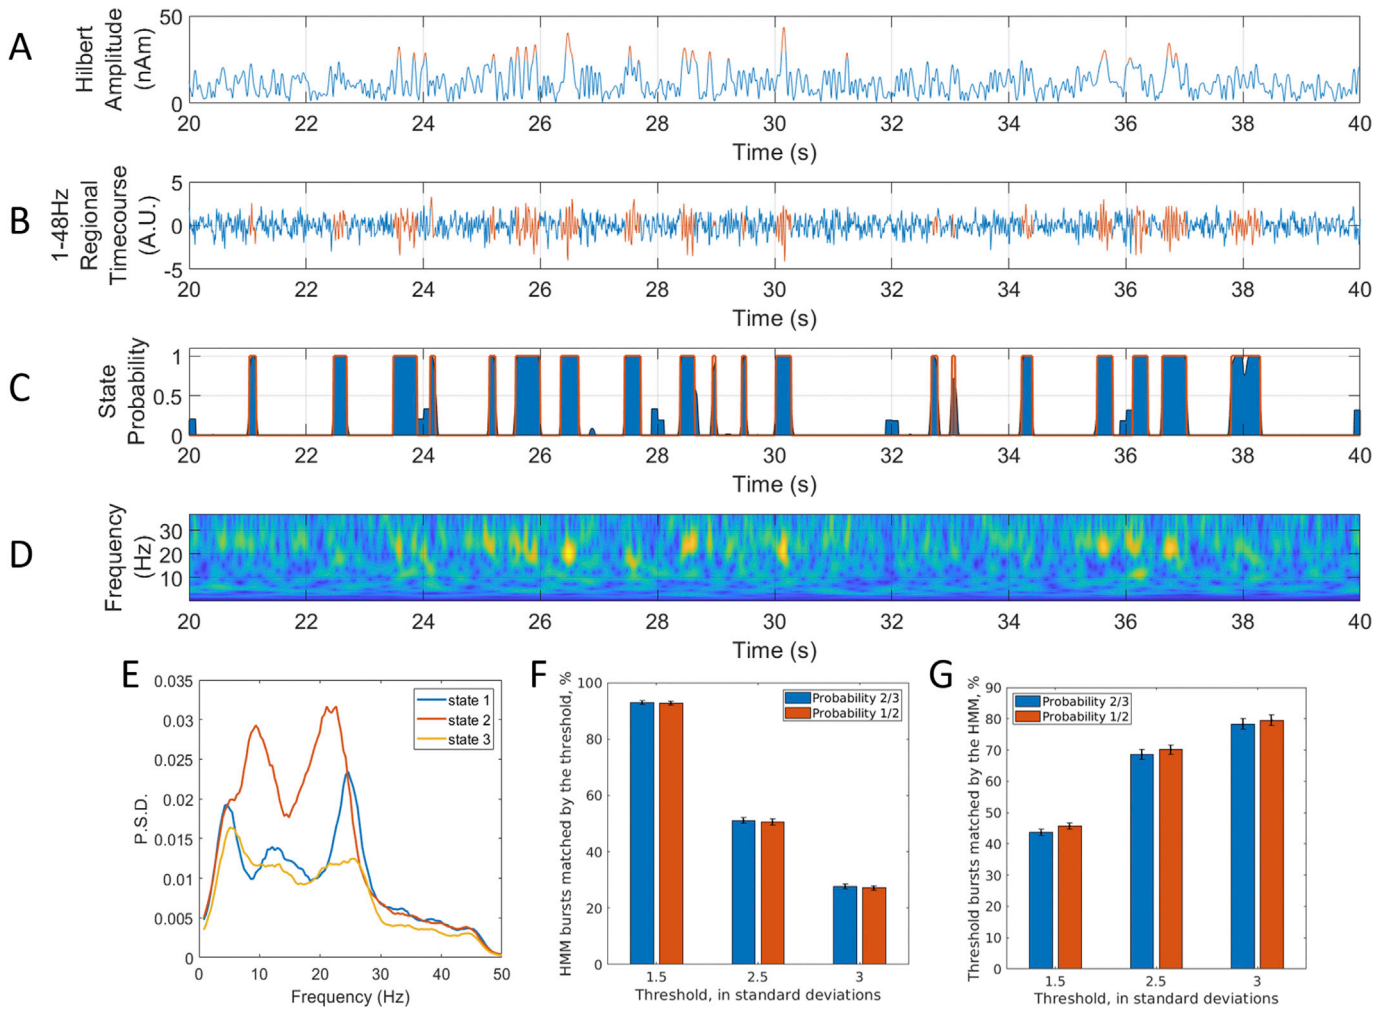

Fig. 2. Burst detection using the Time-Delay Embedded HMM in single region resting state data. A) Beta amplitude with bursts (in red) identified using a simple threshold technique (2.5 times the standard deviation). B) Broadband (1–48 Hz) data with bursts (again in red) identified using the HMM. C) Probability timecourse for the burst state (blue) with binarisation overlaid (red). D) Time frequency decomposition of the data, generated using a wavelet transform. Note that A-D show 20 s of data from a single representative subject. E) Spectra, showing the component frequencies of the three states identified by the HMM from the same single subject. The state most closely correlated with the beta envelope (and hence identified as the beta burst state) is state 2 (red). F) The percentage of HMM identified bursts that were matched by the threshold technique, plotted as a function of threshold itself. Two different probability thresholds on the HMM (two thirds and a half) are shown in blue and red respectively. G) The percentage of bursts identified by the threshold technique and matched by the HMM. The errorbars in F) and G) show the standard error over participants.

snapshot of 20 s of data extracted from the left somatosensory cortex in a single subject. Fig. 2D shows a time-frequency spectrogram (derived from a continuous (Morlet) wavelet transform) of the same data segment, for comparison. Note that despite being applied to broad-band data, the HMM identifies bursts of high amplitude oscillations which are plainly visible in the time-frequency decomposition and also identified using the thresholding technique. However rather than simply using the amplitude envelope, the HMM makes use of the richer information contained in the waveform shape in order to identify the burst. Fig. 2E shows the spectrum of oscillatory frequencies associated with each of the three identified states for this single subject. The burst state was identified as state 2, which in this case peaks in the beta band, but is accompanied by a large alpha band component.

The HMM and thresholding techniques are further compared, across all 66 subjects, in Fig. 2F and 2G. Fig. 2F shows the percentage of HMM bursts which were matched by thresholded bursts, for three different values of threshold. As would be expected, this percentage drops with increasing threshold, but for a low threshold (1.5 times the standard deviation of the envelope) the vast majority (>90%) of HMM-identified bursts were accompanied by a period of above threshold beta amplitude, suggesting that the HMM is a viable way to find punctate periods of high beta envelope. Similarly, Fig. 2G shows the percentage of threshold bursts, which were matched by the HMM. Here, again as would be expected, the percentage increases with threshold, and we see that for a high threshold (3 times the standard deviation) >70% of threshold bursts are also found by the HMM. These statistics suggest that there is agreement between the two methods. However as noted above, the HMM identification of bursts is based on a model of richer spectral content across a broader frequency range, rather than amplitude in a predefined frequency band; and further, the thresholding in the HMM is more principled, by being based on the actual probability of being in a burst state (which has a marked tendency to have values either very close to 0.0 or very close to 1.0). This is demonstrated by a comparison between two different probability threshold values as shown in Fig. 2F and 2G. Changing this probability value from two thirds to a half makes no significant impact.

Fig. 3 shows the properties (duration, amplitude, frequency of occurrence, and interval) of the HMM identified bursts as a function of cortical location, in the resting state. The upper (brain) plots show the mean burst parameters (across subjects). The lower (line) plots show the

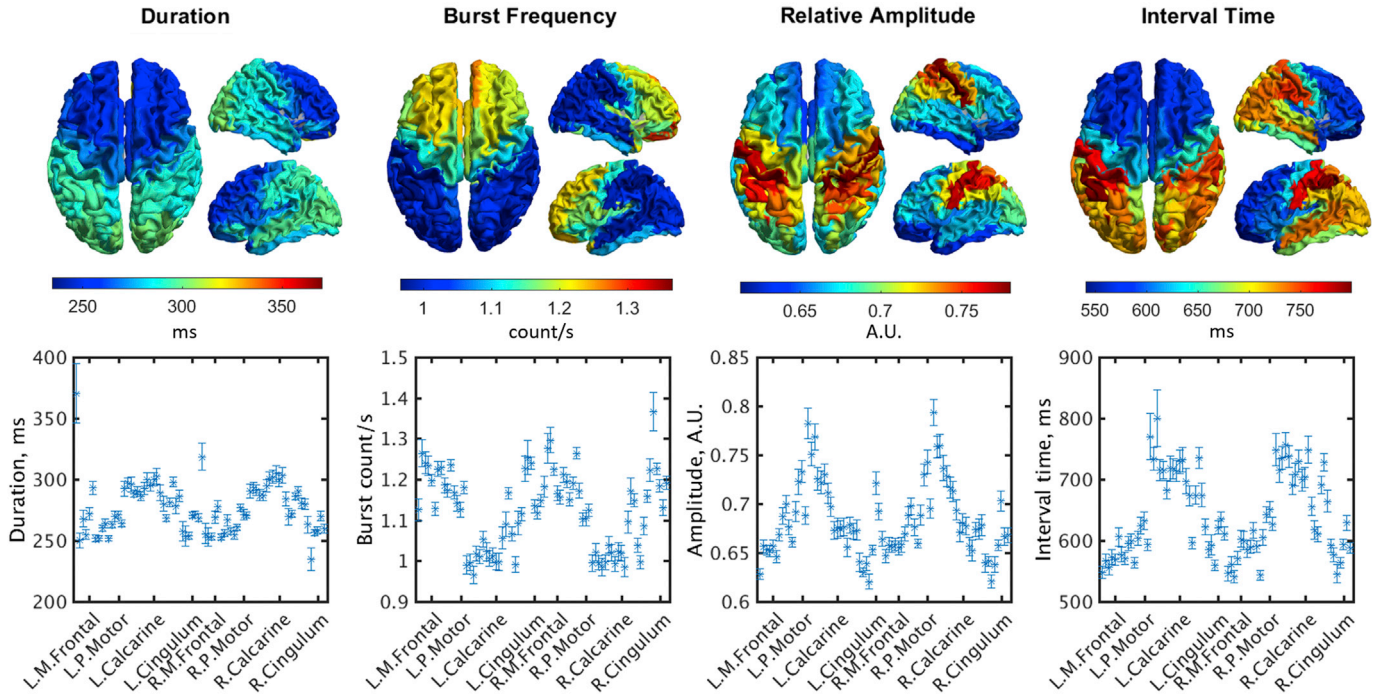

Fig. 3. Resting state burst statistics. Averaged burst duration, frequency of occurrence, amplitude and interval time across the cortex. The brain plots show the mean values as a function of AAL region. The line plots show the same information, but the error bars allow visualisation of standard error in each measurement across subjects.

same thing, but include error bars to enable visualisation of the variation in burst parameters across subjects (note that in general, the variation across subjects is small compared to the regional variation). Generally, we observe that the parietal, occipital and temporal regions generate fewer bursts of longer duration compared to the frontal regions, which tend to generate more bursts of shorter duration. Consistent with the known role of beta in the sensorimotor system, the highest amplitude bursts appear in the primary sensorimotor cortices with the lowest amplitudes in the frontal regions. Interestingly, the time between bursts is also longest in the sensorimotor cortices. On average, the burst durations were found to be of the order 300 ms. However, this is also a function of model parameters (see supplementary information).

3.2. Burst dynamics during a visuomotor task

Fig. 4 shows results of the HMM burst analysis in our visuomotor task data. Fig. 4A shows the difference in burst frequency between a window encapsulating the beta rebound (1s to 3s relative to visual stimulus offset (modulated, on a trial by trial basis, by neighbouring bursts)), and a window capturing the MRBD (0s to 1s (also modulated by neighbouring bursts)). The spatial signature is clear with a peak centred on the contralateral primary sensorimotor area as would be expected. Plots 4B to 4E relate to this peak region. Fig. 4B shows a raster plot of the binary timecourse for the burst state, extracted from the left primary sensory cortex. The x-axis represents time, and the y-axis represents trials. All subjects have been concatenated (vertically) and time zero corresponds to the offset of visual stimulation. As shown, the burst state displays a greater likelihood of occurrence during the PMBR window compared to rest, whilst in the MRBD window, the burst state is less likely to occur. These relationships are formalised in Fig. 4C which show the burst parameters in the three windows associated with the task, as well as in the resting state data. Note here that the distribution represents all bursts in all subjects; the black line shows the mean value, and the red line denotes the median. In agreement with the raster plot, and previous results (Little et al., 2019) we see a marked drop in burst frequency during the MRBD compared to PMBR (or indeed rest and resting state). As would be expected, decreased burst count is accompanied by decreased duration and increased interval between bursts during the MRBD. Interestingly, there is relatively little change in the burst amplitude in the three different windows of the task. There are also substantial differences between the bursts occurring in resting state and the bursts occurring in the visuomotor task data. The bursts detected in the rest window of the visuomotor task are significantly longer than bursts occurring in the resting state; there are also fewer of them, with a shorter interval time. Fig. 4D shows the average frequency spectrum of the burst state across subjects in both the resting state and task data – note the pan-spectral nature of these bursts. Note also the difference between the two spectra; the spectral content which characterises the bursts appears to change with task. Fig. 4E shows the burst-modulated time windows for all trials.

3.3. Transient spectral bursts and their role in functional connectivity

In Fig. 5 we compare burst coincidence with AEC as a means of investigating the whole brain resting state functional connectome. Fig. 5A shows the result of our AEC analysis, applied to beta band data. The matrix shows all-to-all functional connectivity (averaged across subjects), whilst the brain shows the strongest 20% of connections; the colour and width of the lines represent the magnitude of individual connections (i.e. the envelope correlation values between regions) whilst the radius of the grey circles represents connectivity strength between that region and the rest of the brain. As would be expected from previous studies (Hunt et al., 2016), connectivity measured in the beta band is strongest in the sensorimotor, posterior parietal and visual areas. Fig. 5B shows the results of our coincident burst analysis. The elements of the matrix represent a measure of *burst coincidence* (estimated using the (pseudo-z-transformed) Jaccard index; i.e. the intersection over the union (and hence the similarity) between two binary burst timecourses) between AAL regions. The brain plot shows the top 17% of inter-regional coincidence values. There is a clear similarity between the beta AEC and broadband coincident burst results and this is shown quantitatively in

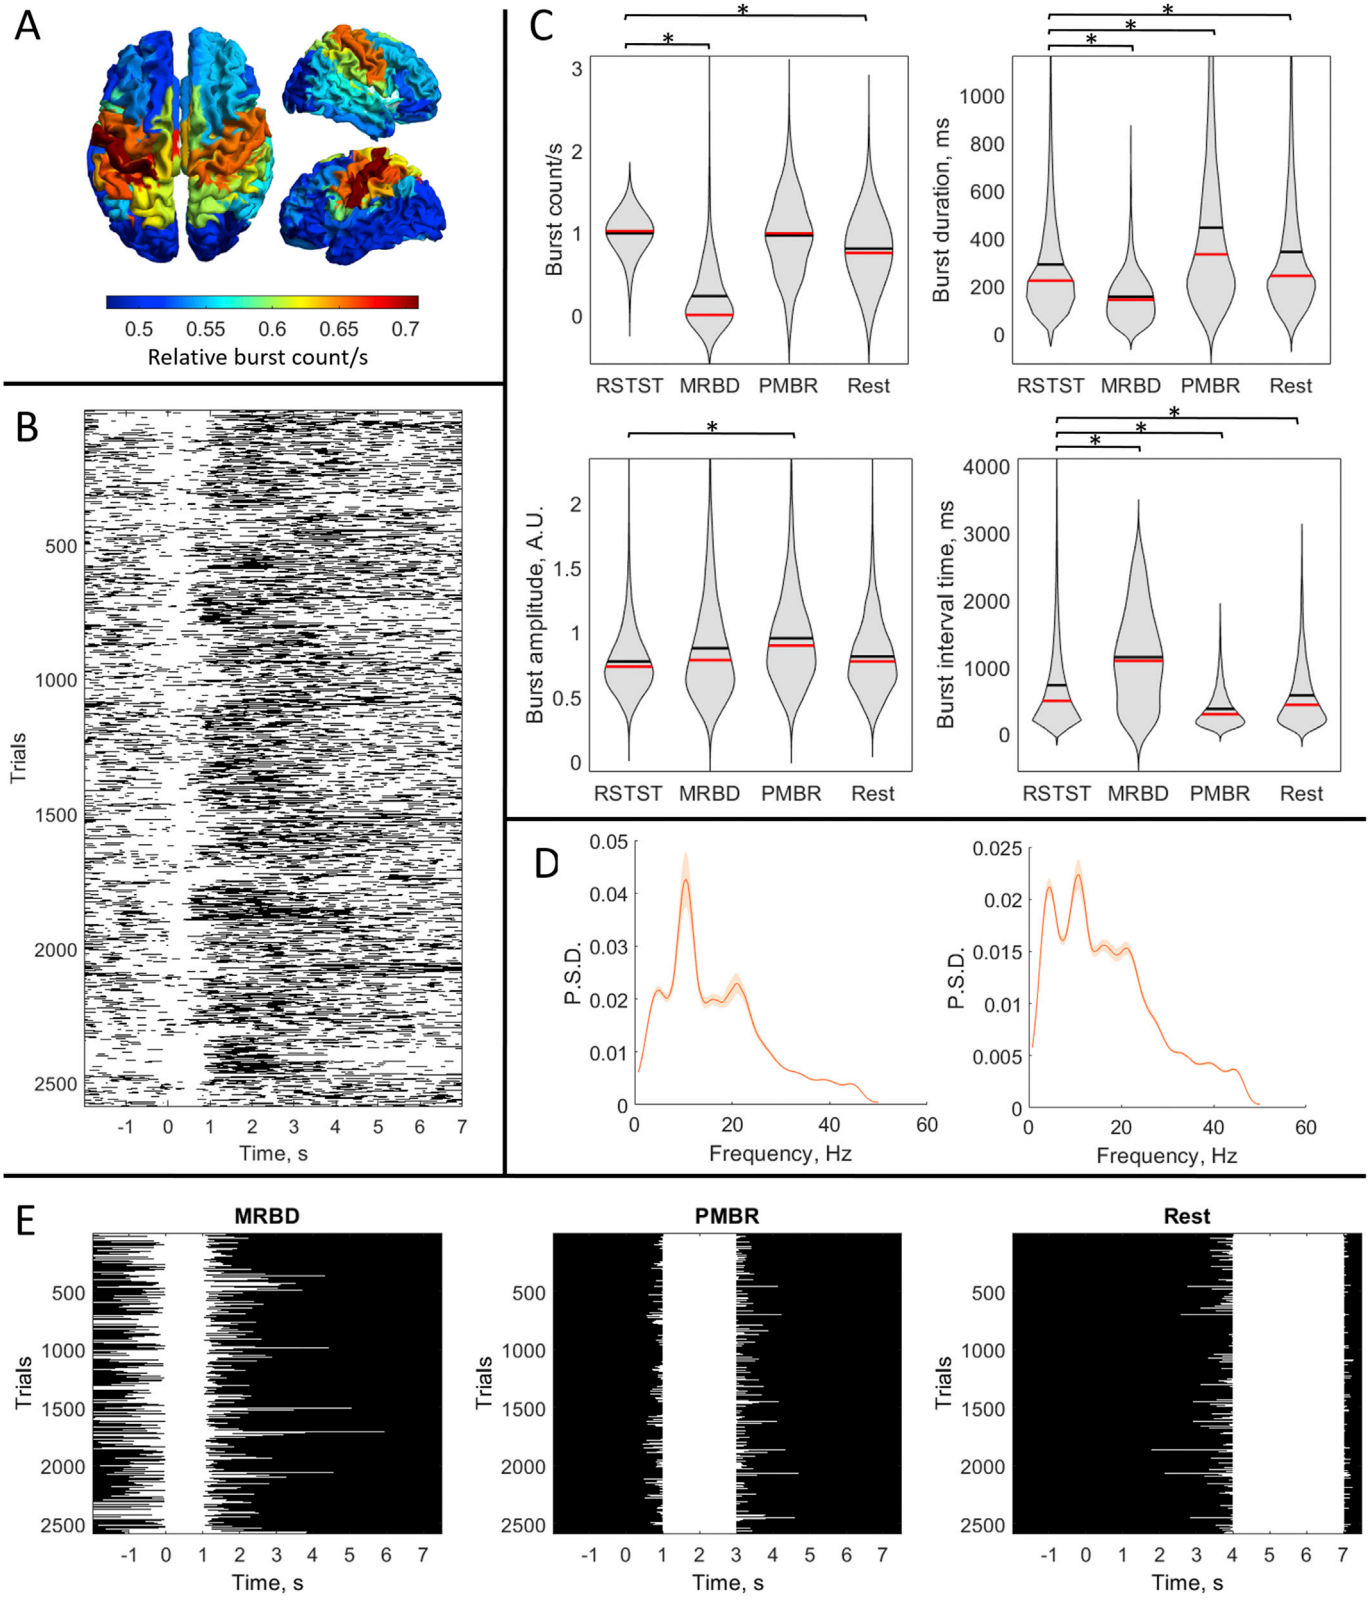

Fig. 4. HMM identified bursts during a visuomotor task. A) Image showing the difference in the frequency of burst occurrence between the PMBR and MRBD windows. Notice that, as expected, burst frequency increases most in left sensorimotor cortex. B) Raster-plot showing the occurrences of the burst state for all trials in all subjects. C) Violin plots showing differences in burst parameters (duration, frequency, amplitude and interval time) in the three separate task windows, and in resting state (RSTST) data. The stars above the violin plots indicate a significant difference (with a corrected p-value$<$0.05) in burst parameters when compared with resting state. D) mean spectrum of the burst state, averaged over subjects, for both resting state (left) and task data (right). The shaded area indicates the standard error over subjects. E) Burst-modulated time windows for the 3 task conditions, extended to include all bursts which fall on the edge of the specified windows. These windows are shown in white and were used in the assessment of burst parameters. Panels B, C, D and E relate to the left sensory cortex.

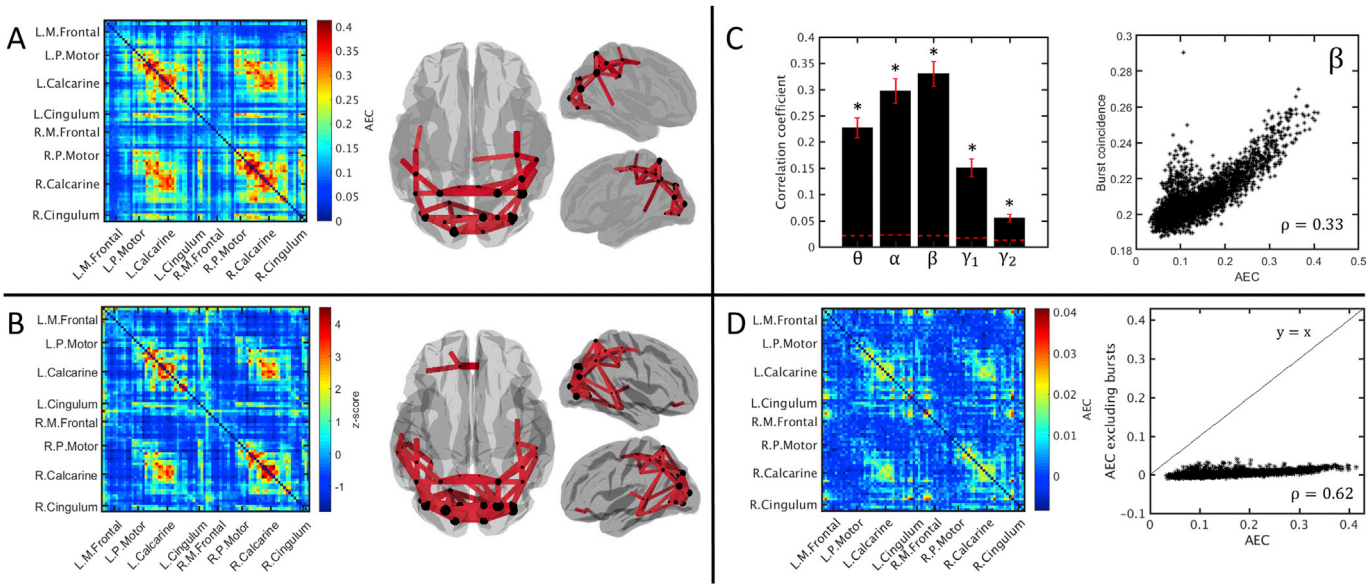

Fig. 5. Comparison of AEC and coincident bursts in resting state functional connectivity. A) Beta band AEC averaged over 66 subjects. The matrix shows all-to-all AEC measurements whereas the brain plot shows the 20% of strongest connections. The thickness and colour of the lines denotes the strength of connectivity and the circles denote connectivity strength between that region and the rest of the brain (i.e. the sum in one direction of the matrix). B) Burst coincidence (estimated using the Jaccard index in the broad (1–48 Hz) band, and transformed to a pseudo-z-statistic). Again, the matrix shows all-to-all connections and the brain plot shows the highest 17% of connections measured. Note similarity between (A) and (B). C) Correlation between the connectome measured using coincident bursts, and AEC measured in all frequency bands. The bar chart shows average values and the error bars show standard error over subjects. The dotted red lines indicate the 'null correlation' values that would be obtained if the structure of the coincident burst matrix were in fact random. The stars indicate where there is a significant difference ($p < 0.05$) between the 'null correlation' and actual correlation values. The scatter plot shows the correlation for the beta band case. D) The matrix shows AEC calculated in beta band data with all bursts removed. Note the different colour scale. The scatter plot shows elements of the matrix in (A) plotted against elements of the matrix in (D).

Fig. 5C. The bar plot shows correlation between connectivity methods (burst coincidence and AEC) for all frequency bands (error bar shows standard error across subjects). An example of this correlation for the averaged matrix is also shown. The dotted line shows correlation between the burst matrix and pseudo-matrices derived from eigenvector phase randomisation; all bands demonstrated correlation beyond chance. The result shows clearly that similar information can be derived from AEC and coincident bursts, with the strongest similarity in the beta band. Fig. 5D shows AEC calculated in envelope data with bursts removed. Note that the colour scale has reduced by a factor of ten when compared with the original AEC matrix (Fig. 5A). The left panel of Fig. 5D shows the beta band AEC matrix and the scatter plot on the right shows AEC values, calculated with and without bursts removed, plotted against each other. Note that whilst some structure remains, the amplitude of the connectivity values is markedly diminished compared to Fig. 5A, highlighting the fact that much of the connectivity information is contained within the bursts. Quantitatively, the AEC values were reduced from 0.146 ± 0.009 to 0.0036 ± 0.0004 as an average over the whole matrix – a reduction of 97.5%. The correlation with burst coincidence was altered from 0.33 ± 0.02 to 0.076 ± 0.009. Given that burst coincidence makes up a fraction of the data (e.g. 10.52% of the timecourse in sensorimotor cortex at resting state) but explains more than 90% of the correlation in the AEC matrix, it is clear that bursts play an important role in driving functional connectivity.

Finally, Fig. 6 shows resting state PDD, measured as a function of time in a -0.5 s to $+0.5$ s time window relative to the centre of coincident (blue) and non-coincident (red) bursts. The figure shows 6 example connections; left-right motor cortex (A), left-right sensory cortex (B), left-right superior parietal cortex (C), left-right visual cortex (D), left-right inferior frontal cortex (E) and left-right orbitofrontal cortex (F). As shown, in the motor, sensory, visual and parietal cortices a coincident burst reveals an increase in broadband coherence between regions, which is not mirrored by non-coincident bursts, with a significant difference ($p < 0.05$) measured in the visual and parietal cortices. This effect was not seen in the frontal cortex. These results support the hypothesis that transient spectral bursts, coincident between regions, offer a means by which to support network connectivity via a brief period of coherent oscillatory activity.

4. Discussion

Recent work has shown that neural "oscillations" are generated, at least in part, via a summation of a set of punctate high amplitude "bursts" whose fundamental frequencies intersect with the traditional bands (e.g. beta). This new finding has paved the way for a new mechanistic interpretation of the role played by "oscillations" in cognition. However, this work is in its infancy, and even the methodological means to detect bursts is not settled. To date, most studies have employed a simple thresholding of band limited data, but this treats individual bands in isolation without reference to a broader pan-spectral picture of bursts. In addition, there has been no consensus reached on the threshold used. For example recent studies have used: 2 standard deviations on the mean (Lundqvist et al., 2016); 6 times the median (Shin et al., 2017); and the 98th percentile (Sherman et al., 2016). Because of this, a more principled and broad band approach to detect bursts would be useful.

In previous work, the HMM has been used on multi-region (i.e. multivariate) brain data to reveal visits to short-lived transient brain states with spatially and spectrally distinct patterns (Quinn et al., 2018; Vidaurre et al., 2018). Here, we showed that the HMM also offers a useful means to detect spectral bursts in single-region (i.e. univariate) data, as an alternative to the simple thresholding approach (Shin et al., 2017). Importantly, our use of a time-delay embedded HMM, where state characteristics are based upon data autocovariance, means that bursts are identified based on the temporal morphology of the signal over a broader frequency range. This is distinct from a simple thresholding technique where only the amplitude of the signal within a prespecified frequency

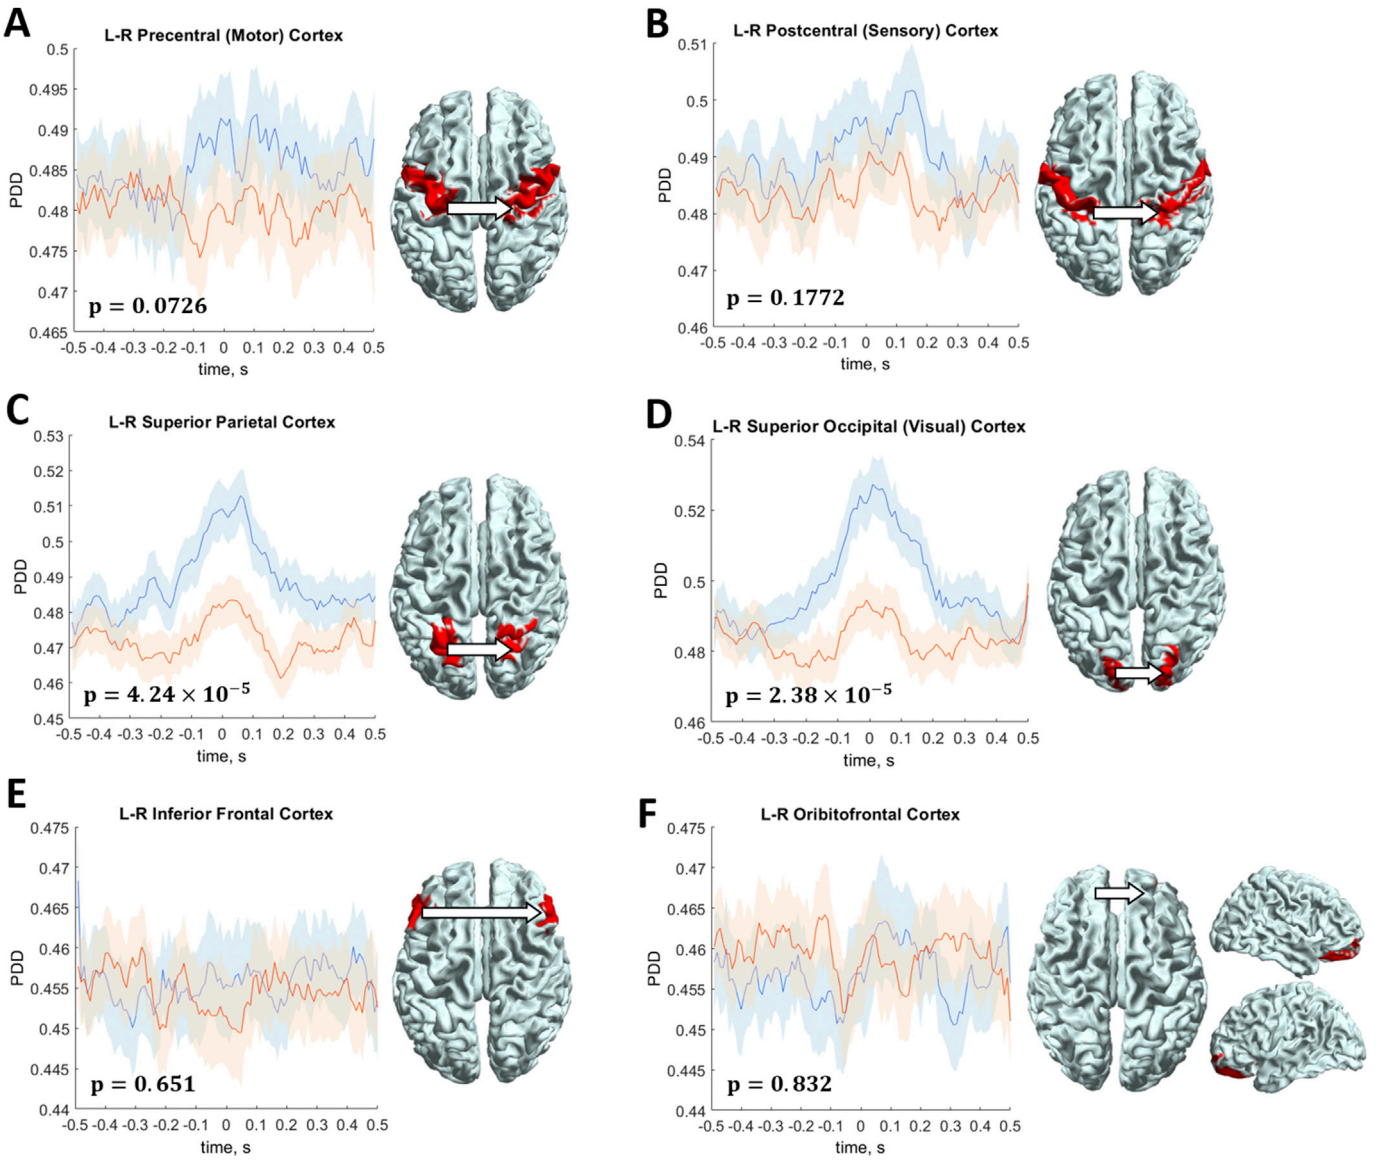

Fig. 6. Coherence underlying coincident bursts in resting state. In all 6 plots, phase difference derivative between two spatially separate brain regions is plotted as a function of time in the −0.5 s to +0.5 s window relative to the centre of a transient spectral burst. The blue lines represent the case where bursts were coincident between regions. The red line represents the equivalent case for non-coincident bursts. The shaded area represents standard error across subjects. 6 example connections are shown; left right motor (A) sensory (B) parietal (C) visual (D) and frontal (E–F) cortices. Brain regions are shown inset.

range is taken into account. The new method therefore allows for the complete spectral content of bursts to be elucidated. In fact, the bursts identified by the HMM were not confined to the beta band, but demonstrate a spectrally-specific shape that includes a prominent alpha peak. This was the case even when the HMM was asked to derive a higher number of states (see supplementary information).

Despite the differences there was still agreement between the HMM and the simple thresholding method, with >90% of the HMM identified bursts being mirrored by a period of high beta amplitude (>1.5 standard deviations from the mean). Additionally, when choosing a suitably high threshold (3 standard deviations from the mean), >70% of the bursts identified by the threshold technique were also found by the HMM. This shows that the transient spectral bursts identified by our HMM are (in the vast majority of cases) equivalent to the traditional beta bursts. This is of significant importance for future studies where, we argue, "beta bursts" should be treated as spectral events that span multiple frequency bands.

In the resting state we see marked variation of burst parameters with cortical region. Generally, the highest burst frequency, and shortest duration was in the frontal regions; the lowest frequencies occurred across the parietal, temporal and occipital regions. This finding could suggest a fundamentally different functional role of bursts in the frontal areas, but could equivalently suggest that the MEG signal in these regions simply has lower signal to noise ratio (likely resulting from the proximity of the MEG sensors to the scalp). Indeed the latter hypothesis is supported by a number of findings that beta band power is highest in parietal occipital and temporal regions. The highest amplitude bursts were found in the sensorimotor regions; this is consistent with the known role of beta in the sensorimotor system. Interestingly the longest duration bursts are observed in the occipital lobe. Whilst the reason for this is unclear, it could relate to the broad band nature of our HMM implementation – recall here that the HMM was applied in the 1–48 Hz band, and the burst state spectra clearly covers both the alpha and beta bands. It is possible that bursts in the occipital cortex are weighted towards an alpha dominance whereas bursts in the sensorimotor cortex are weighted towards beta.

Our task positive data, shown in Fig. 4, were in agreement with previous findings (Little et al., 2019). During movement, we observe a drop in the frequency of bursts; their duration is decreased, and the interval between bursts increased, compared to both the resting phase of the visuomotor task, and the true resting state. Conversely, during the beta rebound, as expected we see a marked increase in burst frequency, and also an increase in duration, with the time between bursts decreasing relative to rest, and resting state. Fig. 4C and 4D show significant differences between the bursts detected in resting state and those in the rest period of the task data. It is likely that this is due to the relatively short inter-trial interval of 8s since recent work has shown that it takes an average of 11s for activity to return to baseline (Pakenham et al., 2019). Interestingly, although the mean burst amplitude was higher during the beta rebound, this increase was relatively small compared to the large standard deviation of burst amplitudes potentially suggesting that burst amplitude is relatively stable, and that it is burst frequency and duration that are most affected by task performance. This finding agrees with previous work based upon a thresholding approach, and so provides further validation of the HMM as a means to detect bursting in MEG data.

Our connectivity findings show that the electrophysiological connectome, derived using amplitude envelope correlation, can also be generated by using a simple measure of the temporal coincidence of bursts. Whilst similarity between burst coincidence (calculated in the broad band) and AEC was apparent for a number of frequency bands, it was strongest for the beta band. The beta band connectome has garnered significant interest in recent years due to i) the resemblance of beta networks to the canonical resting state networks observed in fMRI (Brookes et al., 2011; Hipp et al., 2012), ii) the ability to estimate dynamic spectro-temporal changes in network connectivity (O'Neill et al., 2015; O'Neill et al., 2018), and iii) synergies between electrophysiological and structural data (Hunt et al., 2016). However, an implicit assumption in much of this work was that the oscillatory amplitude varied smoothly over time and correlations in this smooth variation drives connectivity. While there is evidence that the amplitude envelope of oscillations in the resting state is related to transient, bursting amplitude changes from previous HMM work (Baker et al., 2014), the extent to which single-region transient spectral bursts dominate functional connectivity has not previously been investigated.

Here we have shown that, by only taking into account the periods when the single-region bursts from two spatially separate regions overlap, the same static functional connectome emerges. In other words, the known beta connectome can be considered to be driven by coincident bursts. This is further supported by the fact that when periods of identified bursts are removed from an AEC calculation we see a reduction in average AEC values of 97.5%. There is a link between coincident bursts and AEC; high values of AEC are generated by coherent fluctuations in the beta envelope. Given that the HMM nominally pulls out bursts of high beta amplitude, the application of the HMM method effectively binarises the beta envelope, with periods of high amplitude set to one and low amplitude set to zero. Finding the burst overlap effectively amounts to correlating the binarised beta envelopes and therefore it is perhaps not surprising that there is a link between burst overlap and AEC. However, the binarisation removes much of the information content of the envelope - the burst state only accounts for a fraction of the total recording (e.g. 10.52% of the timecourse in sensorimotor cortex in the resting state accounts for overlapping bursts). The fact that we can get practically the same information from a simple burst analysis suggests that AEC is predominantly driven by bursts, rather than e.g. coherent but low amplitude fluctuations which have been lost by the binarisation. However this said, given that some structure remains in the AEC derived matrix even when bursts have been largely removed suggests that bursts are not solely responsible for the observed connectome.

Our final figure shows that when coincident bursts occur, they provide a short period during which the underlying 'oscillations' at the two regions involved are likely to be more coherent. This is consistent with the findings of (Engel et al., 2013), in which resting networks in MEG were shown to be well characterised by visits to short-lived transient brain states, each with spatially distinct patterns of oscillatory power and coherence in specific frequency bands. This obviously lends itself to the communication by coherence hypothesis – i.e. that the brain organises its oscillations in such a way as to provide temporally coincident windows of mutually high electrical potential within which we might expect passage of action potentials, and hence information (Fries, 2015). Bringing together this finding with other work it is tempting to speculate that coincident bursts offer a means for a broad brain network to provide some inhibitory influence on a specific region – e.g. motor cortex; the more frequent the bursts, the more inhibition that region receives. This finding of short periods of coherence also offers some explanation of previous methodological findings. In general it has been shown that amplitude envelope correlation offers a more robust way to measure functional connectivity compared to coherence – at least in individual subjects (Colclough et al., 2016; Liuzzi et al., 2017). Here we show that bursts represent punctate events, occurring on average once per second in the resting state, which are coherent between regions. This means that over all time (for example, in a 300 s resting state recording), we will only see coherence for a very small fraction of that time. It follows that estimated coherence over an entire recording would therefore represent a less reliable measure. Conversely in the amplitude envelope correlation case, it's the beta bursting compared to rest which drives the connectivity measure, and therefore it benefits from both the periods of coherence (with a high beta envelope) and periods of no coherence with a lower envelope.

There are a number of limitations of the current approach which should be understood. First, the HMM requires a number of parameter choices; these include the selection of the number of states, the length of the window used to compute the autocovariance matrices (which characterise each state), and the threshold on the probability timecourses to select which state is "active". It should be noted that parameter choices are typically needed in all analysis approaches. Nonetheless, the HMM output is reasonably robust to changing the number of states (see Supplementary Fig. 1): we found the burst state spectrum remains constant regardless of whether the model employs 3, 6 or 10 states. Also, the spatial signatures of burst frequency, duration, amplitude, and interval remain qualitatively similar regardless of the choice. However, quantitatively the burst duration, frequency and interval times all change significantly when switching from 3 to 6 to 10 states. This is likely due to the mutual exclusivity assumption made by the model – but also means that burst properties are subjective, depending on parameter choice. The length of the window used to define our autocovariance matrices was 230 ms; this value depends on the frequency resolution one wishes to capture. For example, making the window shorter focusses it on higher frequency components (Vidaurre et al., 2018). So here again, a subjective choice is required dependant on the question being asked. The value on the probability threshold is however less subjective since, as shown in Fig. 2, the probabilities are relatively binarised anyway. These limitations on parameter choice should be considered carefully in future studies using this technique.

Aside from parameter choice, one of the key assumptions made by the HMM is that the states are mutually exclusive. This means that the model is unable to cope with the co-occurrence of multiple states. We are not claiming that this represents the true physiological nature of brain activity. Rather, that this assumption and this method provides a useful description of brain activity. Indeed here, we show that despite the simplification, the binarised burst timecourses enable derivation of the electrophysiological connectome which is in good agreement with that derived via AEC. Of course the "true" connectome (i.e. an underlying ground truth) remains unknown. However, it is known that AEC has significant correlation with the resting state networks that are measured using functional magnetic resonance imaging (Hipp et al., 2012; Brookes et al., 2011). Given the vast literature showing the critical role played by these networks in supporting healthy brain function, and their breakdown in a wide variety of disorders, it is reasonable to consider this as a useful description of the connectome. The fact that these network

connectivities can now be understood in the context of transient spectral electrophysiological bursts, will represent an important step in our understanding of how those networks are mediated.

5. Conclusion

We have demonstrated that the time-delay-embedded HMM offers a new way to interrogate MEG data, delineating oscillatory bursts which are in good agreement with the simple thresholding techniques used in the literature. However, unlike arbitrary thresholding approaches applied to oscillatory amplitude in a single frequency band, here we have used the HMM to look for specific morphological (spectral) patterns in single-region timecourse data; in this way it provides an objective means to derive transient spectral bursts with a meaningful threshold in broadband data. We showed how the distribution of bursts changes across the cortex, with the lowest burst frequencies, highest amplitudes, and highest durations in the parietal, occipital and temporal lobes. Analysis of task data was in good agreement with previous work in showing that the 'classical' movement related beta decrease actually corresponds to a lower burst frequency, and shorter burst duration whereas the post movement beta rebound corresponds to a higher burst frequency, and longer duration. Finally, we have shown that the well-known beta connectome, which is typically calculated using amplitude envelope correlation, can also be derived by a simple measure of burst overlap. Our follow up analyses suggest that when bursts are coincident, they facilitate a period of phase locking which likely encourages communication by coherence. In summary, our paper offers a new methodology for both burst identification and connectivity analysis, which will be important for future MEG investigations of neural oscillations and their perturbation by disease.

Author contributions

ZS. Conceptualisation, Experimental design, Data analysis, Data interpretation, Writing paper.

AQ. Conceptualisation, Data analysis, Experimental design, Data interpretation, Writing paper.

DV. Conceptualisation, Experimental design, Data interpretation, Writing paper.

LL. Conceptualisation, Data interpretation, Writing paper.

LG. Conceptualisation, Data interpretation, Writing paper.

BH. Conceptualisation, Data collection, Data analysis, Data interpretation.

GO. Conceptualisation, Data interpretation.

DP. Conceptualisation, Data interpretation, Writing paper.

KM. Conceptualisation, Data interpretation, Writing paper.

PM. Conceptualisation, Data interpretation, Writing paper.

MW. Conceptualisation, Experimental design, Data interpretation, Writing paper.

MB. Conceptualisation, Experimental design, Data analysis, Data interpretation, Writing paper.

Acknowledgements

Data collection for this paper was funded by a Medical Research Council (MRC) New Investigator Research Grant (MR/M006301/1) and a MRC Partnership Grant (MR/K005464/1). We also acknowledge the UK Quantum Technology Hub for Sensors and Metrology, funded by the Engineering and Physical Sciences Research Council (EPSRC) (EP/M013294/1). Funding from EPSRC and MRC (grant number EP/L016052/1) also provided a PhD studentship for ZS through the Oxford Nottingham Biomedical Imaging Centre for Doctoral Training. This research was also supported by the NIHR Oxford Health Biomedical Research Centre, a Wellcome Trust Strategic Award (Grant 098369/Z/12/Z). MWW is supported by a Wellcome Investigator Award (106183/Z/14/Z). The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust (203139/Z/16/Z).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2020.116537.

References

Baker, A.P., Brookes, M.J., Rezek, I.A., Smith, S.M., Behrens, T., Smith, P.J.P., Woolrich, M., 2014. Fast transient networks in spontaneous human brain activity. eLife 2014.

Breakspear, M., Williams, L.M., Stam, C.J., 2004. A novel method for the topographic analysis of neural activity reveals formation and dissolution of 'Dynamic Cell Assemblies'. J. Comput. Neurosci. 16, 49–68.

Brookes, M.J., Groom, M.J., Liuzzi, L., Hill, R.M., Smith, H.J.F., Briley, P.M., Hall, E.L., Hunt, B.A.E., Gascoyne, L.E., Taylor, M.J., Liddle, P.F., Morris, P.G., Woolrich, M.W., Liddle, E.B., 2018. Altered temporal stability in dynamic neural networks underlies connectivity changes in neurodevelopment. Neuroimage 174, 563–575.

Brookes, M.J., Vrba, J., Robinson, S.E., Stevenson, C.M., Peters, A.P., Barnes, G.R., Hillebrand, A., Morris, P.G., 2008. Optimising experimental design for MEG beamformer imaging. Neuroimage 39, 1788–1802.

Brookes, Matthew J., Woolrich, Mark, Henry, Luckhoo, Price, Darren, Hale, Joanne R., Stephenson, Mary C., Barnes, Gareth R., Smith, Stephen M., Morris, Peter G., 2011. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proc. Natl. Acad. Sci. 108, 16783–16788.

Buard, I., Kronberg, E., Steinmetz, S., Hepburn, S., Rojas, D.C., 2018. Neuromagnetic beta-band oscillations during motor imitation in youth with autism. Autism Research and Treatment 2018.

Colclough, G.L., Woolrich, M.W., Tewarie, P.K., Brookes, M.J., Quinn, A.J., Smith, S.M., 2016. How reliable are MEG resting-state connectivity metrics? Neuroimage 138, 284–293.

Colclough, G.L., M.J. Brookes, S.M. Smith, and M.W. Woolrich. 2015. 'A Symmetric Multivariate Leakage Correction for MEG Connectomes', *Neuroimage*, (in press).

de Pasquale, F., Della Penna, S., Snyder, A.Z., Lewis, C., Mantini, D., Marzetti, A., Belardinelli, P., Ciancetta, L., Pizzella, V., Romani, G.L., Corbetta, M., 2010. Temporal dynamics of spontaneous MEG activity in brain networks. Proceedings of the National Academy of Science USA 107, 6040–6045.

Ede, Freek van, Quinn, Andrew J., Woolrich, Mark W., Nobre, Anna C., 2018. Neural oscillations: sustained rhythms or transient burst-events? Trends Neurosci. 41, 415–417.

Engel, A.K., Gerloff, C., Hilgetag, C.C., Nolte, G., 2013. Intrinsic coupling modes: multiscale interactions in ongoing brain activity. Neuron 80, 867–886.

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

Fries, P., 2015. Rhythms for cognition: communication through coherence. Neuron 88, 220–235.

Fry, A., Mullinger, K.J., O'Neill, G.C., Barratt, E.L., Morris, P.G., Bauer, M., Folland, J.P., Brookes, M.J., 2016. Modulation of post-movement beta rebound by contraction force and rate of force development. Hum. Brain Mapp. 37, 2493–2511.

Gaetz, W., Edgar, J.C., Wang, D.J., Roberts, T.P., 2011. Relating MEG measured motor cortical oscillations to resting gamma-aminobutyric acid (GABA) concentration. Neuroimage 55, 616–621.

Gross, J., Kujala, J., Hämäläinen, M., Timmermann, L., Schnitzler, A., Salmelin, R., 2001. Dynamic imaging of coherent sources: studying neural interactions in the human brain. Proc. Natl. Acad. Sci. U. S. A. 98, 694–699.

Hipp, Joerg F., Hawellek, David J., Corbetta, Maurizio, Siegel, Markus, Engel, Andreas K., 2012. Large-scale cortical correlation structure of spontaneous oscillatory activity. Nat. Neurosci. 15, 884–890.

Huang, M.X., Mosher, J.C., Leahy, R.M., 1999. A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. Phys. Med. Biol. 44, 423–440.

Hunt, B.A.E., Liddle, E.B., Gascoyne, L.E., Magazzini, L., Routley, B.C., Singh, K.D., Morris, P.G., Brookes, M.J., Liddle, P.F., 2018. Attenuated post-movement beta rebound associated with schizotypal features in healthy people. Schizophr. Bull. 45 (4), 883–891 sby117-sby17.

Hunt, B.A., Tewarie, P.K., Mougin, O.E., Geades, N., Jones, D.K., Singh, K.D., Morris, P.G., Gowland, P.A., Brookes, M.J., 2016. Relationships between cortical myeloarchitecture and electrophysiological networks. Proc. Natl. Acad. Sci. U. S. A. 113 (47), 13510–13515.

Jones, S.R., 2016. When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning. Curr. Opin. Neurobiol. 40, 72–80.

Jurkiewicz, M.T., Gaetz, W.C., Bostan, A.C., Cheyne, D., 2006. Post-movement beta rebound is generated in motor cortex: evidence from neuromagnetic recordings. Neuroimage 32, 1281–1289.

Little, S., Bonaiuto, J., Barnes, G.R., Bestmann, S., 2019. Motor cortical beta transients delay movement initiation and track errors. PLoS Biol. 17 (10).

Liuzzi, L., Gascoyne, L.E., Tewarie, P.K., Barratt, E.L., Boto, E., Brookes, M.J., 2017. Optimising experimental design for MEG resting state functional connectivity measurement. Neuroimage 155, 565–576.

Lundqvist, Mikael, Rose, Jonas, Herman, Pawel, Brincat, Scott L., Buschman, Timothy J., Miller, Earl K., 2016. Gamma and beta bursts underlie working memory. Neuron 90, 152–164.

Muthukumaraswamy, S.D., Myers, J.F., Wilson, S.J., Nutt, D.J., Lingford-Hughes, A., Singh, K.D., Hamandi, K., 2013. The effects of elevated endogenous GABA levels on movement-related network oscillations. Neuroimage 66, 36–41.

O'Neill, George C., Bauer, Markus, Woolrich, Mark W., Morris, Peter G., Barnes, Gareth R., Brookes, Matthew J., 2015. Dynamic recruitment of resting state sub-networks. Neuroimage 115, 85–95.

O'Neill, George, C., Tewarie, Prejaas, Vidaurre, Diego, Liuzzi, Lucrezia, Woolrich, Mark W., Brookes, Matthew J., 2018. Dynamics of large-scale electrophysiological networks: a technical review. Neuroimage 180, 559–576.

Pakenham, D.O., Quinn, A.J., Fry, A., Francis, S.T., Woolrich, M., Brookes, M.J., Mullinger, K.J., 2019. Post-stimulus beta responses are modulated by task duration. Neuroimage. https://doi.org/10.1016/j.neuroimage.2019.116288. In press.

Pfurtscheller, G., Lopes da Silva, F.H., 1999. Event-related EEG/MEG synchronization and desynchronization: basic principles. Clin. Neurophysiol. 110, 1842–1857.

Quinn, Andrew J., Vidaurre, Diego, Abeysuriya, Romesh, Becker, Robert, Nobre, Anna C., Woolrich, Mark W., 2018. Task-evoked dynamic network analysis through hidden Markov modeling. Front. Neurosci. 12.

Robinson, S., Vrba, J., 1998. Functional neuroimaging by synthetic aperture magnetometry. In: Yoshimoto, T., Kotani, M., Kuriki, S., Karibe, H., Nakasato, N., Tohoku (Eds.), Recent Advances in Biomagnetism. Univ. Press, Sendai, Japan, pp. 302–305.

Robson, S.E., Brookes, M.J., Hall, E.L., Palaniyappan, L., Kumar, J., Skelton, M., Christodoulou, N.G., Qureshi, A., Jan, F., Katshu, M., Liddle, E.B., Liddle, P.F., Morris, P.G., 2016. Abnormal visuomotor processing in schizophrenia. NeuroImage Clinical 12, 869–878.

Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Phys. Med. Biol. 32, 11–22.

Sherman, M.A., Lee, S., Law, R., Haegens, s., Thorn, C.A., Hämäläinen, M.S., Moore, C.I., Jones, S.R., 2016. Neural mechanisms of transient neocortical beta rhythms: converging evidence from humans, computational modeling, monkeys, and mice. Proc. Natl. Acad. Sci. U. S. A. 113, E4885–E4894.

Shin, Hyeyoung, Law, Robert, Tsutsui, Shawn, Moore, Christopher I., Jones, Stephanie R., 2017. The rate of transient beta frequency events predicts behavior across tasks and species. eLife 6.

Tewarie, P.K., Bright, M.G., Hillebrand, A., Robson, S.E., Gascoyne, L.E., Morris, P.G., Meier, J., Van Mieghem, P., Brookes, M.J., 2016. Predicting haemodynamic networks using electrophysiology: the role of non-linear and cross-frequency interactions. Neuroimage 130, 273–292.

Tewarie, P., Hunt, B.A.E., O'Neill, G.C., Byrne, A., Aquino, K., Bauer, M., Mullinger, K.J., Coombes, S., Matthew, J., Brookes, 20192019. Relationships between neuronal oscillatory amplitude and dynamic functional connectivity. Cerebr. Cortex 29 (6), 2668–2681 bhy136-bhy36.

Tinkhauser, Gerd, Alek, Pogosyan, Little, Simon, Beudel, Martijn, Herz, Damian M., Tan, Huiling, Brown, Peter, 2017a. The modulatory effect of adaptive deep brain stimulation on beta bursts in Parkinson's disease. Brain 140, 1053–1067.

Tinkhauser, Gerd, Alek, Pogosyan, Tan, Huiling, Herz, Damian M., Kühn, Andrea A., Brown, Peter, 2017b. Beta burst dynamics in Parkinson's disease OFF and ON dopaminergic medication. Brain 140, 2968–2981.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273–289.

Vidaurre, D., Hunt, L.T., Quinn, A.J., Hunt, B.A.E., Brookes, M.J., Nobre, A.C., Woolrich, M.W., 2018. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. Nat. Commun. 9, 2987.

Vidaurre, D., Quinn, A.J., Baker, A.P., Dupret, D., Tejero-Cantero, A., Woolrich, M.W., 2016. Spectrally resolved fast transient brain states in electrophysiological data. Neuroimage 126, 81–95.

Woolrich, M.W., Baker, A., Luckhoo, H., Mohseni, H., Barnes, G., Brookes, M.J., Rezek, L., 2013. Dynamic state allocation for MEG source reconstruction. Neuroimage 77, 77–92.