



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/166304/>

Version: Accepted Version

Article:

Karapanagiotidis, Theodoros, Vidaurre, Diego, Quinn, Andrew et al. (2020) The psychological correlates of distinct neural states occurring during wakeful rest. Scientific Reports. 21121. ISSN: 2045-2322

<https://doi.org/10.1038/s41598-020-77336-z>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

The psychological correlates of distinct neural states occurring during wakeful rest

Theodoros Karapanagiotidis^{1,*}, Diego Vidaurre², Andrew J. Quinn^{2,3}, Deniz Vatansever³, Giulia L. Poerio⁴, Adam Turnbull¹, Nerissa Siu Ping Ho¹, Robert Leech⁵, Boris C. Bernhardt⁶, Elizabeth Jefferies¹, Daniel S. Margulies⁷, Thomas E. Nichols⁸, Mark W. Woolrich², Jonathan Smallwood¹

1 Department of Psychology, York Neuroimaging Centre, University of York, York YO10 5DD, United Kingdom

2 Oxford Centre for Human Brain Activity, Wellcome Centre for Integrative Neuroimaging, Department of Psychiatry, University of Oxford, Oxford OX3 7JX, United Kingdom

3 Institute of Science and Technology for Brain-inspired Intelligence, Fudan University, Shanghai 200433, PR China

4 Department of Psychology, University of Essex, Colchester, Essex CO4 3SQ, United Kingdom

5 Centre for Neuroimaging Science, Kings College, London SE5 8AF, United Kingdom

6 McConnell Brain Imaging Centre, Montréal Neurological Institute, McGill University, Montréal, QC H3A 2B4, Canada

7 Brain & Spine Institute (ICM), National Center for Scientific Research, Paris 75013, France

8 Oxford Centre for Functional MRI of the Brain, Wellcome Centre for Integrative Neuroimaging, Nuffield Department of Clinical Neurosciences, University of Oxford, Oxford OX3 9DU, United Kingdom

* theodoros.karapanagiotidis@york.ac.uk

Abstract

When unoccupied by an explicit external task, humans engage in a wide range of different types of self-generated thinking. These are often unrelated to the immediate environment and have unique psychological features. Although contemporary perspectives on ongoing thought recognise the heterogeneity of these self-generated states, we lack both a clear understanding of how to classify the specific states, and how they can be mapped empirically. In the current study, we capitalise on advances in machine learning that allow continuous neural data to be divided into a set of distinct temporally re-occurring patterns, or states. We applied this technique to a large set of resting state data in which we also acquired retrospective descriptions of the participants' experiences during the scan. We found that two of the identified states were predictive of patterns of thinking at rest. One state highlighted a pattern of neural activity commonly seen during demanding tasks, and the time individuals spent in this state was associated with descriptions of experience focused on problem solving in the future. A second state was associated with patterns of activity that are commonly seen under less demanding conditions, and the time spent in it was linked to reports of intrusive thoughts about the past. Finally, we found that these two neural states tended to fall at either end of a neural hierarchy that is thought to reflect the brain's response to cognitive demands. Together, these results demonstrate that approaches which take advantage of time-varying changes in neural function can play an important role in understanding the repertoire of self-generated states. Moreover, they establish that important features of self-generated ongoing experience are related to variation along a similar vein to those seen when the brain responds to cognitive task demands.

Introduction

When unoccupied by an external task, humans often engage in complex patterns of self-generated thinking. Often studied under the rubric of mind-wandering^{1;2} or daydreaming³, these patterns of thought often share a common feature of being unrelated to events in the here and now. However, they are also heterogeneous in content⁴ and outcome⁵, while their intrinsic nature constitutes a challenge to their measurement⁶. Although there is a growing understanding that self-generated states constitute an important feature of human cognition, we lack an agreed upon classification on their defining features^{7;8}, as well as the tools with which to study them⁹. The current study examined whether it is possible to gain insight into the repertoire of self-generated states by applying advanced machine learning methods to recordings of neural data during periods of wakeful rest, and using these to predict measures of experience.

Our study builds on an emerging literature using advanced analyses techniques to understand the organisation of self-generated cognition. One common method is to reduce the dimensionality of self-reported data in order to identify latent variables that make up participants' descriptions of their experience. For example, Ruby and colleagues¹⁰ used principal component analysis (PCA) and identified patterns of self-generated thought that were distinguished by their associations with either the past or the future. Using lag analysis, they demonstrated that these two states had unique correlates with subsequent mood – prospective self-generated thought increased positive mood, while past thoughts were associated with the reverse pattern. Other authors have used a similar approach to establish that prospective self-generated thoughts are often structured and realistic¹¹ and that they may play a role in the consolidation of personal goals into concrete plans¹². Another approach involves the application of clustering techniques to experience sampling data. For example, Andrews-Hanna and colleagues¹³ used hierarchical clustering to determine distinct patterns of negative and self-relevant thought and specific thinking, each with unique psychological correlates. Other studies have employed techniques that explored the temporal features of experience sampling data. For example, Zanesco¹⁴ employed a temporal clustering method to experience sampling data generated across several different experimental paradigms. They found several patterns of ongoing cognition that were similar across tasks, and identified regularities in the order with which the hidden cognitive states emerge. Our study builds on these findings by applying a temporal decomposition method to neural data recorded during a period of wakeful rest, a situation which has a high prevalence of self-generated mentation¹⁵ and in which prior studies have successfully linked neural activity to retrospective reports of experience^{16;17;18}. Our aim was to determine whether the application of advanced machine learning methods on neural data revealed states that had reliable associations with an individual's self-reported experience during that period.

In our study, therefore, we measured neural function during wakeful rest in a large cohort of healthy individuals using functional magnetic resonance imaging (fMRI). These participants completed in the scanner a set of questions related to their experiences during the scan, and a subset completed questionnaire measures of physical and mental health at a later session. We applied hidden Markov modelling (HMM) to the neural data to reveal distinct brain states corresponding to specific periods of temporally reoccurring patterns of neural mean activity and functional connectivity¹⁹. We then used the dwell-times of the states as predictors in a multivariate multiple regression with the reports given at the end of the scan as the outcome measures, aiming to explore their particular relationships in detail. Finally, we investigated how these states are distributed across a multi-dimensional space, formed by three well-described neural hierarchies based on the differentiation of whole-brain connectivity

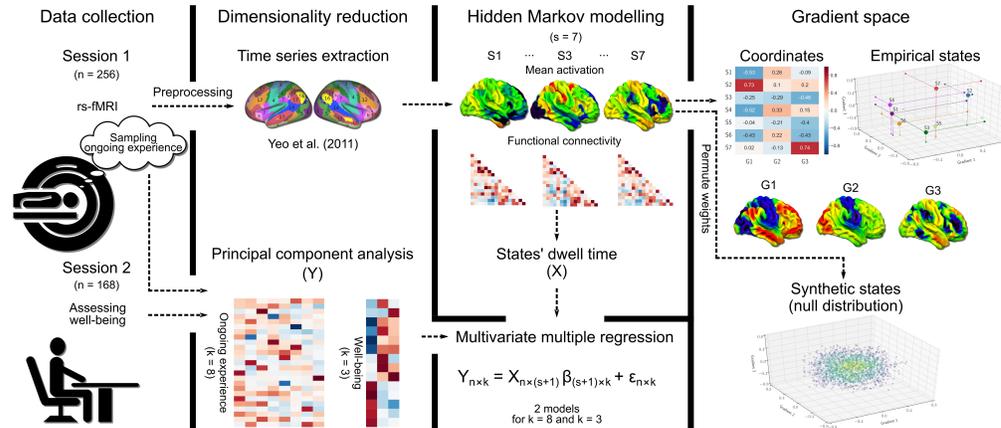


Figure 1. Schematic of the workflow of the current study. We measured neural activity at rest in a sample of $n=256$ individuals who then completed a retrospective measure of the experiences they had during this period. In a subset of the individuals, we also acquired measures of well-being. We then reduced the dimensionality of both data sets and applied hidden Markov modelling to the neuroimaging time series. We performed a multiple regression analysis in which we calculated the association between the mean dwell-time of each state to the patterns of experiences calculated at rest, as well as to trait measures of well-being. Finally, we calculated the association of the identified brain network states to the first three neurocognitive hierarchies identified by²⁰, as well as to a distribution of synthetic states generated through permutation. k : number of PCA components, s : number of states, G : the three gradients from²⁰

patterns²⁰. A number of recent studies have highlighted a composite of large-scale hierarchies as a way to understand human cortical organisation in health and disease^{21;22;23}. These gradients can present information about the non-linear functional interactions between different neural systems which, as recent studies have shown, can be particularly beneficial in the investigation of the multi-dimensional features of higher-order cognition^{24;25}. To that end, our study aimed to examine the possibility that these hierarchies represent a functional landscape that describes how organised transient neurocognitive states emerge. The study workflow is presented in Figure 1.

Results

Identifying reoccurring neural states at rest

We applied hidden Markov modelling to the resting state fMRI data (see Methods). This method uses Bayesian inference to identify reoccurring patterns of brain network activity; see^{26;19} for an empirical validation of this approach. Here, we generated HMM solutions capturing mean activity and functional connectivity patterns within the neural data recorded from our sample. We opted for a 7-state solution, however, our results highlighting associations with behaviour remained broadly similar for a decomposition of 9 states as well. We ran the algorithm 10 times in each case and found that the 7-state solution produced the same decomposition across all iterations. For this reason, we did not ask for a solution with a smaller number of states, so as not to unnecessarily reduce the algorithm’s ability to provide a detailed description of the dataset. Furthermore, for our data and model parameters, the 7-state solution was relatively more stable and had greater split-half reliability than the 9-state solution (see Fig. S2). Accordingly, we focus on the 7-state solution and report a parallel analysis of the 9-state solution in the supplementary Materials (Fig. S6 and S8). Figure 2 shows the spatial description of neural mean activity in each state and the associated cognitive terms

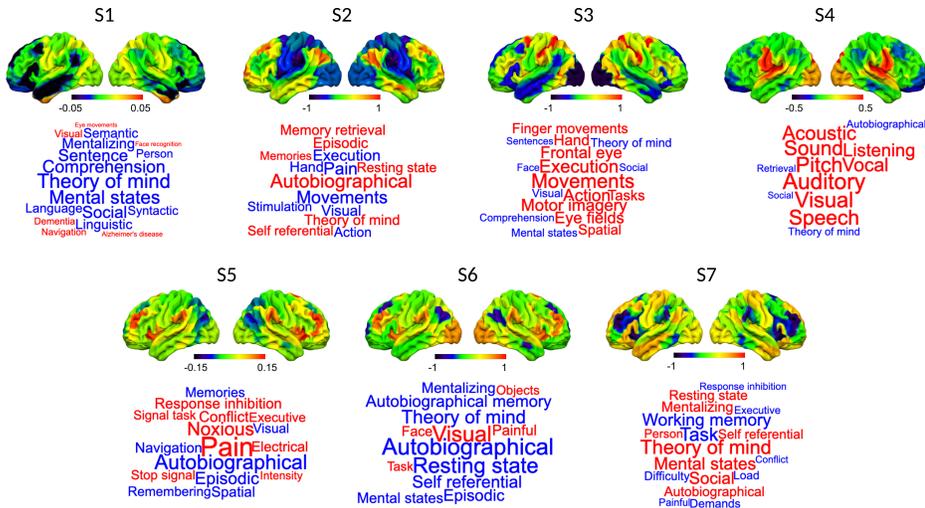


Figure 2. Spatial distribution of mean neural activity characteristic of each state and their meta-analytic associations. The brains show the mean neural activity that is characteristic of each state. Meta-analytic patterns were generated using Neurosynth and are displayed in the form of word clouds. The size of the words reflects the importance of the term, and the colour reflects the direction of the loading (red = positive and blue = negative).

produced by a meta-analysis using Neurosynth²⁷, displayed in the form of word clouds.

The relationship between neural states and experience

Having identified a set of relatively stable recurring patterns within our data, we next explored how these naturally occurring neural states are related to the experiences reported by our participants. To answer this question, we performed a multi-variate analysis of co-variance, in which the mean dwell-time of each state was an explanatory variable and the self-reported data was the dependent variable (see Methods for details of the analyses; the specific questions can be found in Table S1). We included motion, age and gender as co-variates of no interest.

This analysis revealed that the mean dwell-time of states 3 and 7 had a significant multivariate association with patterns of reports made by our participants at the end of the scan ($F(8, 238) = 2.22$, $p = 0.027$, Wilks' $\Lambda = 0.931$, partial $\eta^2 = 0.069$ for state 3 and $F(8, 238) = 2.18$, $p = 0.03$, Wilks' $\Lambda = 0.932$, partial $\eta^2 = 0.068$ for state 7). These results are presented in Figure 3. It can be seen in the radar plot that state 3 is dominated by task-positive systems (frontal parietal and dorsal attention) as well as the medial temporal subsystem of the default mode network (DMN). Experientially, it was associated with a state of autobiographical planning, reflecting high scores on “Future” and “Problem solving” thoughts. In contrast, state 7 was associated with somatomotor, the core and lateral temporal subsystems of the DMN and the limbic system. This state was associated with reports of unintentional, intrusive thoughts about the past. Our analysis therefore identified two distinct neural states that were uniquely associated with states associated with experiential foci on different temporal epochs (the past and the future), a pattern that is routinely seen in experience sampling studies e.g.^{10;28;29}.

Psychological studies have shown that patterns of ongoing thought often have associations with negative affect, and in particular those associated with the past^{10;31;29}. Based on this evidence, our next analysis aimed at identifying whether the mean dwell-time of the states varied with trait measures of well-being recorded in a subset of these participants. This analysis found that two of the states had a multivariate

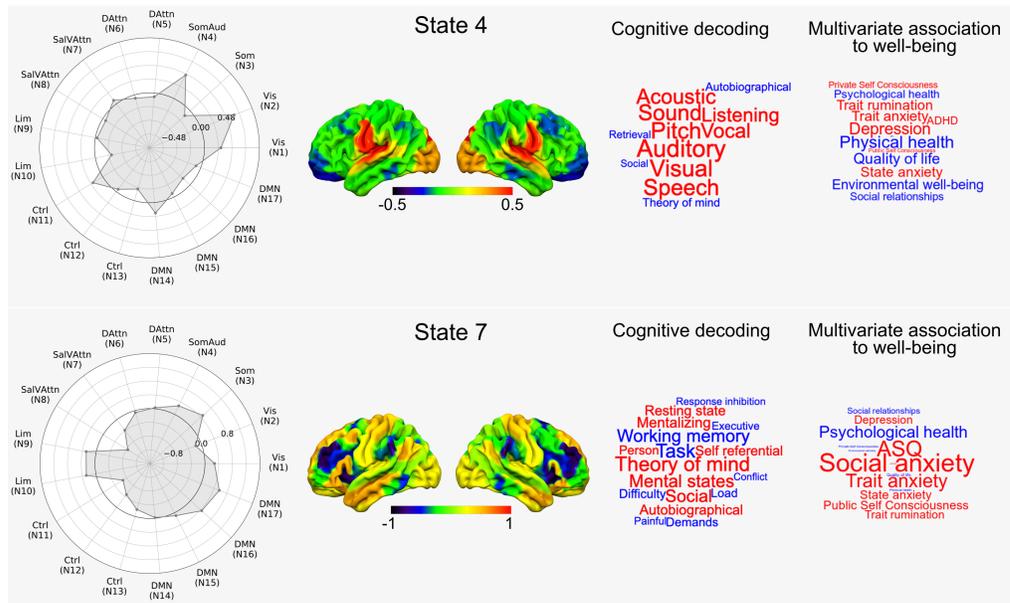


Figure 4. Associations between the mean dwell-time of different neural states and well-being. This figure shows the spatial distribution of the two states as well as a meta-analysis of these spatial maps using Neurosynth (cognitive decoding), which are displayed in the form of word clouds. The word clouds on the right show the pattern of answers associated with each state. In addition, we show for each state, the relative contributions of each of the networks used in the HMM analyses, in the form of radar plots. N1-N17: Network 1 to network 17 from the Yeo network parcellation³⁰. Vis: Visual, Som: Somatomotor, SomAud: Somatomotor/Auditory, DAttn: Dorsal attention, SalVAttn: Salience/Ventral attention, Lim: Limbic, Ctrl: Control, DMN: Default.

generated through HMM.

The relationship between naturally occurring states at rest and neural hierarchies

Having identified network states related to experience and trait measures of well-being, we next explored their association with three well-described neural hierarchies that highlight gradients in connectivity patterns over space. We chose the first three hierarchies from²⁰ which represent the division of unimodal and transmodal systems (gradient 1), the distinction between vision and motor (gradient 2) and the patterns seen during the brain's response to tasks relative to rest (gradient 3). These were generated using diffusion embedding, a non-linear dimensionality reduction technique, on the Human Connectome Project Data, and identified components describing the maximum variance in functional connectivity patterns (see Methods and initial paper for further details). The spatial distribution of each of these gradients is presented in Figure 5a along with a meta-analysis of these maps using Neurosynth.

The aim of this analysis was to understand whether there is a relationship between the states inferred by the HMM (see Fig. 2) and the spatial maps which represent these pre-established hierarchies. To achieve this goal, we calculated the pair-wise similarity between each of the hierarchies and each of the states (see Methods). We compared the associations between the real states inferred in our analyses to a null distribution of synthetic states. The null distribution was generated by randomly permuting the weights describing the relative contribution of each brain region to the mean neural

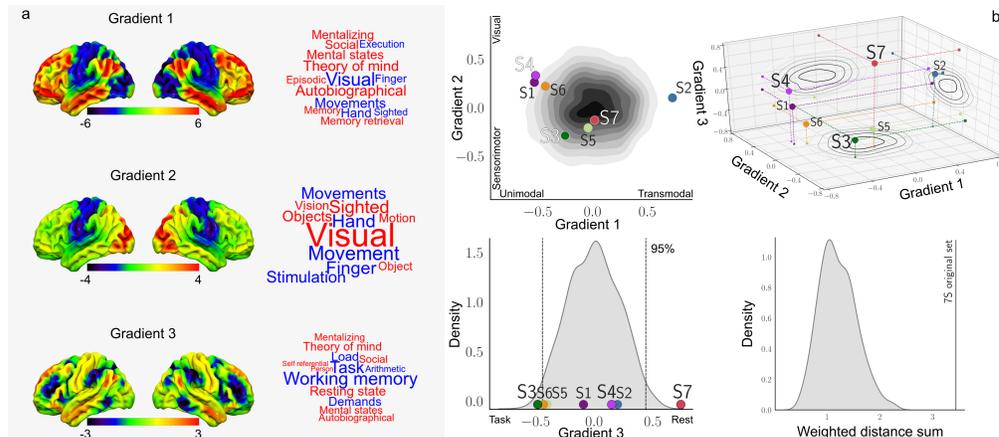


Figure 5. Relationship between naturally occurring states at rest and three well-established cognitive hierarchies. **a)** This panel shows the spatial distribution of the three neurocognitive hierarchies identified by²⁰, as well as a meta-analysis of these spatial maps using Neurosynth, which are displayed in the form of word clouds. Font size represents the strength of the association and font colour its sign (red for positive and blue for negative values). **b)** This panel shows the spatial similarity between the spatial maps describing these neural hierarchies and the set of spatial maps generated from the HMM decomposition of the real data. In these figures, the contour plots projected onto each plane describe the distribution of the synthetic states generated through permutation. The upper and lower left plots show where each empirical state falls on the gradient 1-2 plane, and on the gradient 3 axis respectively. The area between the dotted lines in the density plot shows where 95% of the synthetic states fall on the gradient 3 axis. The lower right-hand histogram shows how far away the states are from the origin of gradient space as a set. This is calculated as the sum of the weighted distance from the origin of each point in each set. The line shows where the set of the empirical states fall, whereas the shaded histogram shows the distribution of the synthetic states sets. States 3, 4 and 7, whose dwell-time was associated with psychological measures, are highlighted.

activity captured by each state (see Methods). The results of these analyses are presented on Figure 5b. The contour plot and histogram on the left show where the real data fall on a subset of the dimensions, while the three-dimensional scatter plot show the same data in a combined manner. In these plots, the null distribution generated through permutation is represented as the shaded area in grey. We also generated a video that shows the data in a 3-d form (https://github.com/tkarapan/hmm-large-scale-hierarchies/blob/master/TS_Null.mp4).

The synthetic states clustered towards the middle of the three dimensions, indicating that states generated in a random manner tend to not be very similar in spatial terms to the three pre-established neurocognitive hierarchies²⁰. In contrast, the states generated based on real data fall, on average, away from the centre and towards the outer edge of the distribution of synthetic states. Notably, the states with a significant association to experience, states 3 and 7, fall towards the mid-point of both gradient 1 and 2 and, in contrast, are maximally dissociated along the dimension describing how the brain responds to a task; state 3 is more similar to the task positive end of gradient 3, while state 7 shows the reverse. Finally, it can be seen that state 4, which alongside state 7 was associated with trait measures of well-being, falls towards the unimodal end of gradient 1.

To quantify whether the apparent difference between the real and permuted data was robust, we calculated the sum of the weighted distance from the origin for each of the sets of states generated through permutation of this space, and compared this to the same value from the real data. The distribution of these values is presented in the

right-hand histogram in Figure 5b, where it can be seen that the weighted distance sum of the real states from the centre of this space is higher than any of the states generated synthetically.

Together, these analyses suggest that states determined through the application of HMM to our real data have an association with the spatial maps describing neural hierarchies that is greater than chance. To identify the robustness of this effect, we performed a parallel analysis based on the 9-state solution to the current data, as well as the published states inferred from applying HMM with the same model parameters on the Human Connectome Project data²⁶, finding comparable results (see Fig. S10). Together, these analyses show that the states that occur naturally at rest, as approximated through HMM, fall at the extremities of a co-ordinate space defined by well-established neural hierarchies. This result implies that macro-scale hierarchies can be thought of as constraining the state space from which the intrinsic network states emerge.

Discussion

The current study set out to understand whether it is possible to shed light on the repertoire of self-generated states that an individual engages in, through the application of advanced machine learning methods to neural data recorded during periods of wakeful rest. Our results indicated a significant association between the amount of time an individual spends in two neural states and the experiential reports gained at the end of the scan. One state had features mimicking patterns of neural activity seen during complex tasks e.g.^{32;33} and was associated with patterns of thoughts focused on future problem solving. A second state highlighted the dominance of the default mode network (DMN)³⁴ and was linked to reports of unintentional intrusive thought about the past. Importantly, this state was also linked to traits associated with negative affect (anxiety, depression and rumination). Comparing these results to those from³⁵, where we ran a time-averaged functional connectivity analysis for a subset of this sample, we found that the time-varying analyses used in this study provided additional information which would be otherwise left unexposed. Together, our study adds to a growing consensus on the utility of methods that explore neural activity from a dynamical perspective, as a tool for understanding patterns of ongoing cognition and behaviour^{9;36}.

Our analyses, therefore, establish that techniques that parse continuous neural data into time-varying states can be used as a tool to empirically constrain accounts of the qualitative patterns of ongoing thought. By relating the amount of time an individual spent in specific neural states to the pattern of thoughts reported at the end of the scan, we highlighted many of the themes of self-generated thought that are common within the literature. For example, patterns of future problem solving identified in our study are consistent with a prospective bias to ongoing thought that develops in young adulthood^{37;38} and is common in multiple cultures (Australia,³⁸ Belgium,¹¹ China,³⁹ Germany,¹⁰ Japan,⁴⁰ U.K.¹⁵ USA,^{41;42}). This may reflect a mode of autobiographical planning that is hypothesised to be a potential beneficial outcome of self-generated thought^{41;11}. The second state we identified emphasises unintentional intrusive thoughts from the distant past and was associated with high levels of trait affective disturbance. This pattern of unhappy rumination may reflect the link between negative affective thought and a focus on the past that is often observed in studies of self-generated thought^{10;31;43;29}. Neurally, this state highlights an association between temporal properties of the default mode network and well-being, complementing previous findings linking DMN to both positive^{44;26} and negative^{45;46} behavioural traits. Examining both states, we also found a distinction related to the spontaneity of thought patterns associated with neural activity. Spending longer periods in the first state was

linked to less spontaneous, more deliberate thoughts about the future, whereas more time spent in the second state was associated with less deliberate, intrusive thoughts about the past. These findings are in line with studies reporting positive, constructive, future thinking^{47;48} and past related, unintentional, intrusive thoughts^{42;29}, highlighting the spontaneity of thought as an important dimension of ongoing experience⁴⁹. Notably, our study estimated states simply based on neural information without the need to repeatedly sample participants' experience. In this context, the synergy between the qualitative motifs associated with our neural states and those identified by prior experience sampling studies helps minimise concerns that experience sampling results are in general a consequence of the meta cognitive demands imposed by this technique².

Finally, our study provides a novel view of the relationships between specific macro-scale hierarchies and patterns of spontaneously occurring neurocognitive states. The two states identified that had experiential correlates fell at opposite extremes of a neurocognitive hierarchy that reflects the neural response to tasks. These data suggest that similar neural hierarchies that support general aspects of how individuals respond to the demands of an external task³² may continue to do so in the absence of any overt external behaviour. Importantly, this neural hierarchy discriminated between patterns of autobiographical planning, which are thought to be advantageous because they help consolidate personal goals¹², and those linked to patterns of rumination, which exacerbate unhappiness²⁹. Our results thus suggest that the same neural hierarchy that determines the brain's response to increasing task demands may also discriminate beneficial and detrimental types of self-generated experiences.

Although our study highlights the utility of advanced machine learning methods in the determination of naturally occurring self-generated states, it has some limitations and leaves several important questions open. For example, previous research has highlighted the potential confounding effects of cardiac and respiratory cycles on the BOLD fMRI signal^{50;51;52}. Since these physiological measures have been shown to relate to behaviour⁵³, acquiring them in future studies can provide a way to control their effect. Furthermore, several studies have shown that changes in arousal and vigilance can explain a meaningful fraction of variance in neural activity and connectivity^{54;55;56}. Given that HMMs can identify distinct neural patterns, it is possible that fluctuations in arousal may be at least partially captured by one or multiple HMM states. Vigilance and arousal, however, can have unique temporal properties, and may relate to cognition in multiple ways^{57;58;59}. Therefore, the magnitude of their effect to neural dynamics, how much can be captured by specific algorithms, different ways to account for it, and the extent that it can be distinguished from the effect of ongoing cognition are nontrivial and are left open for future research^{60;36}. In addition, our application of HMMs allowed us to estimate the occurrence of transient states at rest; however, we only gained a single measure of an individual's experience. This feature of our experimental design allows us to rule out certain meta-cognitive features of experience sampling as confounding our results (see above); however, it also makes unclear the extent to which the observed states are transient or result from more stable trait-like properties of the individual. In the future, studies could overcome this limitation by sampling an individual's experience at rest on multiple occasions. This would allow for a more quantified assessment of whether these states occur in a trait-like manner, or whether they are subject to more transient influences. Moreover, although our study highlights that neural information can be used to provide a quantified assessment of the repertoire of states that individuals engage in, it leaves open the specific experiential features that these states may include. In our study, we sampled individuals' thoughts using 25 questions that were refined through a sequence of empirical investigations, yet, it seems likely that these are not an exhaustive list of experiential states, because at least certain features of self-generated experience are probably not captured by the

specific items we used. Accordingly, our study establishes that neural information can provide an important window into the types of thoughts individuals may engage in. Nevertheless, future work is needed to establish the most appropriate self-report items to fully appreciate the full range of the ontological features of self-generated experience. Finally, our study leaves open the specific duration that states can take. Our study used fMRI to estimate the neural states and we used a repetition time (TR) that was reasonably slow (3 seconds). This feature of our design acts as a hard limit to the duration of states that our analysis can determine since we would be unable to determine states that lasted substantially shorter durations than that of the TR. In the future, it will be possible to identify states with a shorter duration by performing a similar set of analyses, using a method of neuroimaging such as electro/magnetoencephalography that can acquire neural data in a more rapid manner.

Methods and Materials

Participants

277 healthy participants were recruited from the University of York. Informed consent was obtained for all participants, the study was approved by the York Neuroimaging Centre Ethics Committee, and all research was performed in accordance with relevant guidelines and regulations. 21 participants were excluded from analyses, 1 due to technical issues during the neuroimaging data acquisition and 20 for excessive movement during the fMRI scan (mean framewise displacement⁶¹ > 0.3 mm and/or more than 15% of their data affected by motion), resulting in a final cohort of $n = 256$ (169 females, $\mu_{age} = 20.7$ years, $\sigma_{age} = 2.4$).

Behavioural methods

We sampled participants' experience during the resting state fMRI scan by asking them to retrospectively report their thoughts at the end of the scan. Experience was measured using a 4-scale Likert scale, with the question order randomised (all 25 questions are shown in Table S1). In a subset ($n=168$) of the final cohort, we also assessed their physical and mental health by administering well-established questionnaire measures at a later separate session outside of the scanner. Details about each questionnaire are presented in SI.

MRI data acquisition

MRI data were acquired on a GE 3 Tesla Signa Excite HDxMRI scanner, equipped with an eight-channel phased array head coil at York Neuroimaging Centre, University of York. For each participant, we acquired a sagittal isotropic 3D fast spoiled gradient-recalled echo T1-weighted structural scan (TR = 7.8 ms, TE = minimum full, flip angle = 20°, matrix = 256x256, voxel size = 1.13x1.13x1 mm³, FOV = 289x289 mm²). Resting-state functional MRI data based on blood oxygen level-dependent (BOLD) contrast images with fat saturation were acquired using a gradient single-shot echo-planar imaging sequence with the following parameters; TE = minimum full (≈ 19 ms), flip angle = 90°, matrix = 64x64, FOV = 192x192 mm², voxel size = 3x3x3 mm³, TR = 3000 ms, 60 axial slices with no gap and slice thickness of 3 mm. Scan duration was 9 minutes, which allowed us to collect 180 whole-brain volumes.

fMRI data pre-processing

Functional MRI data pre-processing was performed using SPM12 (<http://www.fil.ion.ucl.ac.uk/spm>) and the CONN toolbox (v.18b) (<https://www.nitrc.org/projects/conn>)⁶² implemented in Matlab (R2018a) (<https://uk.mathworks.com/products/matlab>). Pre-processing steps followed CONN's default pipeline and included motion estimation and correction by volume realignment using a six-parameter rigid body transformation, slice-time correction, and simultaneous grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) segmentation and normalisation to MNI152 stereotactic space (2 mm isotropic) of both functional and structural data. Following pre-processing, the following potential confounding effects were removed from the BOLD signal using linear regression: 6 motion parameters calculated at the previous step and their 1st and 2nd order derivatives, volumes with excessive movement (motion greater than 0.5 mm and global signal changes larger than $z = 3$), signal linear trend, and five principal components of the signal from WM and CSF (CompCor approach,⁶³). Finally, data were band-pass filtered between 0.01 and 0.1 Hz. Regarding global signal regression (GSR), studies have shown that it may remove both neuronal and non-neuronal information^{51;64}. We followed multiple pre-processing steps (described above) in an effort to minimise confounding effects and, therefore, opted to not perform GSR in our main analyses, as we were ambivalent about whether the potential benefit from including GSR in our pipeline would justify the risk of removing true neuronal signal. However, in order to mitigate the lack of acquiring physiological measures (see Discussion), and due to the considerable ongoing controversy over the matter⁶⁵, we also ran the same pipeline following GSR. The findings of this analysis, presented in supplementary Materials, were in relative agreement with our original findings.

Dimensionality reduction

In order to reduce the dimensionality structure of the neuroimaging data, we used the 17 functional network parcellation from³⁰, slightly eroded the parcels to avoid signal leakage from neighbouring parcels, masked them with subject specific grey matter masks, and calculated the average signal from the unsmoothed data within each parcel for each volume per participant. The acquired time series were then standardised and concatenated to form a "(256 participants x 180 volumes) x 17 parcels" matrix that was used as input to the HMM algorithm.

For the behavioural data, we applied a principal component analysis (PCA) with varimax rotation to the scores describing the participants' experience at the end of the resting state scan. As in many of our previous studies (i.e.^{66;18;67;68}), using this kind of decomposition and rotation improves the identification and interpretation of distinct patterns of covariance in our multi-dimensional datasets. The PCA revealed eight orthogonal components with eigenvalues greater than 1. We followed the same procedure for the trait measures of well-being that we acquired for a sub-sample of our cohort, which identified three principal components in their structure (see Fig. S1, and Tables S2 and S3 for the loadings of both decompositions, and the variance explained by each component). The component scores were used as dependent variables in two separate multivariate multiple regressions, run at a later stage in order to investigate the relationship between ongoing experience and well-being respectively, with the mean dwell-times of the dynamic neural states as described by our HMM decomposition. In order to help with the interpretation of the results, the beta weights from the regressions were mapped from the principal component space back to the multi-dimensional behavioural space and are shown as word clouds on the right-most column in Figures 3 and 4.

Hidden Markov model

To characterise the dynamics of neural activity, we applied hidden Markov modelling to the concatenated time series of the 17 parcels³⁰. The inference of the model parameters was based on variational Bayes and the minimisation of free energy, as implemented in the HMM-MAR toolbox⁶⁹. The HMM's inference assigns state probabilities to each time point of the time series (i.e. reflecting how likely each time point is to be explained by each state) and estimates the parameters of the states, where each state has its own model of the observed data. Each state was represented as a multivariate Gaussian distribution¹⁹, described by its mean and covariance. Inference was run at the group level, such that the state spatial descriptions are defined across subjects, whereas the temporal activation of the states is defined at the subject level. This allowed us to discover dynamic temporal patterns of whole-brain activity and functional interactions along with their occurrence (state time series) and transition probabilities for the duration of the whole resting state fMRI scan. Detailed information about the HMM implementation and the variational Bayes inference can be found in^{69;26;19}. As with other probabilistic unsupervised learning methods (e.g. independent component analysis), HMM is sensitive to initial values. In order to account for HMM run-to-run variability, we ran the algorithm 10 times and selected the iteration with the lowest free-energy at the end of the inference. As all 10 decompositions were the same for the 7-state solution, our results would have been robust to choosing any of the 10 iterations. This would not have been the case for the 9-state solution, as these decompositions were not as stable (see Fig. S2). The summary statistic describing the states' dwell-times was computed after hard classifying the states as being active or not by using Viterbi decoding. The mean dwell-time of each state (Fig. S5) was then used as an explanatory variable in subsequent multivariate generalised linear models. Dwell-times were standardised and any values $> 2.5\sigma$ were substituted with the mean ($\mu = 0$) to minimise the effect of potential outliers. We also calculated two additional summary metrics describing how often participants switch from one state to any other state (switching rate), and how long participants spend in each state on average (fractional occupancy), that can be seen in supplementary Figure S5. All analyses controlled for age, gender and motion during the resting state fMRI scan.

Projection of state maps to a 3-dimensional space of neural hierarchies

Following the HMM inference, each state S can be characterised by parameters (μ_s, Σ_s) , where μ_s is a vector containing the activation/weights (with respect to the average) of each of the 17 parcels used in our analyses and Σ_s is the covariance matrix describing their functional interactions. Using the weights of each parcel, we obtained each state's "mean activity" spatial map (Fig. 2) and calculated its spatial similarity to the first three gradient maps from²⁰. These maps were produced by a diffusion embedding algorithm, applied on data from the Human Connectome Project, they are orthogonal, and highlight topographies of information exchange based on the differentiation of connectivity patterns. Similarity was calculated as the pairwise correlation between each state spatial map and gradient and was used as the state's co-ordinate for the corresponding dimension in gradient space. Aiming to construct a summary metric describing the topography of the states in this space, we Fisher-z transformed the correlation values (used as gradient space co-ordinates), calculated the distance of each state from the origin weighted by its maximum co-ordinate (as a way to differentiate between points lying on the surface of the same sphere), and added the weighted distances to produce a sum for the set of all states

$$\sum_{i=1}^7 \left(\sqrt{(S_{iG1})^2 + (S_{iG2})^2 + (S_{iG3})^2} \cdot \max(S_{iG1}, S_{iG2}, S_{iG3}) \right)$$

where S_{iG1} , S_{iG2} , S_{iG3} are the Fisher's z correlation values between the spatial maps of state i and gradient 1, 2, and 3 respectively.

Null distribution

In order to compute a null distribution, we randomly permuted the parcel weights, constructed the synthetic states' spatial maps and projected them into gradient space, by calculating their similarity to the gradients in the same way as before. We ran 300 permutations, producing 2100 synthetic states and calculated the weighted distance sum of each set in a similar manner as with the set of empirical states.

Code availability

Code used for the HMM analyses can be accessed on <https://github.com/OHBA-analysis/HMM-MAR>.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments

JS was supported by European Research Council (WANDERINGMINDS - 646927). RL received support from the Wellcome/EPSRC Centre for Medical Engineering (Ref: WT 203148/Z/16/Z) and would also like to acknowledge support from the Data to Early Diagnosis and Precision Medicine Industrial Strategy Challenge Fund, UK Research and Innovation (UKRI). The authors would like to thank Mladen Sormaz, Charlotte Murphy and Hao-Ting Wang for their contribution to data acquisition.

References

- [1] Jonathan Smallwood and Jonathan W Schooler. The restless mind. *Psychological bulletin*, 132(6):946, 2006.
- [2] Jonathan Smallwood and Jonathan W Schooler. The science of mind wandering: empirically navigating the stream of consciousness. *Annual review of psychology*, 66:487–518, 2015.
- [3] Jerome L Singer. *Daydreaming: An introduction to the experimental study of inner experience*. 1966.
- [4] Jonathan Smallwood and Jessica Andrews-Hanna. Not all minds that wander are lost: the importance of a balanced perspective on the mind-wandering state. *Frontiers in psychology*, 4:441, 2013.
- [5] Benjamin W Mooneyham and Jonathan W Schooler. The costs and benefits of mind-wandering: a review. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 67(1):11, 2013.

-
- [6] Jonathan Smallwood. Distinguishing how from why the mind wanders: a process–occurrence framework for self-generated mental activity. *Psychological bulletin*, 139(3):519, 2013.
- [7] Paul Seli, Michael J Kane, Jonathan Smallwood, Daniel L Schacter, David Maillet, Jonathan W Schooler, and Daniel Smilek. Mind-wandering as a natural kind: A family-resemblances view. *Trends in cognitive sciences*, 22(6):479–490, 2018.
- [8] Kalina Christoff, Caitlin Mills, Jessica R Andrews-Hanna, Zachary C Irving, Evan Thompson, Kieran CR Fox, and Julia WY Kam. Mind-wandering as a scientific concept: cutting through the definitional haze. *Trends in cognitive sciences*, 22(11): 957–959, 2018.
- [9] Aaron Kucyi. Just a thought: how mind-wandering is represented in dynamic brain connectivity. *Neuroimage*, 180:505–514, 2018.
- [10] Florence JM Ruby, Jonathan Smallwood, Haakon Engen, and Tania Singer. How self-generated thought shapes mood—the relation between mind-wandering and mood depends on the socio-temporal content of thoughts. *PloS one*, 8(10):e77554, 2013.
- [11] David Stawarczyk, Helena Cassol, and Arnaud D’Argembeau. Phenomenology of future-oriented mind-wandering episodes. *Frontiers in Psychology*, 4:425, 2013.
- [12] Barbara Medea, Theodoros Karapanagiotidis, Mahiko Konishi, Cristina Ottaviani, Daniel Margulies, Andrea Bernasconi, Neda Bernasconi, Boris C Bernhardt, Elizabeth Jefferies, and Jonathan Smallwood. How do we decide what to do? resting-state connectivity patterns and components of self-generated thought linked to the development of more concrete personal goals. *Experimental brain research*, 236(9):2469–2481, 2018.
- [13] Jessica R Andrews-Hanna, Roselinde Henderson Kaiser, Amy EJ Turner, Andrew Reineberg, Detre Godinez, Sona Dimidjian, and Marie Banich. A penny for your thoughts: dimensions of self-generated thought content and relationships with individual differences in emotional wellbeing. *Frontiers in psychology*, 4:900, 2013.
- [14] Anthony P Zanesco. Quantifying streams of thought during cognitive task performance using sequence analysis. *Behavior Research Methods*, pages 1–21, 2020.
- [15] Jonathan Smallwood, Louise Nind, and Rory C O’Connor. When is your head at? an exploration of the factors associated with the temporal focus of the wandering mind. *Consciousness and cognition*, 18(1):118–125, 2009.
- [16] Krzysztof J Gorgolewski, Dan Lurie, Sebastian Urchs, Judy A Kipping, R Cameron Craddock, Michael P Milham, Daniel S Margulies, and Jonathan Smallwood. A correspondence between individual differences in the brain’s intrinsic functional architecture and the content and form of self-generated thoughts. *PloS one*, 9(5): e97176, 2014.
- [17] David Maillet, Roger E Beaty, Megan L Jordano, Dayna R Touron, Areeba Adnan, Paul J Silvia, Thomas R Kwapil, Gary R Turner, R Nathan Spreng, and Michael J Kane. Age-related differences in mind-wandering in daily life. *Psychology and aging*, 33(4):643, 2018.
- [18] Deniz Vatansever, Theodoros Karapanagiotidis, Daniel S Margulies, Elizabeth Jefferies, and Jonathan Smallwood. Distinct patterns of thought mediate the link between brain functional connectomes and well-being. *Network Neuroscience*, pages 1–21, 2020.
-

-
- [19] Diego Vidaurre, Romesh Abeyesuriya, Robert Becker, Andrew J Quinn, Fidel Alfaro-Almagro, Stephen M Smith, and Mark W Woolrich. Discovering dynamic brain networks from big data in rest and task. *Neuroimage*, 2017.
- [20] Daniel S Margulies, Satrajit S Ghosh, Alexandros Goulas, Marcel Falkiewicz, Julia M Huntenburg, Georg Langs, Gleb Bezgin, Simon B Eickhoff, F Xavier Castellanos, Michael Petrides, et al. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, page 201608282, 2016.
- [21] Bertha Vázquez-Rodríguez, Laura E Suárez, Ross D Markello, Golia Shafiei, Casey Paquola, Patric Hagmann, Martijn P van den Heuvel, Boris C Bernhardt, R Nathan Spreng, and Bratislav Misic. Gradients of structure–function tethering across neocortex. *Proceedings of the National Academy of Sciences*, 116(42): 21219–21227, 2019.
- [22] Casey Paquola, Reinder Vos De Wael, Konrad Wagstyl, Richard AI Bethlehem, Seok-Jun Hong, Jakob Seidlitz, Edward T Bullmore, Alan C Evans, Bratislav Misic, Daniel S Margulies, et al. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLoS biology*, 17(5):e3000284, 2019.
- [23] Seok-Jun Hong, Reinder Vos de Wael, Richard AI Bethlehem, Sara Lariviere, Casey Paquola, Sofie L Valk, Michael P Milham, Adriana Di Martino, Daniel S Margulies, Jonathan Smallwood, et al. Atypical functional connectome hierarchy in autism. *Nature communications*, 10(1):1022, 2019.
- [24] Brontë Mckeown, Will H Strawson, Hao-Ting Wang, Theodoros Karapanagiotidis, Reinder Vos de Wael, Oualid Benkarim, Adam Turnbull, Daniel Margulies, Elizabeth Jefferies, Cade McCall, et al. The relationship between individual variation in macroscale functional gradients and distinct aspects of ongoing thought. *NeuroImage*, page 117072, 2020.
- [25] Adam Turnbull, Theodoros Karapanagiotidis, Hao-Ting Wang, Boris C Bernhardt, Robert Leech, Daniel Margulies, Jonathan Schooler, Elizabeth Jefferies, and Jonathan Smallwood. Reductions in task positive neural systems occur with the passage of time and are associated with changes in ongoing thought. *Scientific reports*, 10(1):1–10, 2020.
- [26] Diego Vidaurre, Stephen M Smith, and Mark W Woolrich. Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):12827–12832, 2017.
- [27] Tal Yarkoni, Russell A Poldrack, Thomas E Nichols, David C Van Essen, and Tor D Wager. Large-scale automated synthesis of human functional neuroimaging data. *Nature methods*, 8(8):665, 2011.
- [28] Jonathan Smallwood, Theodoros Karapanagiotidis, Florence Ruby, Barbara Medea, Irene de Caso, Mahiko Konishi, Hao-Ting Wang, Glyn Hallam, Daniel S Margulies, and Elizabeth Jefferies. Representing representation: Integration between the temporal lobe and the posterior cingulate influences the content and form of spontaneous thought. *PloS one*, 11(4):e0152272, 2016.
- [29] Giulia L Poerio, Peter Totterdell, and Eleanor Miles. Mind-wandering and negative mood: Does one thing really lead to another? *Consciousness and cognition*, 22(4): 1412–1421, 2013.

-
- [30] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3): 1125, 2011.
- [31] Jonathan Smallwood and Rory C O'Connor. Imprisoned by the past: unhappy moods lead to a retrospective bias to mind wandering. *Cognition & emotion*, 25(8): 1481–1490, 2011.
- [32] John Duncan. The multiple-demand (md) system of the primate brain: mental programs for intelligent behaviour. *Trends in cognitive sciences*, 14(4):172–179, 2010.
- [33] Jonathan D Power and Steven E Petersen. Control-related systems in the human brain. *Current opinion in neurobiology*, 23(2):223–228, 2013.
- [34] Marcus E Raichle, Ann Mary MacLeod, Abraham Z Snyder, William J Powers, Debra A Gusnard, and Gordon L Shulman. A default mode of brain function. *Proceedings of the National Academy of Sciences*, 98(2):676–682, 2001.
- [35] Theodoros Karapanagiotidis, Elizabeth Jefferies, and Jonathan Smallwood. Population level interactions between the neural correlates of dispositional internal thought and visual imagery. *bioRxiv*, page 857490, 2019.
- [36] Daniel J Lurie, Daniel Kessler, Danielle S Bassett, Richard F Betzel, Michael Breakspear, Shella Kheilholz, Aaron Kucyi, Raphaël Liégeois, Martin A Lindquist, Anthony Randal McIntosh, et al. Questions and controversies in the study of time-varying functional connectivity in resting fmri. *Network Neuroscience*, 4(1): 30–69, 2020.
- [37] Teresa McCormack, Patrick Burns, Patrick O'Connor, Agnieszka Jaroslawska, and Eugene M Caruso. Do children and adolescents have a future-oriented bias? a developmental study of spontaneous and cued past and future thinking. *Psychological research*, 83(4):774–787, 2019.
- [38] Muireann Irish, Zoë-lee Goldberg, Sara Alaeddin, Claire O'Callaghan, and Jessica R Andrews-Hanna. Age-related changes in the temporal focus and self-referential content of spontaneous cognition during periods of low cognitive demand. *Psychological research*, 83(4):747–760, 2019.
- [39] Xiaolan Song and Xiao Wang. Mind wandering in chinese daily lives—an experience sampling study. *PloS one*, 7(9):e44423, 2012.
- [40] Yudai Iijima and Yoshihiko Tanno. The effect of cognitive load on the temporal focus of mind wandering. *Shinrigaku kenkyu: The Japanese journal of psychology*, 83(3):232–236, 2012.
- [41] Benjamin Baird, Jonathan Smallwood, and Jonathan W Schooler. Back to the future: autobiographical planning and the functionality of mind-wandering. *Consciousness and cognition*, 20(4):1604–1611, 2011.
- [42] Paul Seli, Brandon CW Ralph, Mahiko Konishi, Daniel Smilek, and Daniel L Schacter. What did you have in mind? examining the content of intentional and unintentional types of mind wandering. *Consciousness and cognition*, 51:149–156, 2017.

-
- [43] Maitta Spronken, Rob W Holland, Bernd Figner, and Ap Dijksterhuis. Temporal focus, temporal distance, and mind-wandering valence: Results from an experience sampling and an experimental study. *Consciousness and Cognition*, 41:104–118, 2016.
- [44] Stephen M Smith, Thomas E Nichols, Diego Vidaurre, Anderson M Winkler, Timothy EJ Behrens, Matthew F Glasser, Kamil Ugurbil, Deanna M Barch, David C Van Essen, and Karla L Miller. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature neuroscience*, 18(11):1565–1567, 2015.
- [45] Hui-Xia Zhou, Xiao Chen, Yang-Qian Shen, Le Li, Ning-Xuan Chen, Zhi-Chen Zhu, Francisco Xavier Castellanos, and Chao-Gan Yan. Rumination and the default mode network: Meta-analysis of brain imaging studies and implications for depression. *Neuroimage*, 206:116287, 2020.
- [46] J Paul Hamilton, Daniella J Furman, Catie Chang, Moriah E Thomason, Emily Dennis, and Ian H Gotlib. Default-mode and task-positive network activity in major depressive disorder: implications for adaptive and maladaptive rumination. *Biological psychiatry*, 70(4):327–333, 2011.
- [47] George J Huba, Carol S Aneshensel, and Jerome L Singer. Development of scales for three second-order factors of inner experience. *Multivariate Behavioral Research*, 16(2):181–206, 1981.
- [48] Rebecca McMillan, Scott Barry Kaufman, and Jerome L Singer. Ode to positive constructive daydreaming. *Frontiers in psychology*, 4:626, 2013.
- [49] Paul Seli, Jonathan SA Carriere, and Daniel Smilek. Not all mind wandering is created equal: Dissociating deliberate from spontaneous mind wandering. *Psychological Research*, 79(5):750–758, 2015.
- [50] Kevin Murphy, Rasmus M Birn, and Peter A Bandettini. Resting-state fmri confounds and cleanup. *Neuroimage*, 80:349–359, 2013.
- [51] Jonathan D Power, Mark Plitt, Timothy O Laumann, and Alex Martin. Sources and implications of whole-brain fmri signals in humans. *Neuroimage*, 146:609–625, 2017.
- [52] César Caballero-Gaudes and Richard C Reynolds. Methods for cleaning the bold fmri signal. *Neuroimage*, 154:128–149, 2017.
- [53] Joshua S Siegel, Anish Mitra, Timothy O Laumann, Benjamin A Seitzman, Marcus Raichle, Maurizio Corbetta, and Abraham Z Snyder. Data quality influences observed links between functional connectivity and behavior. *Cerebral cortex*, 27(9):4492–4502, 2017.
- [54] Chenhao Wang, Ju Lynn Ong, Amiya Patanaik, Juan Zhou, and Michael WL Chee. Spontaneous eyelid closures link vigilance fluctuation with fmri dynamic connectivity states. *Proceedings of the National Academy of Sciences*, 113(34):9653–9658, 2016.
- [55] Catie Chang, David A Leopold, Marieke Louise Schölvinck, Hendrik Mandelkow, Dante Picchioni, Xiao Liu, Q Ye Frank, Janita N Turchi, and Jeff H Duyn. Tracking brain arousal fluctuations with fmri. *Proceedings of the National Academy of Sciences*, 113(16):4518–4523, 2016.
-

-
- [56] Garth John Thompson, Matthew Evan Magnuson, Michael Donelyn Merritt, Hillary Schwarb, Wen-Ju Pan, Andrew McKinley, Lloyd D Tripp, Eric H Schumacher, and Shella Dawn Keilholz. Short-time windows of correlation between large-scale functional brain networks predict vigilance intraindividually and interindividually. *Human brain mapping*, 34(12):3280–3298, 2013.
- [57] Enzo Tagliazucchi and Helmut Laufs. Decoding wakefulness levels from typical fmri resting-state data reveals reliable drifts between wakefulness and sleep. *Neuron*, 82(3):695–708, 2014.
- [58] Anish Mitra, Abraham Z Snyder, Enzo Tagliazucchi, Helmut Laufs, and Marcus E Raichle. Propagated infra-slow intrinsic brain activity reorganizes across wake and slow wave sleep. *Elife*, 4:e10781, 2015.
- [59] Susan J Sara and Sebastien Bouret. Orienting and reorienting: the locus coeruleus mediates cognition through arousal. *Neuron*, 76(1):130–141, 2012.
- [60] Aaron Kucyi, Arielle Tambini, Sepideh Sadaghiani, Shella Keilholz, and Jessica R Cohen. Spontaneous cognitive processes and the behavioral validation of time-varying brain connectivity. *Network Neuroscience*, 2(4):397–417, 2018.
- [61] Jonathan D Power, Anish Mitra, Timothy O Laumann, Abraham Z Snyder, Bradley L Schlaggar, and Steven E Petersen. Methods to detect, characterize, and remove motion artifact in resting state fmri. *Neuroimage*, 84:320–341, 2014.
- [62] Susan Whitfield-Gabrieli and Alfonso Nieto-Castanon. Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain connectivity*, 2(3):125–141, 2012.
- [63] Yashar Behzadi, Khaled Restom, Joy Liau, and Thomas T Liu. A component based noise correction method (compcor) for bold and perfusion based fmri. *Neuroimage*, 37(1):90–101, 2007.
- [64] Lucina Q Uddin. Mixed signals: on separating brain signal from noise. *Trends in cognitive sciences*, 21(6):405–406, 2017.
- [65] Thomas T Liu, Alican Nalci, and Maryam Falahpour. The global signal in fmri: Nuisance or information? *NeuroImage*, 150:213–229, 2017.
- [66] Delali Konu, Adam Turnbull, Theo Karapanagiotidis, Hao-ting Wang, Lydia Brown, Elizabeth Jefferies, and Jonathan Smallwood. A role for ventromedial prefrontal cortex in self-generated episodic social cognition. *NeuroImage*, page 116977, 2020.
- [67] A Turnbull, HT Wang, C Murphy, NSP Ho, X Wang, M Sormaz, T Karapanagiotidis, RM Leech, B Bernhardt, DS Margulies, et al. Left dorsolateral prefrontal cortex supports context-dependent prioritisation of off-task thought. *Nature communications*, 10(1):1–10, 2019.
- [68] Theodoros Karapanagiotidis, Boris C Bernhardt, Elizabeth Jefferies, and Jonathan Smallwood. Tracking thoughts: Exploring the neural architecture of mental time travel during mind-wandering. *Neuroimage*, 147:272–281, 2017.
- [69] Diego Vidaurre, Andrew J Quinn, Adam P Baker, David Dupret, Alvaro Tejero-Cantero, and Mark W Woolrich. Spectrally resolved fast transient brain states in electrophysiological data. *Neuroimage*, 126:81–95, 2016.