

Towards an Ontology of Ongoing Thought

Hao-Ting Wang

PhD

University of York

Psychology

September 2018

Abstract

Functional outcomes of ongoing thought show both costs and benefits. Yet, the reason for its heterogeneity remains unclear. The executive failure and representational accounts stemmed from different psychological research approaches to understand ongoing thought. The executive failure account examines why changes in ongoing thought happen, while the representational account seeks to explain how humans generate ongoing thought. The attentional system and the default mode network are the common neural processes of both theoretical accounts, but interacting in a contradicting manner. The two accounts can be seen as competing theories of ongoing thought. However, in the family resemblance view (Seli et al., 2018), the two theoretical accounts potentially serve as two component processes of one phenomenon. One possible solution to this conflict could be that under different global neural configurations, the two networks support different cognitive functions. The thesis sets out to present evidence supporting of the family resemblance view and to begin research on the ontology of the component processes in ongoing thought. Neural cognitive hierarchy is the potential explanation of the heterogeneity. The current thesis adopts sparse canonical correlation analysis to incorporate the neural and behavioural aspects of ongoing thought. The data suggests ongoing thought is a collective phenomenon with many types of experience driven by the connectivity patterns in the default mode network. Each type of experience associated with their unique functional outcomes and neural hierarchies at the whole-brain level. Cognitive flexibility and the balance of segregation and integration between the transmodal systems and the rest of the cortex determines the immersive details. The current findings suggested the importance of whole-brain neural hierarchies to ongoing thought. The confirmation of these trait level findings at a state level are necessary to gain more insights into the architecture of the component processes.

Contents

Abstract	ii
Contents	iii
List of Figures	vii
List of Tables	ix
Acknowledgements	xi
Author's Declaration	xii
1 Introduction	1
1.1 Heterogeneity of ongoing thought	1
1.1.1 Heterogeneity in definitions	2
1.1.2 Heterogeneity in functional outcomes	3
1.1.3 Heterogeneity in experiential profiles	4
1.2 Theoretical accounts of mind-wandering	5
1.2.1 Executive failure account	5
1.2.2 Representational account	7
1.3 Neural hierarchies	8
1.3.1 Historical perspective	8
1.3.2 Abstract rule governing	9
1.3.3 Sensory integration/segregation	10
1.3.4 Integration of hierarchical configurations	11
1.4 The current thesis	12
1.4.1 Method of acquiring experience	14
1.4.2 Content of experience	15

1.4.3	Conjoined decomposition of brain and cognition	16
1.5	Summary and thesis outline	17
2	Canonical Correlation Analysis	21
2.1	Abstract	21
2.2	Motivation	22
2.3	Modelling intuitions	23
2.3.1	Joint information compression	25
2.3.2	Symmetry	25
2.3.3	Multiplicity	26
2.3.4	Interim summary	26
2.4	Examples	27
2.5	Interpretations	30
2.5.1	Utility of CCA	31
2.5.2	Relation to other commonly used methods	33
2.6	Practical considerations	35
2.6.1	Preprocessing	35
2.6.2	Model selection	36
2.7	Summary	37
3	Exploring the Heterogeneity of Ongoing Thought	39
3.1	Abstract	39
3.2	Introduction	40
3.3	Method	43
3.3.1	Participants	43
3.3.2	MRI acquisition	44
3.3.3	Questionnaires	44
3.3.4	Behavioural testing sessions	44
3.3.5	Neuroimaging data preprocessing and analysis	47
3.4	Results	53
3.4.1	Determining consistent categories of experience	53
3.4.2	Validating the categories of experience	55
3.5	Discussion	60

4	Population Variation in the Associations Between Large-Scale Networks and Experiences at Rest	64
4.1	Abstract	64
4.2	Introduction	65
4.3	Method	69
4.3.1	Participants	69
4.3.2	Cognitive measures and questionnaires	69
4.3.3	ongoing cognition measure	70
4.3.4	MR data processing	70
4.3.5	Conjoint decomposition of connectivity and experience . .	72
4.3.6	Test of component robustness	75
4.3.7	Code availability	76
4.4	Results	76
4.4.1	Determining constituent category of experience	76
4.4.2	The relationship between neuroexperiential components and cognitive functions	79
4.5	Discussion	81
5	Inhibition of Prior Information Contributes to Internal Content Representation	87
5.1	Abstract	87
5.2	Introduction	88
5.3	Method	89
5.3.1	Participants	89
5.3.2	MRI acquisition	90
5.3.3	Resting state data preprocessing	90
5.3.4	ROI-ROI functional connectivity.	91
5.3.5	Behavioural data	91
5.3.6	Multivariate pattern analysis	93
5.3.7	Group analysis	95
5.4	Results	96
5.4.1	Dimensions of ongoing thought	96
5.4.2	Neurocognitive component selection	96
5.4.3	Determining constituent category of cognitive functions .	97

5.4.4	The relationship between neurocognitive components and self-reports on thoughts	100
5.4.5	The relationship between neurocognitive components and ongoing thought patterns	101
5.5	Discussion	102
6	General discussion	106
6.1	Empirical findings	107
6.1.1	Heterogeneity	107
6.1.2	Integration and segregation in transmodal networks	110
6.2	Limitations	112
6.3	Future directions	114
6.4	Concluding remarks	114
	Appendix	116
A	Chapter 3 Supplemental Materials	116
A.1	Questionnaires	116
A.1.1	Health Organization Adult ADHD Self-Report Scale	116
A.1.2	Autism Spectrum Quotient	116
A.1.3	British Dyslexia Association Dyslexia checklist	117
A.1.4	World Health Organization Quality of Life	117
A.1.5	CES-Depression scale	117
A.1.6	State-Trait Anxiety Inventory	117
A.1.7	Ruminative Response Scale	118
A.2	Cognitive tasks	118
A.2.1	General apparatus of the laboratory session	118
A.2.2	Semantic tasks	119
A.2.3	Fluency Task	120
A.2.4	Word pair memory task	120
A.2.5	Digit span	121
A.2.6	Flanker task	121
A.2.7	Task-switching task	122
A.2.8	Four mountains task	122
A.2.9	Ravens advanced progressive matrices	122

A.2.10 Unusual uses task	123
A.3 Supplementary analysis and figures	124
B Nested K-Fold Cross-Validation	126
Bibliography	127

List of Figures

1.1	Schematic of the thesis	13
2.1	An example of CCA on behavioural data.	25
2.2	The analysis pipeline of Smith et al. (2015). The arrows represent analysis performed.	27
2.3	The analysis pipeline of Wang, Poerio et al. (2018). The arrows represent analysis performed.	28
2.4	A flowchart illustrating the choices when considering the application of CCA to a dataset.	32
3.1	Schematic of the procedure and analysis strategy employed in the current study.	42
3.2	Results of the sparse canonical-correlation analysis.	49
3.3	Results from principal component analyses of (a) behavioural tasks and (b) questionnaires.	54
3.4	Relationship between the different neural-cognitive components and the laboratory and questionnaire measures.	57
4.1	A diagram of the nested k-fold cross-validation with model selection.	74
4.2	Unique neurocognitive dimensions of population variation revealed by sparse canonical correlation analysis of measures of whole brain connectivity and self-reported descriptions of ongoing experience.	77
4.3	The relationship between the different neural-cognitive components and the measures assessed in the cognitive battery.	80
4.4	The principle component and its relationship to the different neurocognitive components.	80
5.1	Analysis pipeline.	95

5.2	Dimensions of ongoing thought.	96
5.3	Grid search result.	97
5.4	Significant components from SCCA.	98
5.5	Group level analysis on neurocognitive components and self-report on thoughts.	101
5.6	The relationships between the neurocognitive components and the ongoing thought.	102
A.1	Restricted temporal sampling and bootstrapping resampling dis- tribution with 1000 iteration.	124
A.2	Scree plots of the principle component analysis.	124
A.3	Full set of components.	125
A.4	Decomposition with motion outlier subjects excluded.	125

List of Tables

3.1	Experience-Sampling Questions in the 0-Back/1-Back Task. . . .	47
4.1	The New York Cognition Questionnaire (NYC-Q).	71
A.1	Correlation between motion parameter (Mean FD Jenkinson) and variable of interests	124

Acknowledgements

The most exciting chapter of my academic journey would not be possible without my supervisors Prof. Jonathan Smallwood and Prof. Elizabeth Jefferies. Jonny and Beth are the kindest people I have ever known. Especially, I cannot be more thankful for Jonny's trust in me when I was doubtful of my own ability. That offer really changed my life. Their help and support in both science and life have made this PhD a truly rewarding experience. From them, I learned to be a better scientist and also a good person.

I express my gratitude to Prof. Dr. Danilo Bzdok for the guidance and discussions on the canonical correlation analysis project. Our discussions at various Brainhacks helped tremendously in the research methods of this PhD.

I thank my thesis advisory panel, Dr. Tom Hartley, Dr. Aidan Horner, and Dr. Cade McCall for providing a friendly environment to discuss science and my skill development.

All the research projects presented here will not be possible without the current and past members of Beth and Jonny's lab. Good colleagues are difficult to come by. Their contributions lie not just in the research, and also moral supports on my personal life.

Finally, thanks to all my family and friends for all their assistance. The special mentions go to Thomas Hardman and Rebecca Jones for making my life more fun in general.

Author's Declaration

I, Hao-Ting Wang, declare that this thesis is a presentation of original work and I am the sole author. I undertook the research at University of York during 2015 – 2018, under the joint supervision of Professor Jonathan Smallwood and Professor Elizabeth Jefferies. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Some parts of this thesis have been published in peer-reviewed journals or is currently under preparation for publication. Author contributions are noted at the start of each chapter.

- Chapter 2: Wang, H.-T., Smallwood, J., & Bzdok, D. (2018). Finding the needle in high dimensions: A tutorial on CCA in biomedicine. Manuscript preparing for publication.
- Chapter 3: Wang, H.-T., Poerio, G. L., Murphy, C. E., Bzdok, D., Jefferies, E., & Smallwood, J.(2018). Dimensions of Experience: Exploring the Heterogeneity of the Wandering Mind. *Psychological Science*, 29 (1), 56–71. doi: 10.1177/0956797617728727
- Chapter 4: Wang, H.-T., Bzdok, D., Margulies, D. S., Craddock, C., Milham, M., Jefferies, E., & Smallwood, J.(2018). Patterns of thought: Population variation in the associations between large-scale network organisation and self-reported experiences at rest. *NeuroImage*, 176 (1), 518–527. doi: 10.1016/j.neuroimage.2018.04.064

Copyright © 2018 by Hao-Ting Wang

The copyright of this thesis rests with the author. Any quotations from it should be acknowledged appropriately.

Chapter 1

Introduction

With the advance of functional magnetic resonance imaging (fMRI) and other neuroimaging techniques, the study of ongoing thought has gained wider interest in psychology and neuroscience in the past decade. The increasing research has given rise to heterogeneous views on how and why mind-wandering occurs, however, the detailed neural basis of the heterogeneity remains largely unclear. The current thesis explores the patterns of ongoing thought extending from well-studied off-task thought—mind-wandering—to task-related thought. The aim is to gain an understanding of the component processes of ongoing thought at task and rest. A flexible family resemblance view (Seli et al., 2018) of ongoing thought will jointly identify patterns of the shared similarities, as well as unique features that drive the heterogeneity.

In this chapter, I walk through the conflicting behavioural literature of mind-wandering and discuss the theoretical accounts of ongoing thought. Next, I introduce the emerging neural evidence on the hierarchical organisation of cognition (Margulies et al., 2016; Mesulam, 1998) and argue how overlapping component processes might facilitate ongoing thought. In closing, I introduce the overall methods used in the current research and how a multivariate approach helps to examine the family resemblance view of ongoing thought.

1.1 Heterogeneity of ongoing thought

Mind-wandering is particularly well-studied among all phenomena related to ongoing thoughts. Researchers aim to understand how the mind shifts between the

external environment and internal thoughts unrelated to the here-and-now. The executive failure account is concerned with why some mind-wandering episodes occur to the detriment of the integrity of an ongoing task (McVay & Kane, 2010). In contrast to the executive failure view, the representational account seeks an understanding of how the mental content is generated (Smallwood et al., 2016). The two approaches have led to conflict in the mind-wandering literature. Seli et al. (2018) has recently proposed a family resemblance view to incorporate different theoretical accounts under a common framework. In a family resemblance view, members of the ongoing thought family can have shared similarities along with unique features resulting in heterogeneity. The following section will introduce the evidence for construct overlap and shared processes of ongoing thought at task and rest.

1.1.1 Heterogeneity in definitions

Among types of ongoing thought, mind-wandering attracted the most interest as it concerns the ability to focus on the task at hand. Mind-wandering has been studied in a variety of related psychological domains, such as cognition, emotion, and neuroscience. Various lines of research have addressed the basic phenomenal characteristics of mind-wandering—

a shift in the contents of thought away from an ongoing task and/or from events in the external environment to self-generated thoughts and feelings.

(Smallwood & Schooler, 2006, 2015).

We are all familiar with moments when the train of thought shift away from the tasks at hand, and sometimes get annoyed by the mind-wandering episode. This intuition has led to researches describing mind-wandering as an ‘attentional lapse’ (McVay & Kane, 2009, 2012a), implying the occurrence of mind-wandering is an unintended failure. When a study explicitly instructs the participant to perform a task, the time not focused on the task is considered as ‘mind-wandering’. Such research designs dismissed the possibility of voluntary engagement in the mind-wandering state.

Recent investigations have found that mind-wandering can occur with or without intention (see the review from Seli, Risko, Smilek, & Schacter, 2016).

The participant can intentionally mind wander if they lack a motivation to engage in the experiment. When a simple yes/no question is asked about the mind-wandering state, the response cannot access the nature of the occurrence. When participants are asked about the nature of mind-wandering periods in a laboratory scenario, less than half of the mind-wandering is intentional (Seli, Cheyne, Xu, Purdon, & Smilek, 2015) due to the lack of motivation to complete the task (Seli, Cheyne, et al., 2015), or the task is not mentally demanding enough to have all attentional resources allocated to the task (Seli, Risko, & Smilek, 2016). The occurrence of intended and unintended mind-wandering can also be down to individual differences. Intentional and unintentional mind-wandering have been found to be differentially associated with attention-deficit/hyperactivity disorder (ADHD; Seli, Smallwood, Cheyne, & Smilek, 2015) and obsessive-compulsive disorder (OCD; Seli, Risko, Purdon, & Smilek, 2017). The work on the intention of mind-wandering demonstrates that there is indeed overlap in the various definitions and other components that contribute to the heterogeneity.

1.1.2 Heterogeneity in functional outcomes

The family resemblance view suggests that complex thought can emerge from the combination of multiple overlapping processes. A myriad of mind-wandering research concerns functional outcomes. The current thesis proposes that the heterogeneous functional outcomes are evidence in support of a variety of component processes underlying ongoing thought.

Mind-wandering has been associated with poor executive control during working memory tasks (McVay & Kane, 2009). Individuals who mind-wandered more during fluid intelligent testing perform less well (Mrazek et al., 2012). Mind-wandering leads to bad reading comprehension due to failure in the construction of the mental models of ongoing events (Smallwood, McSpadden, & Schooler, 2008). Comprehension ability is related to working memory capacity and mediated by the ability to suppress mind-wandering (McVay & Kane, 2012b; Unsworth & McMillan, 2013). Mind-wandering has been linked to unhappiness (Killingsworth & Gilbert, 2010) and is an indicator of depression (Smallwood, O'Connor, Sudbery, & Obonsawin, 2007). The evidence above supports the highly disruptive nature of mind-wandering and its potential costs to cognitive

performance.

In addition to exploring the costs of mind-wandering, researchers have discovered its potential benefits. Mind-wandering may facilitate a creative solution to an old problem (Baird et al., 2012; Smeekens & Kane, 2016) and recovery from negative emotional states (Ruby, Smallwood, Engen, & Singer, 2013; Poerio, Totterdell, Emerson, & Miles, 2016). Mind-wandering relies on mental time travel—the mental capacity of remembering the past and imagining the future (Stawarczyk & D’Argembeau, 2015)—and relies on neural mechanisms associated with the memory function (D’Argembeau & Van der Linden, 2006; D’Argembeau, Jeunehomme, Majerus, Bastin, & Salmon, 2015). Mind-wandering can also refine personal goals (Medea et al., 2016), potentially through mental time travel.

In summary, different functional associations arise from the same type of experience—mind-wandering. To reconcile this contradictory evidence, researchers have suggested that mind-wandering may encompass multiple states with differential contents and underlying cognitive architectures (Smallwood & Andrews-Hanna, 2013). Complex thought can emerge from the combination of multiple overlapping processes.

1.1.3 Heterogeneity in experiential profiles

Self-report is commonly used to understand the content of mind-wandering thoughts and ongoing experience. The content of mind-wandering covers a wide variety of topics and modalities. The questions in the report address a number of dimensions of the ongoing experience, ranging from the state of attention, temporal content, social content and modality (i.e. thinking in words or images). Principle component analysis (PCA) formalised the statistically shared association between different aspects of the reports.

Studies using PCA on such experience reports have revealed detailed experiential profiles. Temporal information is one common theme (Ruby, Smallwood, Sackur, & Singer, 2013; Ruby, Smallwood, Engen, & Singer, 2013). The content of mind-wandering is mainly future-focused (Baird, Smallwood, & Schooler, 2011), therefore mind-wandering often involves planning for the future goals of the individual. On the contrary, when the mind wanders in an unhappy mood,

the content is drawn to events from its past (Smallwood & O'Connor, 2011). The form of spontaneous thoughts is likely to be imagery or verbal (Gorgolewski et al., 2014; Smallwood et al., 2016).

In conclusion, the positive/negative-valence of the emotion of thought has the tendency to accompany with different temporal directions. These unique associations discovered through PCA suggested component processes at an experiential level. Investigation in experiential profiles is the first step to explore the commonality of various type of ongoing thought.

1.2 Theoretical accounts of mind-wandering

The heterogeneity of mind-wandering has been formalised into two theoretical accounts. The executive failure account aims to understand the conditions that trigger or associate with ongoing thought (Kane & McVay, 2012; McVay & Kane, 2010). The representational account examines the mechanisms that give rise to different patterns of ongoing thought (Smallwood et al., 2016). In other words, the executive failure account examines why changes in ongoing thought happen, while the representational account seeks to explain how humans generate ongoing thought while mind-wandering. The two accounts can be seen as competing theories of mind-wandering. However, in a family resemblance view, the two theoretical accounts potentially support a singular phenomenon that is composed of multiple underlying component processes.

1.2.1 Executive failure account

The executive failure account is concerned with a single aspect of mind-wandering, namely understanding why some mind-wandering episodes occur to the detriment of the integrity ongoing task. Mind-wandering occurs during attention-demanding tasks when control processes are insufficient to deal with the interference created by off-task thoughts (Kane & McVay, 2012; McVay & Kane, 2010). Under this view, mind-wandering results from a failure of attention to external tasks, rather than from the consumption of executive resources by internally generated thoughts. This research focuses on the negative effect of mind-wandering on the development of negative mood and task performance. Mind-wandering

thoughts are mostly unhappy in ecologically valid scenario (Killingsworth & Gilbert, 2010). Depressive thinking correlates with the frequency of mind-wandering (Smallwood et al., 2007).

In the executive failure view, mind-wandering reflects the momentary lapse in attention. The definition of attention lapse is a relatively slow response time to the task at hand, which is consistent with the mind-wandering indicator used in working memory capacity research (McVay & Kane, 2012a). Mind-wandering has been considered the consequence of poor executive control during working memory task (McVay & Kane, 2009). Executive deficits and slow reaction times correlate with individual differences in working memory capacity and mind-wandering (McVay & Kane, 2012a). The capacity to avoid mind-wandering during demanding tasks is a potentially important source of success on measures of fluid intelligence (Mrazek et al., 2012).

Task-based fMRI studies of attentional lapses have contributed to the functional neural processes to support the executive failure account. Weissman, Roberts, Visscher, and Woldorff (2006) have described the neural mechanism associated with attentional lapses during a global/local selective-attention task. Brief attentional lapses are related to early activity in frontal control region including anterior cingulate cortex (ACC), right middle frontal gyrus (MFG), and right inferior frontal gyrus (IFG). Attentional lapses also suggest a failure to maintain perceptual representations. Reduced activity is found in the primary visual area. Activation of the default mode network (DMN; Raichle et al., 2001; Shulman et al., 1997) has also been observed during a brief attention lapse. DMN is a set of brain regions composed of the medial prefrontal cortex (MPFC), posterior cingulate cortex (PCC) and angular gyrus (AG) as the core, plus subsystems within medial and lateral temporal lobe (Andrews-Hanna, Reidler, Sepulcre, Poulin, & Buckner, 2010). DMN is commonly referred as a task-negative network (M. D. Fox et al., 2005), associating with task-unrelated thought and mind-wandering (Mason et al., 2007; Christoff, Gordon, Smallwood, Smith, & Schooler, 2009). The lapses increase demands on the frontal-parietal control network to redirect attention. Ventral frontal-parietal regions, including right temporal-parietal junction (TPJ) and right IFG, respond during recovery from lapses (Corbetta & Shulman, 2002; Weissman et al., 2006).

1.2.2 Representational account

The representational account concerns the generation of content during mind-wandering, suggesting that mind-wandering is not merely a mindless state. The ability to generate information without task constraints are consistent with the productive functional outcomes of mind-wandering, such as creativity (Baird et al., 2012; Smeekens & Kane, 2016) and social-temporal problem-solving (Ruby, Smallwood, Engen, & Singer, 2013; Poerio et al., 2016; Medea et al., 2016).

The internal representation of semantics and episodic memory is associated with brain regions highly overlapping with DMN. The brain regions involved in semantic processing include left AG, lateral and ventral temporal cortex, left dorsal MPFC, left IFG, left ventral MPFC, and PCC (Binder, Desai, Graves, & Conant, 2009; Lambon-Ralph et al., 2017). Dorsal and ventral MPFC show high activity at rest and are associated with personally relevant information (Gusnard, Akbudak, Shulman, & Raichle, 2001). Studies of spontaneous thought suggest that PCC is an integrational hub of information from medial and lateral temporal lobes (Smallwood et al., 2016). Integration of the hippocampus with the DMN facilitates mental time travel (Karapanagiotidis, Bernhardt, Jefferies, & Smallwood, 2017). This representational process is not unique to mind-wandering. Vatansever, Menon, and Stamatakis (2017) demonstrated that the application of a newly acquired rule is associated with memory representation related brain regions such as hippocampus and PCC.

DMN demonstrates both integrative and segregating pattern with the sensory systems to support the representational process, such as semantic processing (Binder et al., 2009; Krieger-Redwood et al., 2016), episodic recollection (Rugg & Vilberg, 2013), mental time travel (Schacter, Addis, & Buckner, 2007). A common requirement of these cognitive processes is the focus on previously-encoded knowledge, as opposed to information in the external environment. The integrative and segregating modes of DMN are closely allied to perception-decoupling and conceptually-guided cognition (Murphy et al., 2018). The ability of DMN to functionally decouple from perceptual dominant systems allows DMN to operate in an offline manner dissociated from the external input (Smallwood, 2013). This is consistent with recent observations of the functional organisation of the cortical surface (Margulies et al., 2016) where the DMN is far from primary visual

and motor cortex in terms of Euclidean distance and functional connectivity. In addition, the integrative pattern between DMN and sensorimotor regions might be supported by increased abstraction and integration along the ventral visual system, ending in a conceptual hub in the ATL (Lambon-Ralph et al., 2017). This view is also consistent with the observation that a gradient from unimodal to transmodal cortex (Margulies et al., 2016) corresponds to increasingly abstract and complex cognitive tasks, where the influence of specific features linked to stimuli in the immediate environment is reduced (Mesulam, 1998; Buckner & Krienen, 2013; Margulies et al., 2016).

1.3 Neural hierarchies

The executive failure and representational accounts stemmed from distinctively different psychological research approaches to understand mind-wandering. The attentional system and the DMN are the common neural processes of both theoretical accounts, but interacting in a seemingly contradicting manner. In the executive failure account, the attention-related system deactivates with poor task performance, accompanied by the activation in the DMN; while in the representational account, the attention system and the DMN work together to maintain the internal representation of memory. One possible reason for this conflict could be that under different global neural configurations, DMN might support different cognitive functions. In the current section, I discuss the progress of functional neuroimaging study towards a hierarchical view of neural systems corresponding to the family resemblance account.

1.3.1 Historical perspective

Research on brain organisation has been dominated by two opposing views—*functional specialisation* and *functional integration*. Functional specialisation emphasises that small, distinguishable brain regions are solving distinct problems (Kanwisher, 2010). Studies of cognitive impairments in people with focal brain lesions provides the extensive evidence for the localisation of some functions in the human brain, such as the role of mid-fusiform gyrus in processing faces (Iaria, Fox, Waite, Aharon, & Barton, 2008) and left Brodmann areas 44 and

45 (left IFG) in speech production (Broca, 1861). Single-cell recordings and microscopic tissue examination revealed the segregation of occipital visual cortex (Zeki, 1978). Overall, the approach used to yield localise brain function shares one important similarity. The methodology leads to interpretations based on non-overlapping, discrete region as the basic compartments of brain organisation.

The recent focus in neuroscience has shifted from restricted regions to network organisation. Functional integration emphasises that cognitive function is enabled by a complex interplay between these distinct brain regions (Sporns, 2014). Biological neural network properties are an important source of electrophysiological oscillations. Independent component analysis (ICA) became the workhorse of network discovery in neuroimaging (Beckmann, DeLuca, Devlin, & Smith, 2005). Functional connectivity (Friston, 1994) and graph theory (Rubinov & Sporns, 2010) application to functional neuroimaging provided non-biophysical models of brain organisation. Recent advances in the field of human connectomics have revealed multiple large-scale networks, each characterised by distinct functional profiles (e.g. Yeo et al., 2011). In contrast to the specialisation of regions, cross-regional integration is the central approach to understanding the basic architecture of brain organisation.

Discovery from functional specialisation and integration has both revealed spatial gradients in brain organisation. Advances in mapping local streams such as vision (Zeki, 1978) have revealed spatial gradients extending along adjacent cortical regions. Stepwise functional connectivity analysis demonstrated transitions from primary sensory cortices to higher-order brain systems for perceptual integration in the human brain (Sepulcre, Sabuncu, Yeo, Liu, & Johnson, 2012). The diffusion embedding analysis on connectivity data in humans and the macaque monkey reveals the principal gradients of whole brain topographical organisation (Margulies et al., 2016). The discovery of multiple whole-brain functional configurations provides speculations on hierarchical relations among inter-plays of large-scale networks.

1.3.2 Abstract rule governing

Duncan (2010) discovered a common pattern of activity in the prefrontal and parietal activity of the human brain in response to a diverse cognitive challenge.

He revealed that the governance of cognitive control is the multiple-demand network (MDN) covering regions in the attention system and the frontoparietal network (FPN). This involves cortices in and around the posterior part of the inferior frontal sulcus (IFS), in the anterior insula (AI) and adjacent frontal operculum (FO), in the pre-supplementary motor area (pre-SMA) and adjacent dorsal ACC, and in and around the intraparietal sulcus (IPS). Similar multiple-demand patterns are identified in resting state data described as ‘task-positive’ (M. D. Fox et al., 2005), as opposed to the ‘task negative’ pattern—the DMN. The MDN activity is consistent with the neural model (Weissman et al., 2006) proposed to explain the executive failure account of mind-wandering (McVay & Kane, 2012a). The antagonistic roles of DMN and MDN seem to be essential to *abstract rule governing*.

The past research paradigms segregated complex cognition into isolated operations such as working memory capacity (Vogel & Machizawa, 2004) and response inhibition (Aron, Robbins, & Poldrack, 2004). Complex, multi-component behaviour should be examined to understand the central role of control in realistic behaviours. Tasks with multiple-demand properties, such as intelligence tasks, examine abstract thinking, multi-modal or feature integration skills, and working memory. The current thesis speculates that the multiple-demand neural patterns exhibit family resemblance property in executive control.

1.3.3 Sensory integration/segregation

Mesulam (1998) observed that the primary visual and auditory cortices form a spatially continuous organisation towards the hetermodal cortices of the frontal and parietal lobes. He hypothesised that the hetermodal regions are selectively converging the input from unimodal regions to form abstract information, forming a hierarchical polarity. This viewpoint was examined by Margulies et al. (2016) through a meta-analysis of cognitive function of the first principle gradient of the human brain. The first gradient anchors the unimodal regions at one end and the transmodal regions in FPN and DMN at the other. The continuum characterises a spectrum from unimodal to transmodal activity in a meta-analysis on cognitive function tasks, with sensory-driven tasks on the unimodal end and the abstract reasoning task on the transmodal end.

In this *sensory integration/segregation* view, the FPN and DMN are functionally adjacent. Empirical research in cognitive neuroscience has found a similar hierarchy in mental scene construction (Villena-Gonzalez et al., 2018) and higher-order conceptual representations (Murphy et al., 2018), which are essential functions supporting the representational account of mind-wandering. Recent research on the role of DMN has provided support for a global integrative view in which DMN forms the highest level in a neural hierarchy (Margulies et al., 2016). Investigation of the activity of DMN indicates major revisions of cognitive context when conducting an explicit task. Vatansever, Menon, and Stamatakis (2017) recently demonstrated that the integrative role of DMN with primary visual cortices and hippocampus facilitate a rapid and adaptive rule learning. Patterns of cognition with neural activity located along the gradient between the sensorimotor system and DMN will tend to share key characteristics, giving rise to a family resemblance in memory-guided and perception-guided representations.

1.3.4 Integration of hierarchical configurations

The neural hierarchy of *abstract rule governing* and *sensory integration/segregation* have described contradictory views of the organisation of DMN and FPN. In the *sensory integration/segregation* view, reflecting sensory integration/segregation, these two networks are near-neighbours. Yet in the *abstract rule governing* view, they tend to be in opposition. The commonality of the two views lies in the attention regulation role of FPN. DMN represents the lack of control in the abstract rule governing, while it serves as the integrational hub of information from FPN in the sensory integration view. Interestingly, these two neural hierarchies are consistent with the contradictions between the *executive failure* and *representational* account in ongoing thought, where mind-wandering results from poor task performance in executive failure account, but the representational account can explain the benefits and generation of mental representations in ongoing thought. The two neural hierarchies may describe the complementary whole-brain activity of these two theoretical accounts of ongoing thoughts.

The study of ongoing thought is in need of an integrative approach to pool related cognitive functions and whole brain patterns under a cohesive narrative.

The hidden family resemblance may resolve the conflict in theoretical views of ongoing thought. Recent studies of the MDN supports the family resemblance view that a component can possess multiple states. Crittenden, Mitchell, and Duncan (2015, 2016) demonstrated the involvement of both MDN and DMN in a rule switching task. Using multi-voxel pattern analysis, two tightly knitted subprocesses with distinct role have been revealed in multiple-demand tasks (Crittenden et al., 2015), while the coupling of the two networks has shown a board representation of abstract rules (Crittenden et al., 2016). The context dependency of the function of DMN implies the importance of the whole brain pattern to understand complex behaviour. Further development of the neuro-cognitive model is crucial to achieving a more granular view of ongoing thought (Mittner, Hawkins, Boekel, & Forstmann, 2016; Smallwood & Andrews-Hanna, 2013).

1.4 The current thesis

The conflicts in the mind-wandering literature arise from the heterogeneous, unconstrained nature of ongoing thought. To date, research on ongoing thought consists of investigations on three important aspects: experience, neural profile, and cognition (Figure 1.1). A detailed description of spontaneous thought is needed to confirm the experience of ongoing thought. The neural organisation serves as the intrinsic biological basis of cognition. Finally, established cognitive measures link functional outcomes to the experiential profiles.




Figure 1.1. Schematic of the thesis

Studies targeting relationships within a single aspect of ongoing thought (i.e. experiential, neural or cognitive) might only capture few potential members of the family of ongoing thought. When the family members sampled possess non-prototypical features, the discovery could be presented as rival theories. For instance, the study of ongoing thought associated with creativity and working memory are presented as conflicts of mind-wandering profiles. With an understanding of the component processes composing the family property, the heterogeneity of ongoing thought may be interpreted as exemplars of component processes rather than conflicts.

The current thesis adapts a multivariate approach as the first step towards a family resemblance view of ongoing thought. A multivariate approach can include observations on a wider variety of behavioural profiles, thus preventing over-representing an exemplar as the whole category. Multidimensional experience sampling (MDES; Medea et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013; Smallwood et al., 2016) is the main technique of experience profile assessment to capture the various aspects of thought. Resting-state functional connectivity is used to describe the trait-like neural feature of each individual. The tasks selected measure cognitive functions documented in the past mind-wandering literature, including executive control, fluid intelligence, episodic memory, semantic memory, and information generation. Finally, canon-

ical correlation analysis (CCA; Hotelling, 1936) is the conjoined-decomposition method of choice to explore the multivariate patterns of mind-wandering. Here I present the overview of the ongoing thought measure and the benefit of CCA used in the thesis.

1.4.1 Method of acquiring experience

In the current thesis, both online and retrospective method is used to acquire the content of ongoing thought in laboratory scenario. In the online measure, MDES uses thought probes to record participants' experiences during a given task. Thought probes appear during the task in a semi-random fashion. The online measure captures spontaneous thought in, possibly, both off-task and task-focused moments. In the first and third empirical studies, the averaged momentary report from MDES is used to capture the trait-like features of spontaneous cognition. In the second empirical study, New York Cognition Questionnaire (Gorgolewski et al., 2014) is used to acquire the trait-like multidimensional mind-wandering thought content during a 9-minute resting state fMRI session. The retrospective method measures the summary of the mind-wandering experience during a specified time period. The benefit of retrospective measure is that no interruption during the given task or the ongoing thoughts will happen. The trait-level features of mind-wandering are accessed in the retrospective report.

A hybrid of go/no-go task and n-back task was used to manipulate working memory capacity while recording online experiences (Konishi, McLaren, Engen, & Smallwood, 2015; Medea et al., 2016). The earlier version has used numbers as the test items (Smallwood, Tipper, et al., 2013; Smallwood et al., 2011). The majority of the experiment consists of nontarget presented in neutral colour and a small proportion of targets. Participants are instructed to judge whether the target number is odd or even. In the *choice reaction time* (i.e. 0-back) condition, the judgement is made when the number changes colour; in the *working memory* (i.e. 1-back) condition, participants judge the number on the previous screen when presented with a question mark. The improved version is proposed by Konishi and colleagues (2015), replacing numbers with two 2-dimensional geometric shapes separated by a vertical line. Each pair consists of two shapes among a circle, a triangle, and a square, each in two different left/right config-

urations. In the 0-back condition, the target is flanked by one of two shapes, and participants indicate which shape matches the target shape. In the 1-back condition, the target is flanked by two question marks, and participants match the target shape to the prior trial.

1.4.2 Content of experience

For the purpose of capturing the momentary evolution of thoughts during experiments, experience sampling (Kahneman, Krueger, Schkade, Schwarz, & Stone, 2004) is a commonly used technique. Experience sampling is conducted during an external task, such as reading (Franklin, Smallwood, & Schooler, 2011), go/no-go task (Christoff et al., 2009), and n-back task (Kane, Conway, Miura, & Colflesh, 2007). The list of questions needs to be short and concise to minimize interruption of the external task. To access the complex, heterogeneous content of spontaneous thoughts, the current thesis employs MDES (Medea et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013; Smallwood et al., 2016). MDES expanded the thought probe from an on/off task question to a collection of dimensions related to a wide range of questions about the form and content of thoughts. The idea of MDES is based on various questionnaires to understand the content of mind-wandering thoughts retrospectively, such as the Dundee Stress State Questionnaire (Matthews et al., 1999), Amsterdam Resting-State Questionnaire (Diaz et al., 2013), the resting state questionnaire (Delamillieure et al., 2010), and the New York Cognition Questionnaire (Gorgolewski et al., 2014). The questionnaires above include more than 20 items, providing comprehensive coverage of thought content. Direct implementation of the retrospective questionnaires listed above is not practical for experience sampling.

The current version of MDES is based on the 10-question set used in Medea et al. (2016) and Smallwood et al. (2016). In the previous work, the questions are separated into content and form aspects of the ongoing experience. PCA was used to extract latent linear structure in the spontaneous thought report. Each participant has a set of identified principle components concluding the average momentary state of all the sampled time period. The current version includes both form and content questions in one set with three extra questions. The average score of each thought dimension indicates the average momentary state

of ongoing thought. Unique associations among the questions are expected to capture the family resemblance at an experiential level.

1.4.3 Conjoined decomposition of brain and cognition

Despite being invented in the 30s, until fairly recently CCA has not aroused researchers' interests due to the lack of practicality. With the advance in computing resource and enriched data size, CCA has gained popularity in neuroimaging research. Three characteristics of CCA make it the method of choice. Firstly, CCA can be understood as a natural extension of PCA to two variable sets, but mutually linked by a joint correlation criterion (see 2.3.1 Joint information compression). Such extension enables the exploration of neuro-experiential component processes of the ongoing thought family. Secondly, CCA does not distinguish between the two variable sets during the information compression process (see 2.3.2 Symmetry). The identified dual-component dimensions correlate two aspects of the data together without indication of causality between neural function and cognition. Finally, CCA is capable of estimating more than one corresponding component pair from the two variable sets (see 2.3.3 Multiplicity). More than one pair of meaningful decomposition can be found, giving it the potential to examine the heterogeneity of mind-wandering content and their functional outcome.

While CCA provided useful features to explore ongoing thought, its sparsity variation overcomes two technical issues of CCA application. Performing feature selection with sparsity improves the interpretability of the data and model fit. Data with a higher number of features than samples is accompanied with consequences of the so-called curse of dimensionality—the more dimensions are added to a data set, the less explanatory value a sample would have (Domingos, 2012). The number of functional connectivity measure can easily exceed the number of samples. In sum, SCCA allows decomposition on functional connectivity measures without sacrificing data richness or introducing difficulties in interpretation following data compression. The advantage and disadvantage of CCA and its variation are discussed in the next chapter.

The current thesis adopts the sparse variation of CCA (SCCA; Witten & Tibshirani, 2009) to resolve arguments related to the component processes under the

family resemblance view. Most studies of ongoing thought have only focused on its relationship with one cognitive outcome at a time. As a consequence, mind-wandering is treated as a singular construct of all off-task ongoing thought. Studies on contents of ongoing experience have shown the diversity of information using self-report methods. Based on the heterogeneity of content, ongoing thought can be a collection of various type of spontaneous thoughts. Adopting a multivariate method has the potential to identify family resemblance among the heterogeneity of ongoing thought, and to begin the research on the ontology of the component processes of ongoing thought.

1.5 Summary and thesis outline

Human cognition has the capacity to generate thoughts loosely related to the external world. However, researchers have not understood the mechanism behind the consequences of ongoing thought. The heterogeneity of functional outcomes has been a controversial and much disputed subject within the field of mind wandering research. Extensive research has shown that both costs and benefits to cognitive functions are associated with mind wandering. Negative consequences include reduced attention, poor task performance, unhappiness, and depression. The related positive outcomes include creative problem solving, planning personal goals, and recovery from negative emotion.

Most studies of mind wandering have only focused on one cognitive outcome at a time. Mind-wandering is treated as a singular construct. Studies of the contents of spontaneous ongoing thoughts have shown the diversity of information using self-report methods. The temporal content of ongoing thoughts can be future- or past-focused. The topic can be on personal issues or task related. Based on the heterogeneity of content, ongoing thoughts can be a family of various type of spontaneous thoughts with essential shared features. The lack of understanding of family resemblance results in conflicts in ongoing thought literature.

In the past decade, the neural basis of unconstrained processes has become the centre of the topic. DMN is the commonly engaged large-scale network. The extensive literature has documented the task-negative trait of DMN. The executive failure account of mind-wandering is in line with the task-negative

view of DMN. However, the role of DMN in the memory representation account of ongoing thought is unresolved. Recent literature on neural hierarchy has cast a different light on the function of DMN. The highest-level abstract and heteromodal cognitive functions are associated with DMN, whereas perception and motor-related regions are related to sensory processing. This new view provides a clue to investigate the representational account of ongoing thought.

The conflict in the mind-wandering literature is related to the one-to-one matching between ongoing thought and the brain or behaviour. To paint the full picture of ongoing thought, a multivariate method will help incorporate the three aspects mentioned above: cognition, experience, and neural basis. The current thesis adopts SCCA to resolve arguments related to heterogeneity in mind wandering. With advances in computing resources and enriched data size, CCA has gained popularity in neural imaging research. The motivation of the current thesis is to resolve the conflict in the ongoing thought literature. Based on recent research on neural hierarchies of cognition, the current thesis argues that regions beyond DMN contribute to the cognitive process underlying ongoing thought. The analysis, therefore, incorporates the functional organisation of large-scale neural networks as the main neural measure. The aim is to present evidence supporting of the family resemblance view and to begin research on the ontology of the component processes in ongoing thought. An outline of the remaining chapters is listed below:

Chapter 2: Canonical correlation analysis

Canonical correlation analysis is introduced as the main method of the thesis. This review focuses on the potential applications in neuroimaging research. The features and applications of this multivariate method are outlined, followed by a discussion of the methods' interpretations and limitations.

This chapter is under preparation for publication.

Chapter 3: Exploring the heterogeneity of ongoing thought

Heterogeneity of mind-wandering leads to conflicting views about its functional outcomes in behavioural studies. Default mode network is commonly associated with the emergence of mind-wandering. SCCA conjointly decomposed functional

connectivity patterns of DMN and thought reports, revealing unique neuro-experiential components. The study then revealed that the neuro-experiential components each associated with unique cognitive task measures. The various connectivity configurations within DMN are associated with different types of ongoing thought and specific functional outcomes.

This chapter is published in Psychological Science.

Chapter 4: Population variation in the associations between large-scale networks and experiences

Unconstrained cognitive processes have two faces. The representational account argues that the primary sensory brain regions decouple from DMN to facilitate memory representation; whereas the executive failure account shows lapses in attention are related to the demands on attention system and activation of DMN. We used SCCA to extract related whole-brain functional connectivity patterns, profiling the neuro-experiential components of unconstrained cognitive processes. Examining the association between demanding cognitive tasks and neuro-experiential components, the study revealed evidence supporting both the representational and executive failure accounts.

This chapter is published in NeuroImage.

Chapter 5: Inhibition of prior information contributes to internal content representation

Various cognitive functions are involved in the generation of internal experiences. The study of the functional outcomes is often conducted in a manner of one-to-one mapping. The heterogeneous cognitive outcomes are discussed as conflicts rather than the complex details driving the diversity of ongoing thought. In Chapter 5, we explore the intrinsic whole-brain neural basis of the cognitive functions supporting the unconstrained generation of spontaneous thoughts. With SCCA we describe the conjoined decomposition of cognitive function and resting state functional connectivity. Similar to Chapter 4, this study explored the unconstrained neuro-cognitive mechanism underlying the various dimensions of ongoing thought.

Chapter 6: General discussion

The overarching themes of the thesis are discussed and linked to specific results throughout the thesis. Future research directions are inspired by the findings and limitations of the current thesis based upon the key questions which this thesis attempted to answer:

- Why does ongoing thought show both costs and benefits?
- Can functional neural hierarchy explain the heterogeneity?
- Is the family resemblance view viable for ongoing thought?

Chapter 2

Canonical Correlation Analysis

The following chapter has been adapted from: Wang, H.-T., Smallwood, J., & Bzdok, D. (2018). Finding the needle in high dimensions: A tutorial on CCA in biomedicine. Manuscript preparing for publication. ¹

2.1 Abstract

Since the beginning of the 21st century, the sample size of studies in medicine and neuroscience has grown rapidly. For example, data sets with thousands of subjects are becoming more common and they often entail extensive neural and behavioural phenotyping yielding datasets with tens of thousands of variables. The size and complexity of these big data sets pose new challenges to researchers hoping to use them to understand relationships between brain, cognition and disease. Canonical correlation analysis (CCA) is a promising method for dealing with and harvesting insight from these large data sets. CCA allows two input variable sets to be simultaneously considered and extracted, such as descriptions of the brain and behaviour. The present tutorial paper introduces rationale, promises, and pitfalls of CCA.

¹D. Bzdok and H.-T. Wang planned the structure of the manuscript. H.-T. Wang. drafted the manuscript under the supervision of D. Bzdok and J. Smallwood.

2.2 Motivation

Large biomedical data sets and increasing computational power have opened up novel ways to conceive of understanding the relationships among brain, cognition and disease. Similar to the advent of microarrays in genetics, brain-imaging and extensive behavioural phenotyping yield datasets with tens of thousands of variables. Since the beginning of the 21st century, the popularity and feasibility of technologies, such as functional magnetic resonance imaging, (fMRI) have made it more practical to collect large neuroscience data sets. At the same time, problems in reproducing the results of key studies in neuroscience and psychology have highlighted the importance of these large data sets. Accordingly, there has been a staggering increase in the collection of large cohort datasets (Efron, 2010). For instance, UK Biobank is a prospective population study with 500,000 participants and comprehensive imaging data, genetic information and environmental measures on mental disorders and other diseases (Allen et al., 2012; R. L. Miller et al., 2016). The Human Connectome Project (HCP; van Essen et al., 2013) has recently completed brain-imaging of more than 10,000 young adults, with 4 hours of body scanning per subject, and utilising vast improvements in the spatial and temporal resolutions of the acquired data. Both the Enhanced Nathan Kline Institute Rockland Sample (Nooner et al., 2012) and the Cambridge Centre for Aging and Neuroscience (Cam-Can; Taylor et al., 2017; Shafto et al., 2014) reflect large ($N \geq 700$), cross-sectional adult lifespan (18–87 years old) population-based samples. These datasets describe changes in cognition and brain structure and function, with raw and preprocessed brain imaging data and cognitive behavioural experiments and demographic and neuropsychological data. While extensive phenotypes and big sample size provide opportunities for more robust descriptions of key population variation, these are not without associated costs. Classical statistical tools struggle to resolve datasets with more variables than observations, and even large samples of participants are smaller than the number of voxels that are possible in state of the art high-resolution brain imaging scans. On the other hand, in large samples, standard statistical techniques often yield highly significant associations that only account for a very small fraction of the variance to be explained. The growing interest in big data sets, therefore, requires that researchers must seek alternative tools

to gain the benefit provided by big data sets.

The present tutorial paper considers the suitability of Canonical correlation analysis (CCA) as a tool for charting and generating understanding from big data sets. One key feature of CCA is that it can simultaneously evaluate two matrices of information, such as many brain measurements and many behavioural measurements. In particular, CCA simultaneously identifies the main sources of variation that are common to both sources of variation. CCA is a multivariate statistical method that was introduced in 1936 (Hotelling, 1936). However, CCA is computationally expensive and so has only become a useful tool for biomedicine relatively recently. One important feature of CCA is that it describes dimensions that unravel the correspondence between two different sets of variables. In cognitive neuroscience, this often allows the determination of variation that links patterns of brain activity to patterns of behaviour. Moreover, the multivariate nature of CCA allows the identification of patterns that describe many-to-many relations and so provides a utility that goes beyond techniques that map one-to-one relationships (e.g., Pearson correlation) or many-to-one relationships (e.g., linear support vector machines). With the advent of larger datasets, researchers in neuroscience have begun to take advantage of these features of CCA to address novel questions regarding the links between brain, cognition and disease (Marquand, Haak, & Beckmann, 2017; Smith et al., 2015; Tsvetanov et al., 2016; Vatansever, Bzdok, et al., 2017; Wang, Poerio, et al., 2018; Wang, Bzdok, et al., 2018).

Our guide to CCA proceeds in four parts. We first introduce the model in detail and the circumstances of use with recent applications of CCA in existing research. Next, we consider the quantitative conclusions that can be drawn from the application of the CCA algorithm, with special attention to the limitations of this technique. Finally, we provide a set of practical guidelines about how the analysis can be used moving forward.

2.3 Modelling intuitions

One way to appreciate the idea behind CCA is by viewing this procedure as an extension of the widely applied principal component analysis (PCA). PCA produces a set of dimensions that act as a close approximation of the variance in the

original data set, except in a compressed form. In other words, PCA converts a set of correlated variables into a smaller number of hidden factors that were not directly observable in the original data, but explain the structure of the observations in an efficient way. As a prominent example, the Big Five personality is a psychological construct of human personality traits discovered through PCA (Barrick & Mount, 1991). In this case, personality survey data is entered into a PCA, which produces five components that explain a substantial amount of meaningful variation within the data. The advantage of decomposition methods such as PCA is its ability to reduce the original data sets to fewer dimensions that are more amenable to psychological interpretation. Such re-expression of the original data in a compressed, more parsimonious form has computational-statistical and interpretational appeal, while still capturing a large amount of the variability in the original large variable array. Unlike PCA, CCA maximises the linear correspondence between linear combinations of two variable sets, by seeking dimensions of variance that described shared variance across both sets. CCA, therefore, is particularly useful when describing observations that bridge two domains, for example, (i) genetics and behaviour, (ii) brain and behaviour, or (iii) brain and genetics. There are three characteristics for modelling data using CCA: joint-information compression, symmetry and multiplicity.




Figure 2.1. An example of CCA on behavioural data.

Figure from Bailey et al. (2013). Consider a case of exploring the relations between movie genre and personality with CCA. On the left-hand side, we can input the data related to the movies the participants watched, such as the number of action/ documentary/comedy etc. movies watched. The right-hand side is the personality rating of the participants, for example, extrovert vs not, openness vs not etc. The CCA can then find the related personality traits with the type of movies they are likely to watch.

2.3.1 Joint information compression

The purpose of CCA is to find sources of variability that are common to two sets of variables. The relations of a pair of factors across sets of variables, their canonical correlation, indicates the conjoined explained variance across both domains. In a similar manner to PCA, CCA aims to find the most compact linear patterns, canonical variates, based on the variance explained under the contained of uncorrelated hidden dimensions. Also, similar to PCA these dimensions are ranked by their explained variance, with earlier dimensions accounting for more variance than later dimensions. An example of the application of CCA to behavioural data is presented in Figure 2.1.

2.3.2 Symmetry

CCA does not distinguish between the left and right variable sets during the information compression process. Instead, the canonical correlation indicates that a unit change in a component in one set of observations is consistently

associated with an equivalent change in the other set of observations. As CCA does not privilege one of the variable sets, numerically identical decomposition results are returned regardless of the input order of the variable sets. This feature of CCA is known as symmetry and is a key feature that distinguishes CCA from other linear regression methods. The symmetrical nature of CCA can be contrasted with linear-regression models, in which dependent and independent variables play different roles in the analysis. Regression indicates the impact of a unit change in the independent variable on the dependent variables, therefore dependent and independent variables cannot be exchanged to obtain an identical result. As CCA is a correlation-based method, it describes the co-relationship of the two variable sets, thus the exchange of the two variable sets produces identical results.

2.3.3 Multiplicity

In CCA, a pair of dimensions that share variance in both sets of observations is known as a mode. A mode contains a pair of canonical variates that describe the linear structure of the two variable domains. After finding the mode that describes the most variation, CCA will next determine the next pair of dimensions that remains in the unexplained variance of both data sets. Since every new mode was found in the residual variance, the modes are optimised to be uncorrelated with each other. In this manner, CCA produces a set of mutually orthogonal modes naturally ranked by explained variance. The orthogonality constraint ensures the modes represent unique linear patterns that describe different features in the data. When the modes are theoretically meaningful, the researchers can potentially use these to formulate a component process approach to interpret the data.

2.3.4 Interim summary

In conclusion, CCA uncovers effective, symmetric linear relations that compactly summarize doubly-multivariate data. We introduced three important characteristics of CCA. First, CCA provides more effective hidden representation that captures most variance in original variables. Next, the CCA model is symmetrical in the sense that no numerical difference happens in the exchange of the two

variable sets. Finally, we can estimate several modes of correspondence between the two variable sets. In the next section, we would like to explore examples of CCA applications.




Figure 2.2. The analysis pipeline of Smith et al. (2015). The arrows represent analysis performed.

2.4 Examples

Smith et al. (2015) employed CCA to uncover brain-behaviour modes of population co-variation in the HCP (van Essen et al., 2013). Smith et al. (2015) aimed to discover whether any specific patterns of functional brain connectivity, on the one hand, are associated with specific sets of correlated demographics and behaviour on the other hand (see Figure 2.2 for the analysis pipeline). Functional brain connectivity was retrieved from resting state functional scans which measure the brain activity in the absence of a task or stimulus (Biswal, Zerrin Yetkin, Haughton, & Hyde, 1995). Independent component analysis (ICA; Beckmann et al., 2005) was used to identify 200 networks from the resting state scans. ICA identifies independent networks by separating the spatial sources of the resting state data. Next, functional connectivity matrices were calculated based on the pair-wise correlation of the 200 networks. The behavioural measures ranging from cognitive function to demographic information were entered into the CCA as one set of observations and the functional connectivity matrices were

the second set. The robustness of the modes was determined via permutation tests on the canonical correlations. One significant mode demonstrated strong population-level co-variation of network connectivity and behavioural measures. The behavioural measures varied along a positive-negative axis with intelligent, memory and cognition tests and life-satisfactory on the positive end and negative lifestyle measures anchoring the other end. The brain regions highly contributing to the connectivity resembles the default mode network (DMN; Buckner, Andrews-Hanna, & Schacter, 2008). The positive-negative dimensions in the behavioural component and the emergence of DMN in the brain component may seem trivial on their own, however, CCA formalised the relation of the underlying biology and the correlation among the general behavioural measures that captures intelligence. Regions composing DMN has been associated with episodic and semantic memory, scene construction, and complex social reasoning such as the theory of mind (Andrews-Hanna, Smallwood, & Spreng, 2014). The finding of Smith et al. (2015) provided evidence that the DMN is important for higher-level cognition, especially intelligence—one of the perhaps most important indices so far identified by psychologists.




Figure 2.3. The analysis pipeline of Wang, Poerio et al. (2018). The arrows represent analysis performed.

Another use of CCA has been to understand the relationship between patterns of brain activity and variations of experience. Wang, Poerio, et al. (2018,

see Figure 2.3 for the analysis pipeline) used CCA to examine links between the DMN and patterns of ongoing thought. In both the laboratory and in daily life, ongoing thought can often shift from the task at hand, a phenomenon that is characterised by the experience of mind-wandering (Smallwood & Schooler, 2006, 2015). These shifts have been associated with poorer performance on attention-demanding tasks (McVay & Kane, 2009; Mrazek et al., 2012), yet studies of problem-solving suggest that mind wandering may promote creativity (Baird et al., 2012; Smeekens & Kane, 2016) and future planning (Baird et al., 2011; Medea et al., 2016). Despite the heterogeneity of functional outcomes, task-unrelated thoughts during mind wandering are linked to changes in DMN activity (see the review from Smallwood & Schooler, 2015). Wang, Poerio, et al. (2018) used CCA to examine the hypothesis that the reason why patterns of off-task thought can have opposing links with behaviour is that there are distinct patterns of the population of variance that link different types of ongoing thought to activity in the DMN. Their analysis used patterns of connectivity within the DMN as one set of observations and self-reported descriptions recorded in the laboratory, recorded across multiple days, as the second set. The connectivity among 16 DMN regions and 13 self-report questions on thoughts were entered using a sparse version of CCA (Witten & Tibshirani, 2009). Two stable modes corresponded to traits of positive-habitual thoughts and spontaneous task-unrelated thoughts both with unique patterns of neural connectivity patterns within the DMN. Importantly, subsequent analyses identified that the modes were uniquely related to aspects of cognition, such as executive control and the ability to generate information in a creative fashion, and independently distinguished well-being measures. These data suggest that the DMN can contribute to ongoing thought in multiple ways, each which have unique behavioural associations. Wang, Poerio, et al. (2018), therefore, suggest that mind wandering is a collective term for various types of spontaneous thought (see Seli et al., 2018). The different configurations of DMN also demonstrated its possible role as an integrator in cognition, rather than a task-unrelated network (Margulies et al., 2016).

2.5 Interpretations

The symmetrical data compression feature of CCA makes it particularly useful to help researchers handle the complexity of two sets of variables. However, whether the process of the analysis is an exploration of hidden structures or mutually constrained predictive component remains a matter of debate. CCA can be viewed as a supervised predictive algorithm or as an unsupervised exploratory algorithm. A supervised algorithm relies on the predefined labels/relationships in the data to form prediction; whereas an unsupervised algorithm aims to extract patterns in the data with unlabelled data. CCA has some properties of both supervised and unsupervised modelling approach. The more the dimensionality of one of the variable sets resembles the single output of linear-regression-type methods, the more CCA application approaches output is similar to a supervised modelling approach. In contrast, with larger variable sets on both sides, the more CCA resembles an unsupervised modelling approach.

The main difference between supervised and unsupervised method depends on whether there is a set of determined variables as the goal of modelling. To examine whether the models prediction matches the desired result, a supervised learning method contains a learning target or loss function. The difference between the prediction generated by the model and the real label is the objective of a so-called loss function. CCA has the objective to maximise the linear correlation between the latent dimensions from two variable sets. While most supervised learning methods estimate loss between real data and predictions, CCA has an unusual objective that we rarely see in supervised estimators. Symmetry is another reason that makes CCA an unusual case of supervised learning. In supervised learning, the model learns the pattern in the data to predict a set of targets. The symmetrical nature of CCA does not distinguish the two sets of input variables. The data compression and hidden structure inspection aspect put CCA in line with unsupervised methods. The conjoined decomposition, therefore, captures the relations among the variables. The found relations are used to construct fewer factors that capture the variance of the original data. CCA has the strength of both supervised and unsupervised methods. Like unsupervised methods, CCA can search through candidate patterns to find structure in data. The accurate predictions formed by CCA highlights the trait of a supervised

method. In conclusion, CCA is a special case that sits in-between the supervised and unsupervised methods. The flexibility in interpretation offers people multiple ways to utilise CCA in research.

Statistical methods can be categorised into three categories based on their goals: estimation, prediction, and inference. Estimation represents ways or a process of learning and determining the population parameter based on the model fitted to the data. Prediction is making an inference of an unknown data points based on information obtained from a sample. Inferential statistical analysis utilises hypothesis testing to draw conclusions about populations or scientific truths from data. CCA falls into the family of estimation. The focus of CCA is to establish statistical associations (i.e. the latent linear relations among variables, the association between the two latent linear relations). Predicting some variables based on other variable is not the optimisation goal. CCA does not seek to establish ‘statistically significant links between variables’. The null hypothesis that is really tested around the robustness of the latent space correlation (i.e., the canonical correlations of the latent variables extracted from the two variable sets) across modes, not so much particular variable-variable links. CCA is often used to rigorously evaluate whether overall linked covariation patterns can be found in two variable sets, rather than pinpointing and ‘putting the finger on’ certain specific relations that should be interpreted with more caution.

2.5.1 Utility of CCA

CCA does not come without limitations and here we discuss several issues that researchers should keep in mind when evaluating whether the data set is suitable for CCA. We summarise these choices in the form of a flowchart (see Figure 2.4). As with many statistical approaches, the sample size is an important factor when considering whether CCA is appropriate. CCA can handle data with more observations than the number of variables of the smaller variable set (i.e. $p < \min(m, n)$). However, smaller data set does not fully utilise the strength of CCA since they tend to have less variability. Relationships in areas like neuroscience are often small, and so a large number of samples is required to correctly infer the variability in data. On the other hand, if the number of variables of either side of the equation exceeds the sample size, CCA does not generate unique linear

combinations for each variable set. In such a scenario, a PCA rank-reduction preprocessing step is commonly performed before applying CCA (Smith et al., 2015). The use of CCA, therefore, is more appropriate when samples sizes are reasonably large relative to the sets of variables being analysed.




Figure 2.4. A flowchart illustrating the choices when considering the application of CCA to a dataset.

A second issue is the nature of the relationships that CCA describes in the underlying data. CCA is essentially a linear model, and so makes a set of assumptions regarding the normality of the distribution of the observed data, as well as the linearity of the underlying relationships. CCA can accommodate any metric variable without the strict assumption of normality. However, normality is desirable as it allows for the highest correlation among variables, and this makes the task of identifying the underlying dimensions easier. It is recommended to evaluate the normality of all variables and apply data transformation where appropriate before CCA is applied to the data. CCA also assumes that relationships within the data are linear and this introduces two limitations. Since only linear effects can be captured, any patterns within the data that are non-linear (e.g. quadratic or cubic relationships) will not be captured by this analysis. Finally, in CCA relationships are optimised in terms of how effectively they describe linear correlations between variables within the observed data. Relationships discovered by CCA, therefore, should not be considered as predictive accounts

of relationships with future data. Any interpretation and application related to predictability of the canonical components in new data should be treated with caution. If the prediction of future observations is important, a small fraction (20 – 25%) of the full data can be left out as a test data set when the sample size is sufficiently big. The analysis will be applied on the training set first to retrieve the canonical modes. The predictability of the canonical variates from the modes will be examined by calculating the explained variance on the test set.

2.5.2 Relation to other commonly used methods

CCA can be framed as a general example of other statistical procedures that are derived from the **general linear model** (GLM). Virtually all of the parametric tests most often used by behavioural scientists (e.g., ANOVA, MANOVA, multiple regression, Pearson correlation, t-test) can be subsumed by CCA as special cases in the GLM (Knapp, 1978; Thompson, 2015). Because these techniques are intricately related and fundamentally the same in many respects of CCA, learning CCA may help facilitate conceptual understanding of statistical methods throughout the GLM.

CCA is related to other feature learning methods in neuroscience. As mentioned in the modelling intuition section, **PCA** is a similar method performed on one set of variables only. The objective of PCA is a transformation of several possibly correlated variables into a smaller number of uncorrelated variables known as principal components. PCA compresses data into a smaller number of factors that carry most of the variance of the data.

Independent component analysis (ICA) performs a linear transform that makes the resulting variables as statistically independent from each other as possible. The basic assumption behind ICA is that the data is composed of independent sources of information. In contrast to PCA and CCA, all components are equally important. ICA helps when you want to find a representation of your data as independent sub-elements

Partial least squares (PLS) regression and CCA are both techniques for feature extraction from two sets of multidimensional variables. The fundamental difference between CCA and PLS is that CCA maximizes the correlation while

PLS maximizes the covariance. PLS is a supervised approach that attempts to find directions that help explain both the response and the predictors. Most of the comparison between PLS and CCA is focused on the generation of the first components. There is no one way to compute the other components in PLS (De Bie, Cristianini, & Rosipal, 2005).

Several useful extensions of CCA has emerged in the past decade to overcome the limitation of the linear version of CCA. The first was the nonlinear version of CCA, with **kernelization of CCA** (KCCA; Haroon, Szedmak, & Shawe-Taylor, 2004) as the representative. Kernels are methods of implicitly mapping data into a higher-dimensional feature space with kernel function, a method known as the kernel trick. KCCA first projects the data into a higher-dimensional feature space before performing CCA in the new feature space. While KCCA allows learning of nonlinear representations, the drawback is that the representation is limited by the fixed kernel. KCCA is a nonparametric method, hence, the time required to train KCCA or compute the representations of new data points scales poorly with the size of the training set.

Sparse CCA (SCCA; Witten & Tibshirani, 2009) is a method for identifying sparse linear combinations of the two sets of variables that are highly correlated with each other. It has been shown to be useful in the analysis of high-dimensional data when the variable number of either array is higher than the number of samples. The sparse feature reduces some coefficients to 0 in the linear structure depending on the penalty parameters. The benefit of sparsity is an improvement in interpretation and feature selection. However, sparsity violates the orthogonality of CCA, meaning the components in different modes can correlate with each other. The explained variance of each mode will not follow the rank order either. Recently, we have used k-fold cross-validation to identify the ideal level of sparsity in a study that explores the relationship between patterns of thought at rest and the associated brain organization (Wang, Bzdok, et al., 2018).

With the recent advance in the deep neural network, **deep CCA** (DCCA; Andrew, Arora, Bilmes, & Livescu, 2013) has been proposed as an alternative to KCCA as a non-linear method for CCA. A deep neural network is an algorithm that learns the representation of data through multiple non-linear transforma-

tions. The name came from the architecture that loosely follows the connections of neurons. DCCA simultaneously learns two deep neural network mappings of two variable sets that are maximally correlated. The main advantage is the faster performance over KCCA because DCCA directly learns the data without re-mapping data into a higher dimension.

2.6 Practical considerations

Having introduced the features and interpretations of CCA, we close by considering the practical considerations that determine how it should be implemented. The computation of CCA is available in the in-built library of MATLAB (`canocorr`) and R (`cancor`), and Python machine-learning library `scikit-learn` (`sklearn.cross_decomposition.CCA`). All implementations above provides comprehensive documentation for how to implement CCA and we describe the additional steps that a researcher may wish to consider before applying CCA to their data.

2.6.1 Preprocessing

Some minimal data preprocessing is usually required for most machine-learning methods. CCA is **scale-invariant** in that applying some standardising data transformation on the columns of the variable sets should not change the resulting canonical correlations. This property is inherited from Pearson's correlation defined by the degree of simultaneous unit change between two variables, with implicit standardisation of the data. Nevertheless, z-scoring of each variable of the measurement sets is still recommended before performing CCA to facilitate the model estimation process and to enhances interpretability. To avoid outliers skewing the results, application of **outlier detection techniques** and **outlier cleaning** are recommended before the analysis, such as imputation with the mean or median of the variables. Aside from outliers, some confound variables can also introduce unwanted effects. Confound variable removal is recommended as a preprocessing step to reduce the risk of finding non-meaningful associations. The same rules are also commonly applied for the GLM in the neuroimaging data. The subsequent analysis would reflect the data without influences of the

confound. In neuroimaging, for example, motion is thought to be an important confound (Power et al., 2014) and as consequence, it is common to remove the influence of this variable prior to conducting CCA (see Wang, Bzdok, et al., 2018).

When the number of variables exceeds the number of samples, PCA is recommended as a preliminary **dimension reduction** step before performing CCA. An example is work by Smith et al. (2015). The application of PCA compresses the number of variables in each matrix to the most explanatory dimension of variation. However, a downside of this method is the difficulty to directly map the resulting dimensions to the original data. To interpret the CCA solutions in the original data, Smith et al. (2015) correlate the canonical variates to the original data to recover the relevant variate captured by the CCA component pair.

CCA component can be challenging to interpret, especially when a PCA dimension reduction is applied. An alternative solution is a **CCA+ICA** method in addition to the PCA (K. L. Miller et al., 2016; Sui et al., 2010) has been proposed to overcome the issue of projecting the PCA-compressed data back to the original space. In the original CCA+ICA approach, the assumption is that CCA extracted components are an incomplete decomposition with multiple possible sources (i.e. patients vs controls). CCA first finds the correlated variance of the two variable set. After CCA, the canonical components are concatenated into one array. ICA is then applied to the canonical components to recover the source of the variance. The ICA step can be done in the full feature space by projecting the CCA components to the PCA components (K. L. Miller et al., 2016). The CCA+ICA approach achieves both high estimation accuracy and provides the correct connection between two variable sets. The ICA step is especially useful in the detection of independent components that contribute to the common solution extracted from the two variable sets.

2.6.2 Model selection

CCA allows multiple modes to be calculated from the observed data, however, it is necessary to specify the appropriate number of these latent sources of variation. To select the number of modes (canonical component pair, see Section 2.3), we

can use **explained variance** metrics to determine a useful number of modes. Since the canonical components are the compressed information of the original data, the canonical component would be expected to be related to the original data. The calculation can be done by predicting the canonical components with the original scores. Another solution is to calculate the family-wise error rate through **permutation tests**. The incentive of a permutation test is to assess the robustness of results when comparing to results from randomly re-arranged data. The permutation is done by randomising one of the variable sets to break the unique relationship between the two variable sets in each observation instance. The extracted modes from the randomised sample will serve as the chance level results. Permutation tests establish robust above-chance correspondences between variable sets, but no null hypothesis significance testing for an individual variable is tested this way. The first canonical correlation of the permuted sample is compared with all the CCA modes extracted from the real data. The p-value for each mode is calculated as the number of permuted sample canonical correlation higher than the given mode from the real sample, divided by the number of permutation.

The variations of CCA might need an extra step for **hyperparameters selection**, such as the knobs of kernel type and penalty in KCCA, penalty strengths in SCCA, and layer number in DCCA. A permutation or cross-validation scheme is recommended for hyperparameters selection. The permutation test on hyperparameter selection is set up in the same way as model selection, but focusing on the first canonical correlation only (For example, see Appendix A in Witten & Tibshirani, 2009). In terms of cross-validation, the objective function for model selection can be the out-of-sample explained variance or the variance loss between the training set and the testing set.

2.7 Summary

In biomedicine research, the relationships among brain, cognition and disease are often complicated. Focusing on a small selection of measures will risk in ignoring potential factors. CCA is a doubly multivariate pattern analysis on two variable sets. With no directionality implied on either variable set, CCA enables more flexibility on the research questions. As the interest in multitask data and

rich cognitive phenotyping in large datasets grows, CCA fulfills the need of a method that can consider a large set of possible variables in one analysis. With its ability to reduce the data to meaningful and concise information, CCA is a promising method for scientists who are interested in exploration of multivariate patterns in large data sets.

Chapter 3

Exploring the Heterogeneity of Ongoing Thought

*The following chapter has been adapted from: Wang, H.-T., Poerio, G. L., Murphy, C. E., Bzdok, D., Jefferies, E., & Smallwood, J. (2018). Dimensions of Experience: Exploring the Heterogeneity of the Wandering Mind. *Psychological Science*, 29 (1), 56-71. doi:10.1177/0956797617728727¹*

3.1 Abstract

The tendency for the mind to wander to concerns other than the task at hand is a fundamental feature of human cognition, yet the consequences of variations in its experiential content for psychological functioning are not well understood. Here, we adopted multivariate pattern analysis to simultaneously decompose experience-sampling data and neural functional-connectivity data, which revealed dimensions that simultaneously describe individual variation in self-reported experience and default-mode-network connectivity. We identified di-

¹ J. Smallwood, E. Jefferies, H.-T. Wang, and C. Murphy designed the study. H.-T. Wang, C. Murphy, and G. Poerio collected the data. The connection-strength and sparse canonical-correlation analysis pipeline was constructed by D. Bzdok and H.-T. Wang. Data were analyzed by H.-T. Wang, C. Murphy, and G. Poerio under the supervision of D. Bzdok, J. Smallwood, and E. Jefferies. H.-T. Wang and J. Smallwood drafted the manuscript. G. Poerio and D. Bzdok provided critical revisions. All the authors approved the final version of the manuscript prior to submission.

mensions corresponding to traits of positive-habitual thoughts and spontaneous task-unrelated thoughts. These dimensions were uniquely related to aspects of cognition, such as executive control and the ability to generate information in a creative fashion, and independently distinguished well-being measures. These data provide the most convincing evidence to date for an ontological view of the mind-wandering state as encompassing a broad range of different experiences and show that this heterogeneity underlies mind wandering’s complex relationship to psychological functioning.

3.2 Introduction

Although people’s minds frequently wander from events in the here and now, or any task being performed, the functional consequences of this state remain poorly understood (Mittner et al., 2016; Seli, Risko, Smilek, & Schacter, 2016; Smallwood & Andrews-Hanna, 2013). Some studies link mind wandering to unhappiness (Killingsworth & Gilbert, 2010); others suggest it facilitates recovery from negative emotional states (Poerio et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013). Mind wandering is associated with poorer performance on tasks that place high demands on executive functions (McVay, Kane, & Kwapil, 2009; Mrazek et al., 2012), yet studies of problem solving suggest that mind wandering may promote creativity (Baird et al., 2012; Smeekens & Kane, 2016). This wide range of associated functional outcomes is puzzling—if mind wandering is a homogeneous construct, then it is unclear why it should be associated with such a complex array of often opposing outcomes. To reconcile this contradictory evidence, researchers have suggested that mind wandering may be heterogeneous, encompassing multiple states with differential contents and underlying cognitive architectures (Smallwood & Andrews-Hanna, 2013). According to this ontological perspective, different functional associations arise from different ‘types’ of experience, which explains the range of functional outcomes observed in the literature.

In the current study, we recruited 165 participants and obtained data on (a) the organization of the brain at rest using functional MRI (fMRI), (b) the content and form of experience recorded across different days, (c) cognitive functions assessed by a comprehensive battery of tasks (including memory, creativity,

and executive control), and (d) psychological well-being via questionnaires. Our procedure is presented in Figure 3.1. These data allowed us to use novel multivariate analysis methods to test the hypothesis that there are different types of mind wandering, with unique neural and experiential patterns accounting for unique variance in the psychological profile of our sample.

We used functional connection strength to characterize the neural organization of each individual. We selected regions for our analysis on the basis of evidence that task-unrelated thoughts are linked to concurrent increases in activity in medial prefrontal cortex (mPFC), posterior cingulate cortex (pCC), and lateral parietal cortex (for meta-analyses, see K. C. R. Fox, Spreng, Ellamil, Andrews-Hanna, & Christoff, 2015; Stawarczyk & D’Argembeau, 2015). regions that make up the core of the default mode network (DMN; Buckner et al., 2008). During mind wandering, it is believed that these regions interact with other areas of the cortex, in particular, temporal lobe regions associated with memory representation that are also allied to the DMN. For example, the hippocampus activates early during mind wandering (Ellamil et al., 2016). whereas connectivity between lateral and medial aspects of the temporal lobe and the DMN core predicts individual variation in features of mind wandering, such as its episodic content (Karapanagiotidis et al., 2017). Contemporary accounts of mind wandering posit that the DMN may be important for automatic aspects of cognition (Christoff, Irving, Fox, Spreng, & Andrews-Hanna, 2016). Other studies have highlighted links with the lateral prefrontal cortex, which is important for executive control when mind wandering is more deliberate (e.g., Golchert et al., 2017).

We applied multivariate pattern analysis to the neurocognitive and experiential data to identify different types of mind wandering. If the DMN is important for automatic aspects of cognition (Christoff et al., 2016), states linked to high levels of connectivity within this system may have experiential features reflecting more automatic types of cognition. Our a priori decision to focus on the DMN core to derive patterns of experience limited our ability to observe interactions with regions outside of this system, so we used whole-brain functional connectivity to characterize these links for each type of experience. On the basis of prior studies (e.g., Ellamil et al., 2016; Golchert et al., 2017; Smallwood et al., 2016),




Figure 3.1. Schematic of the procedure and analysis strategy employed in the current study.

We collected resting-state functional MRI (fMRI) data, cognitive-function measures, and questionnaires on personal traits for each participant. These were submitted to analysis (rectangles with dashed borders), which created latent components or features (rectangles with solid borders) for each subject, and these variables were subsequently passed to the main analysis (ovals). WHOQOL = World Health Organization Quality of Life assessment (WHO, 2002).

we expected this analysis to identify connections with regions in the temporal lobe or the executive system. This pattern would confirm the hypothesised accounts of the DMN as important in integrating neural information (Margulies et al., 2016; Smallwood et al., 2016). Having characterised different types of mind wandering in both brain and experience, we used these to test the hypothesis that different categories of experience are related to different functional outcomes. We performed an individual differences analysis to understand whether our characterised types of mind wandering have unique functional associations, including better creativity, worse executive control, and lower levels of well-being. We expected different patterns of experience to capture different psychological profiles explaining the heterogeneous pattern of functional outcomes that have been linked to the mind-wandering state in previous studies (Smallwood, Ruby, & Singer, 2013).

3.3 Method

3.3.1 Participants

One hundred sixty-five healthy participants were recruited from the University of York (99 females, 66 males; age range = 18–31 years, $M = 20.43$, $SD = 2.63$). We preselected a sample size approximately double those used in our prior studies (e.g. Smallwood et al., 2016). A sample size of at least 125 is recommended in order to have 95% confidence that a correlation of typical size ($r = .20-.30$) is present and greater than 0 (Hemphill, 2003). Participants were right-handed native English speakers with normal or corrected-to-normal vision and no history of psychiatric or neurological illness. Participants underwent MRI scanning, completed an online questionnaire, and then attended three 2-hr behavioural testing sessions to complete a battery of cognitive tasks. The behavioural sessions took place within a week of the scan. Eight participants were excluded from the multivariate pattern analysis because they failed to complete all of the behavioural testing sessions. In total, 157 participants were included in the multivariate pattern analysis and the comparison with cognitive performance. One hundred forty-two participants completed both the behavioural testing sessions and questionnaires and were included in the analysis associated with well-being.

Participants were rewarded with either a payment of £80 or a commensurate amount of course credit. All participants provided written consent prior to the fMRI session and the first behavioural testing session. Approval for the study was obtained from the ethics committee of the University of York Department of Psychology and the University of York Neuroimaging Centre.

3.3.2 MRI acquisition

Structural and functional data were acquired using a 3T HDx Excite MRI scanner (GE Healthcare, Little Chalfont, United Kingdom) utilising an eight-channel phased-array head coil tuned to 127.4 MHz at the York Neuroimaging Centre, University of York. Structural MRI acquisition in all participants was based on a T1-weighted 3-D fast-spoiled gradient-echo sequence repetition time (TR) = 7.8 s, echo time (TE) = minimum full, flip angle = 20° , matrix size = 256×256 , 176 slices, voxel size = $1.13 \times 1.13 \times 1 \text{ mm}^3$. Resting-state activity was recorded from the whole brain using single-shot 2-D gradient-echo-planar imaging TR = 3 s, TE = minimum full, flip angle = 90° , matrix size = 64×64 , 60 slices, voxel size = $3 \times 3 \times 3 \text{ mm}^3$, 180 volumes. Participants viewed a fixation cross for the duration of the 9-min fMRI resting-state scan. A fluid-attenuated inversion-recovery (FLAIR) scan with the same orientation as the functional scans was collected to improve co-registration between subject-specific structural and functional scans.

3.3.3 Questionnaires

We administered a battery of questionnaires to comprehensively assess a diverse range of trait-level individual differences that have been previously related to mind wandering. These questionnaires captured the trait-like features of participants' psychological states, particularly aspects of well-being. The complete details of the questionnaires are presented in Appendix A.1.

3.3.4 Behavioural testing sessions

The trait profiles captured by the questionnaires were complemented by measures of performance on a range of cognitive tasks. Behavioural tasks were selected to measure a broad range of cognitive attributes, including semantic and episodic memory, executive control, fluency, and creativity. These measures were assessed

in three sessions. Each session began with a task to index the content and form of mind wandering (0-back/ 1-back task), followed by the other cognitive measures. The order of sessions and the order of tasks was counterbalanced across individuals. Details of the 0-back/ 1-back task are presented in the following paragraph. The complete details of the other cognitive tasks are described in Appendix A.2.

Using a block design, we assessed the contents of experience during mind wandering in the context of a simple task that manipulated working memory load (see Konishi et al., 2015; Medea et al., 2016, for prior published examples of this task). This task was performed at the beginning of each laboratory session to minimise participant fatigue. Measuring experience over 3 days provided us with a more comprehensive description of participants' trait-level mind wandering than would have been possible in a single experimental session.

In both tasks, participants completed target and nontarget trials. In nontarget trials, a pair of shapes appeared on screen; the two shapes were separated by a vertical line. The pairs consisted of a circle and a square, a circle and a triangle, and a square and a triangle, each in two different left/right configurations for a total of six possible pairs. Following an unpredictable sequence of nontarget trials, a target trial was presented in which participants had to make a manual response. The target was a small stimulus presented in either blue or red across conditions, with the colour counterbalanced across participants. In the 0-back condition, the target was flanked by one of two shapes, and participants had to indicate which shape matched the target shape by pressing the appropriate button. In the 1-back condition, the target was flanked by two question marks, and participants had to respond depending on which side the target shape had been on during the prior trial. Responses were made using the left and right arrow keys. Presentation times for fixation crosses ranged from 1.3 to 1.7 s in steps of 0.05 s, and nontarget presentation times varied from 0.8 to 1.2 s in steps of 0.05 s. Target presentation times always ranged from 2.1 to 2.5 s in steps of 0.05 s, and a response from participants did not end the target presentation.

There were eight blocks in one session, and each block consisted of two to four miniblocks. Each block contained either the 0-back or 1-back condition. The change of task was signalled by the presentation of the word SWITCH,

which remained on screen for 5 s. The order of the tasks was counterbalanced across participants, and the eight blocks lasted around 25 min. In each mini-block, there was one target trial, and the number of nontarget trials preceding the targets varied between one and six. Participants' performance was measured by their efficiency, which was calculated by dividing their average response time by their accuracy. For ease of interpretation, efficiency scores were reversed, so that higher scores indicated better performance.

In order to sample different features of participants' ongoing experiences, we used multidimensional experience sampling (MDES; Medea et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013; Smallwood et al., 2016). This technique uses self-report data to assess the contents of experience on a number of dimensions. The first thought probe asked participants to rate their level of task focus (My thoughts were focused on the task I was performing) on a sliding scale from 0 (*completely off task*) to 1 (*completely on task*). Participants then answered 12 randomly presented questions regarding the content and form of their experience at the moment just before they answered the first thought probe (on level of task focus). These questions (described in Table 3.1) were based on those used in prior studies adopting this approach to measure self-generated thought (Medea et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013; Smallwood et al., 2016). At the moment of target presentation, there was a 20% chance of a thought probe being presented instead of a target, with a maximum of one probe per block of the 0-back/1-back task. In each session, an average of 14.07 ($SD = 3.30$, range = 6–25) MDES probes occurred; in the 0-back condition, an average of 7.02 ($SD = 2.36$, range = 2–14) MDES probes occurred; and in the 1-back condition, an average of 7.04 ($SD = 2.24$, range = 1–15) occurred. In total, we sampled 7,006 examples of experience in this study. We calculated the mean scores of each question across the three sessions for each participant. The MDES scores were first transformed into z scores for mean-centring and univariance scaling. The scores described the average momentary experience in each dimension. We used this score in the multivariate analysis later.

Table 3.1. Experience-Sampling Questions in the 0-Back/1-Back Task.

Dimension	Question	Response scale	
		0	1
Focus	My thoughts were focused on the task I was performing.	<i>Not at all</i>	<i>Completely</i>
Future	My thoughts involved future events.	<i>Not at all</i>	<i>Completely</i>
Past	My thoughts involved past events.	<i>Not at all</i>	<i>Completely</i>
Self	My thoughts involved myself.	<i>Not at all</i>	<i>Completely</i>
Other	My thoughts involved other people.	<i>Not at all</i>	<i>Completely</i>
Emotion	The content of my thoughts was:	<i>Negative</i>	<i>Positive</i>
Images	My thoughts were in the form of images.	<i>Not at all</i>	<i>Completely</i>
Words	My thoughts were in the form of words.	<i>Not at all</i>	<i>Completely</i>
Vividness	My thoughts were vivid as if I was there.	<i>Not at all</i>	<i>Completely</i>
Detail	My thoughts were detailed and specific.	<i>Not at all</i>	<i>Completely</i>
Habit	This thought has recurrent themes similar to those I have had before.	<i>Not at all</i>	<i>Completely</i>
Evolving	My thoughts tended to evolve in a series of steps.	<i>Not at all</i>	<i>Completely</i>
Deliberation	My thoughts were:	<i>Spontaneous</i>	<i>Deliberate</i>

3.3.5 Neuroimaging data preprocessing and analysis

3.3.5.1 Resting-state fMRI.

Functional and structural data were preprocessed and analysed using the Oxford Centre for Functional MRI of the Brain’s (FMRIB’s) Software Library (FSL Version 4.1, <http://www.fmrib.ox.ac.uk/fsl>). Individual FLAIR and T1-weighted structural brain images were extracted using FSL’s Brain Extraction Tool (BET). Structural images were linearly registered to the MNI152 template using FMRIB’s Linear Image Registration Tool (FLIRT). The resting-state functional data were preprocessed and analysed using FSL’s FMRI Expert Analysis Tool (FEAT). The individual-subject analysis involved motion correction using FSL’s MCFLIRT, slice-timing correction using Fourspace time-series phase shifting, high-pass temporal filtering (Gaussian-weighted least-squares straight-line fitting, $\sigma = 200$ s), and Gaussian low-pass temporal filtering ($\sigma = 2.8$ s). In addition, we regressed out six motion parameters (as estimated by MCFLIRT) and regressing out cerebrospinal fluid and white-matter signal (top five components in the principal component analysis, PCA; CompCor method). No spatial smoothing and no global signal regression were applied.

3.3.5.2 Network-strength analysis.


To describe the functional architecture of the DMN, we transformed the resting-state blood-oxygen-level-dependent time series into connection strength values of the selected regions for each participant. The regions of interest (ROIs) were obtained from connectivity-based functional parcellation studies of the DMN

by Bzdok and colleagues (Bzdok, Langner, et al., 2013; Bzdok et al., 2015, 2016; Eickhoff, Laird, Fox, Bzdok, & Hensel, 2016; Eickhoff et al., 2016). There were 16 selected target network nodes, including subregions located in the bilateral temporoparietal junction (TPJ), ventromedial prefrontal cortex (vmPFC), dorsomedial prefrontal cortex (dmPFC), and posteromedial cortex (PMC; see Figure 3.2a). The ROI masks and the related functional-connectivity network produced with Neurosynth core tools (http://github.com/neurosynth/neuro_synth) can be found on NeuroVault (<http://neurovault.org/collections/2275/>). First, we extracted and then averaged the time series of all voxels within the 6-mm sphere masks of the given regions. Second, we created 16×16 symmetrical correlation matrices representing the network of the regions that was computed for all the individual subjects. The off-diagonal of each correlation matrix contained 120 unique region-region connection strengths. This approach provided a measure of connection strength of the region-region coupling of the DMN for each participant.

3.3.5.3 Multivariate pattern analysis.

We performed a sparse canonical-correlation analysis (SCCA) on the connection strength data and MDES scores to yield different dimensions that simultaneously described neural organisation and experience. Canonical correlation analysis (CCA) is an advanced multivariate technique that identifies distinct components between two variables spaces (Hardoon et al., 2004)—in our case, brain-region connection-strength values and experiential reports obtained through MDES. This modelling approach allows linear combinations of the two variable vectors with correlations among variables to be determined and, unlike in PCA and independent component analysis, produces dimensions in which the biological data are simultaneously constrained by psychological measures (and vice versa). To enhance the interpretability of the decomposition solutions, we used a variant of CCA penalised by L_1 regularisation, SCCA (see Hastie, Tibshirani, & Wainwright, 2015). This was achieved by setting a maximum number of brain or behaviour variables to exactly zero, which resulted in a regularised version of the singular value decomposition. A reliable and robust implementation of the SCCA method was retrieved as an R package from CRAN (penalized multivari-

C



a

b





Figure 3.2. Results of the sparse canonical-correlation analysis.

The regions of interest (ROIs) of the default mode network (DMN) from which the network connection strength was calculated are shown in (a). The correlation between the two canonical components (positive-habitual experience, 1, and spontaneous off-task experience, 2) and the two analyses is shown in (b). The analyses were restricted temporal sampling (RTS), which describes the canonical components produced when the data from 1 day of each participant were randomly removed from the decomposition, and bootstrapping (BOOTS), which describes the solution produced using bootstrapping. Panel (c) shows the results of the sparse canonical-correlation analysis (SCCA) conducted on the network-connection-strength values of key nodes of the DMN at rest and self-reports of experience during a laboratory task. Results are shown separately for the two components of experience for each analysis. Also shown are the common findings between the two analyses. The numbers indicate the subregions of each ROI, as indicated in (a). For the questions associated with each self-report dimension, see Table 1.

ate analysis, or PMA). In the current analysis, the L_1 penalty was set to 0.3 on resting-state functional connectivity and to 0.5 for the MDES results. Other parameters were set to the default. In this way, our analysis performed low-rank (i.e., described an overall network pattern by a parsimonious set of connectivity causes), conjoint (i.e., respected variance in brain and behaviour at once), and sparse (i.e., automatically found unimportant variables) decomposition of experiential and neural data.

3.3.5.4 Stability analyses.

We performed two analyses to assess the stability of the solutions produced by SCCA. First, for each participant, we excluded the MDES data of 1 random day and then recalculated the average scores for these questions. We repeated the decomposition on this new set of MDES data and the network connection strength. This corroborative quantitative assessment provided insight into the robustness of the obtained findings by a permutation analysis that left 1 day out at a time. In particular, this procedure addressed whether either the first day (when participants may be learning how to respond to the experience-sampling method) or the last day (when participants may have lower levels of motivation) might unduly bias the decomposition solutions. We reasoned that if the average momentary MDES responses are stable across three sessions, then they should yield similar latent components. Second, we acquired bootstrap samples as a permutation analysis to estimate the variance and generalisability of the sample to the population. The bootstrap resamples, each reflecting an alternative data sample that we could have obtained from the same distribution, was created by random sampling with replacement. The identical SCCA computation was then reiterated individually on each of the 1,000 perturbed versions of the actual data sample. This approach enables quantitative assessment of the quality of the original SCCA estimates by inferring confidence intervals (see Figure A.1 in Appendix A.3 for the distributions). We selected latent components that were consistent across the decomposition of the original sample, a leave-1-day-out sample, and a bootstrap sample, as those are the stable components that were less biased by the session effect and closer to our best estimation of population. We formalised the similarity of these two types of resampling by conducting a

formal conjunction of the solutions generated through these different methods of resampling. To quantify the similarity between the components, we performed a conjunction that highlights the common elements of each solution. The feature conjunctions were calculated as follows:

$$\text{conjunction} = \begin{cases} 1, & \text{when } \sqrt{\text{weight}_{\text{LODO}} \times \text{weight}_{\text{BOOTS}}} > 0.1 \\ 0, & \text{otherwise.} \end{cases} \quad (3.1)$$

where LODO refers to leave 1 day out and BOOTS refers to bootstrapping. In addition, because bootstrapping produces a population estimation of our sample, we used the latent component weights produced by this method to compute component scores. This set of scores was used in all subsequent analyses. The source code for this analysis is available at <https://github.com/htwangtw/DimensionsOfExperience>.

3.3.5.5 Whole-brain analysis.

A limitation in our analysis is that we focused on the DMN to describe patterns of thought. To overcome this limitation, we generalised the types of experience provided by the SCCA by assessing their associations with areas outside of the DMN using a process conceptually similar to dual regression (Beckmann, Mackay, Filippini, & Smith, 2009). To perform these analyses, we preprocessed and analysed the resting-state functional data using FEAT. For the individual-subject preprocessing procedure, see the Resting-State fMRI section.

Following these preprocessing steps, we used a mask produced by the average of the DMN ROIs to determine the time series that described this neural system. This time series was used in a whole-brain functionality analysis for each participant. This allowed us to produce a subject-specific spatial map based on the selected ROIs, and these maps were used as dependent measures in our group-level analysis. To test whether the functional connectivity of the DMN ROIs was associated with the canonical components, we conducted a group-level analysis using FMRIB’s Local Analysis of Mixed Effects Stage 1 (FLAME 1). To control for spurious correlations that might emerge from movement, we included the two canonical components on thought reports only, group mean and Jenkinson’s mean framewise displacement (FD Jenkinson, Bannister, Brady, & Smith,

2002), as explanatory variables in the full model. The Jenkinson’s mean FD was calculated by the motion power statistic function in Configurable Pipeline for the Analysis of Connectomes (C-PAC; <https://fcp-indi.github.io/>). A 50% probabilistic gray-matter mask was applied to the results maps, and the results were thresholded at the whole-brain level using cluster-based Gaussian random-field theory, with a cluster-forming threshold (Z) of 2.6 and a familywise-error-corrected cluster significance level (p) of .05. Unthresholded maps were uploaded onto Neurovault (<http://neurovault.org/images/43189/>).

3.3.5.6 PCA.

To summarise the questionnaire and task data, we performed an initial data-reduction step using PCA in SPSS (Version 24). This analysis was performed separately for the questionnaires and task measures. One hundred forty-five participants’ data were included in the analysis of the questionnaire items, and 157 participants’ data were included in the analysis of the behavioural tasks. The behavioural-task measures were converted into z scores to avoid data distortions derived from the difference in score means. Missing data were imputed by mean scores in both analyses. The Kaiser-Meyer-Olkin (KMO) measure and Bartlett’s test of sphericity were used to measure the sampling adequacy of the model. Components were selected on the basis of the elbow in a scree plot (see Figure A.2 in Appendix A.3), and varimax rotation was used to maximise the distinctiveness of each solution.

In the PCA of the phenotypical variation measured by behavioural tasks, Bartlett’s test of sphericity was significant, $\chi^2(210) = 775.01$, $p < .001$, which indicates that it is appropriate to apply PCA to these data. The KMO measure of sampling adequacy indicated that the current sample was acceptable for PCA (KMO = 0.79). The PCA of task performance revealed three principal components with a clear elbow after the third component observed in the scree plot. The three orthogonal components accounted for 40.7% of the total variance; the component loading patterns are shown in Figure 3.3a. The three components, which accounted for 24.9%, 8.3%, and 7.5% of the variance, respectively, can be interpreted as the three aspects of cognitive functioning: (a) semantic memory, (b) executive control, and (c) the generation of information (including letter or

category fluency and the generation of creative solutions).

In the PCA of the questionnaire data, Bartlett’s test of sphericity was significant, $\chi^2(105) = 919.78$, $p < .001$, which indicates that PCA is an appropriate model for the data. The KMO measure of sampling adequacy indicated that there were strong relationships among the variables (KMO = 0.82). The application of PCA to the questionnaire data revealed four components with a clear elbow after the fourth component observed in the scree plot in Fig. S2. The four orthogonal components accounted for 65% of the total variance (produced component loading patterns are shown in Figure 3.3b). The four components accounted for 35%, 14%, 9%, and 7% of the variance, respectively. The first component, affective disturbance, was anchored at one end by high levels of depression and rumination and at the other by high levels of well-being. The second component was associated with high scores on four of the five autism subscales, excluding the attention-to-detail subscale. The third component loaded on components of both attention-deficit/hyperactivity disorder (ADHD) and dyslexia. The fourth component loaded on trait anxiety and high levels of attention to detail, as measured by the Autism Spectrum Quotient (Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001). We analysed these data using a multivariate analysis of variance (MANOVA) in which the dependent variables were the PCA loadings produced by the decomposition of the questionnaires, and the independent variables were the canonical component loadings.

3.4 Results

3.4.1 Determining consistent categories of experience

We applied SCCA to the network-connection-strength values among ROIs in the DMN and the average scores on the experiential reports gained in the laboratory. We accepted 13 canonical components generated by SCCA (see Figure A.3 in Appendix A.3 for the complete set). Of these initial components, two were consistent when we randomly removed the MDES reports of 1 day per participant and when bootstrapping was used to provide a more comprehensive description of the sample (see Section 3.3). The consistency of these patterns across the three different analyses indicates that, in qualitative terms, they were not




Figure 3.3. Results from principal component analyses of (a) behavioural tasks and (b) questionnaires.

In the analysis of behavioural tasks, the components were semantic memory (SEM), executive control (EXE), and the generation of information (GEN). In the analysis of questionnaire data, the components were affective disturbance (AD), social interaction (SOC), dyslexia (DYSL), and attention to detail (ATT). The heat map indicates the loadings of each measure. In (a), an asterisk indicates that measures were relatedness tasks from a semantic battery. In (b), a single asterisk indicates measures drawn from the World Health Organization Quality of Life assessment (WHO, 2002), and two asterisks indicate measures drawn from the Autism Spectrum Quotient (Baron-Cohen et al., 2001). ADHD = attention-deficit/hyperactivity disorder; RAPM = Raven's Advanced Progressive Matrices (Raven et al., 1998). For the scree plots describing the eigenvalues for each dimension, refer to Figure A.2 in Appendix A.3.

unduly biased by a particular session of our study and were likely to provide adequate estimation of the population (Figure 3.2b). These stable components are presented in Figure 3.2b, in which we show both the bootstrapping results, the analysis that randomly excluded one session (restricted temporal sampling), and the common elements of each solutions.

Canonical Component 1 reflects a pattern of stronger coupling within the mPFC, as well as between the left inferior parietal cortex (subregion 2 in the TPJ; see Figure 3.2a). This pattern of integration within key nodes of the DMN was associated with descriptions of experience as positive, evolving, and habitual. We will refer to this as *positive-habitual* experiences. Canonical Component 2 was associated with relatively weak patterns of coupling between the pCC bilaterally (subregions 2 and 4 in the TPJ; see Figure 3.2a) and regions of the mPFC (subregions 1, 5, and 6 in the vmPFC; see Figure 3.2a). This component was associated with thoughts that were task unrelated and nondeliberate. We will refer to this component as *spontaneous off-task* experiences.

3.4.2 Validating the categories of experience

Having identified two reliable dimensions of neurocognitive experience, we tested whether these patterns accounted for additional variance in the measures that we collected in our experiment. We first conducted a whole-brain analysis to determine whether the different patterns of experience were associated with differential communication from the DMN to other areas of the brain. In this analysis, we first employed dual regression to calculate the subject-specific spatial maps describing the correlation of the DMN and the whole brain and then used these spatial maps as dependent measures in a group-level multiple regression in which participants' variation in positive-habitual and spontaneous off-task experiences were both explanatory variables of interest (see Section 3.3). This analysis revealed a pattern of regions in which connectivity was differentially related to the dimensions of *positive-habitual* and *spontaneous off-task* experiences. These regions were the left temporoparietal cortex, left hippocampus/entorhinal cortex, left lateral middle temporal gyrus, and the left pre-supplementary region. Extracting the connectivity in this network and plotting these against the different types of experience revealed that these regions showed a pattern of connectivity

that was linked to the expression of positive-habitual experiences but was unrelated to levels of spontaneous off-task experiences. These data are consistent with those found in previous studies that show that medialtemporal connectivity with the DMN is linked to aspects of spontaneous experience, such as episodic thought (Karapanagiotidis et al., 2017), and on-line studies that show that activity in this region is important during mind-wandering states (e.g., Ellamil et al., 2016). It also confirms theoretical accounts of states of mind wandering as relying on regions that fall outside of the core of the DMN, such as the pre-supplementary motor area (pre-SMA; Christoff et al., 2016).

Next, we explored whether the different canonical components had specific implications for performance on the tasks in which we assessed experience (i.e., the 0-back/1-back task). Because the SCCA depends on resting-state data recorded independently of the task, we were unable to estimate the canonical components separately for each task. Consequently, in these analyses, we explored whether overall differences in canonical component loadings across participants were associated with performance efficiency on the 0-back/ 1-back task. We used a repeated measures analysis of variance in which the dependent variable was the efficiency with which participants performed the 0-back/ 1-back task, respectively. This analysis revealed a significant interaction between task efficiency and variation in our spontaneous-off task component, $F(1, 154) = 6.43$, $p = .012$, $\eta_p^2 = .04$. Decomposition of this interaction showed that participants who scored higher on spontaneous off-task experience performed better on the 0-back condition, $b = 0.06$, 95% confidence interval (CI) = [0.01, 0.11], $t(151) = 2.38$, $p = .019$, $\eta_p^2 = .04$, and worse on the 1-back condition, $b = -0.09$, 95% CI = [-0.15, -0.02], $t(151) = -2.55$, $p = .012$, $\eta_p^2 = .04$. The differential relationship between the levels of spontaneous off-task experience and performance on the 0-back/1-back task is shown in Figure 3.4. These data confirm accounts that suggest that attentional lapses linked to mind wandering are context dependent, tending to have more negative effects as tasks become more demanding (Smallwood, Ruby, & Singer, 2013); they are also consistent with prior studies suggesting that context regulation may be more problematic for spontaneous than deliberate mind wandering (see also Seli, Risko, Smilek, & Schacter, 2016).




Figure 3.4. Relationship between the different neural-cognitive components and the laboratory and questionnaire measures.

For the whole-brain analysis, the brain diagrams show clusters of the default-mode-network mask, and the graphs show the correlation between their functional connectivity and the two experience components. For the behavioral analysis, the graphs show the relationship between the two canonical components and measures of well-being and task performance.

Finally, we used MANOVA to determine how the patterns of experience revealed by SCCA were related to the decompositions of the battery of cognitive performance and questionnaire measures. In this analysis, PCA scores describing either phenotypical variation or questionnaire measures on each of the components of cognitive function were the independent variables, and the individual loadings for each of the two canonical components describing experience from the SCCA were the dependent variables. For the analysis of phenotypical variation, this produced two significant results with the executive-control component, $F(2, 152) = 5.84$, $p = .006$, $\eta_p^2 = .065$, and the generation-of-information component, $F(2, 152) = 3.41$, $p = .007$, $\eta_p^2 = .065$. Higher loadings on the positive-habitual component, $F(1, 153) = 9.84$, $p = .002$, $\eta_p^2 = .060$, were associated with worse performance on tasks requiring executive control, $b = 0.19$, 95% CI = [0.32, 0.07], $t(153) = 3.14$, $p = .002$, $\eta_p^2 = .060$, and higher loadings on the spontaneous-off task experience component, $F(1, 153) = 10.15$, $p = .002$, $\eta_p^2 = .062$, were associated with better performance on tasks involving the generation of information (such as creativity), $b = 0.20$, 95% CI = [0.08, 0.33], $t(153) = 3.19$, $p = .002$, $\eta_p^2 = .062$. This indicates that two of the experiential components identified by the SCCA were uniquely associated with poor performance on executively demanding tasks and better performance on measures of creativity: both aspects of psychological functioning that have previously been linked to mind wandering (e.g., Baird et al., 2012; McVay et al., 2009). The relationships for both neurocognitive dimensions are shown in Figure 3.4.

In terms of the relationship to the questionnaire decomposition, we found a significant association with the first principal component, $F(1, 151) = 3.76$, $p = .026$, $\eta_p^2 = .05$, which captured affective disturbance. This revealed two significant relationships: (a) a strong association with the positive-habitual component, $F(1, 152) = 6.13$, $p = .014$, $\eta_p^2 = .04$, which suggests a negative association between positive-habitual thought and levels of affective disturbance, $b = 0.16$, 95% CI = [-0.29, 0.03], $t(152) = 2.48$, $p = .04$, $\eta_p^2 = .062$, and (b) an association with the spontaneous-off-task-experience component, $F(1, 152) = 4.55$, $p = .035$, $\eta_p^2 = .03$, which suggests that higher loadings on this component were associated with higher levels of affective disturbance, $b = 0.15$, 95% CI = [0.11, 0.28], $t(152) = 2.13$, $p = .035$, $\eta_p^2 = .03$. This analysis demonstrates that the dif-

ferent canonical components have dissociable associations with respect to well-being, capturing aspects of the bidirectional relationship between the mind-wandering state and affective disturbance highlighted by prior research (e.g., Killingsworth & Gilbert, 2010; Ruby, Smallwood, Engen, & Singer, 2013). Importantly, our analysis demonstrates that the different canonical components have dissociable associations with respect to well-being, which shows that our method captured both elements of the apparently contradictory analysis linking the mind-wandering state to well-being that has been highlighted by prior research.

One concern with resting-state functional connectivity arises from the possibility that the connectivity matrices are unduly affected by individual differences in motion (Power et al., 2014). Consistent with this possibility, our results showed a correlation at the group level between the positive-habitual component, $r(155) = .363$, $p < .001$, but not the spontaneous-off-task-experience component, $r(155) = .097$, $p = .229$. Hence, we assessed the contribution of this association to our results linking positive-habitual thought to our measured phenotypes. We performed a series of stepwise analyses to identify the contribution that motion made to the phenotypical associations with positive-habitual thought. In these analyses, the canonical component was the dependent variable. We entered the principal components describing cognition or well-being in the first step and Jenkinson’s mean FD in the second step. Including motion significantly improved the predictive value of the model for well-being— Model 1: $R^2 = .06$, $F(4, 152) = 2.21$, $p = .07$, $\eta_p^2 = .06$; Model 2: $R^2 = .19$, $F(5, 151) = 6.95$, $p < .001$, $\eta_p^2 = .19$; model change: $R^2 = .13$, $F(1, 151) = 24.51$, $p < .001$ —as well as of the model for cognition: Model 1: $R^2 = .07$, $F(3, 153) = 3.92$, $p = .010$, $\eta_p^2 = .07$; Model 2: $R^2 = .18$, $F(4, 152) = 8.22$, $p < .001$, $\eta_p^2 = .18$; model change: $R^2 = .11$, $F(1, 152) = 19.65$, $p < .001$. In the case of well-being, the explained variance of the affective disturbance component was not improved with the inclusion of motion— Model 1: affective-disturbance $\beta = -0.20$, $t(152) = 2.48$, $p = .014$, $\eta_p^2 = .04$, 95% CI = [0.29, 0.03]; Model 2: affective-disturbance $\beta = 0.20$, $t(151) = 2.59$, $p = .011$, $\eta_p^2 = .05$, 95% CI = [0.28, 0.03], Model 2: mean-FD $\beta = 0.36$, $t(151) = 4.94$, $p < .001$, $\eta_p^2 = .14$, 95% CI = [3.29, 7.67]. Thus, the relationship between affective disturbance and positive-habitual thought remained

largely unchanged by the inclusion of motion as a nuisance variable. In the case of cognition, executive control accounted for less variance in the positive-habitual component when mean FD was included— Model 1: executive-control $\beta = 0.24$, $t(153) = 3.14$, $p = .002$, $\eta_p^2 = .06$, 95% CI = [0.32, 0.07]; Model 2: executive-control $\beta = 0.16$, $t(152) = 2.17$, $p = .032$, $\eta_p^2 = .03$, 95% CI = [0.25, 0.01]; Model 2: $\beta = 0.34$, $t(152) = 4.43$, $p < .001$, $\eta_p^2 = .11$, 95% CI = [4.82, 12.56].

Unlike in the well-being analysis, motion explained a substantial amount of variance that was shared in the relationship between executive control and positive-habitual thought. To explore whether the positive-habitual component reflected an artefact of motion, we selected participants for whom movement greater than 0.2 mm occurred on less than 5% of the resting-state data ($N = 134$) and reran the SCCA with the identical pipeline. This produced similar solutions for both positive-habitual and spontaneous off-task thought (see Figure A.4 in Appendix A.3). Importantly, positive-habitual thought was not significantly correlated with motion, $r(132) = .10$, $p = .236$, but was correlated with poor executive control, $r(155) = -.26$, $p = .001$ (see Table A.1 in Appendix A.3 for the full set of correlations). This final analysis shows that in a more restricted sample in which motion did not correlate with either latent component, we still observed a relationship between positive-habitual thought and poor executive control.

3.5 Discussion

Using multivariate pattern analysis, our study demonstrated that the content of the mind-wandering state is heterogeneous and confirmed hypotheses that different types of experience have differing functional associations (Smallwood, Ruby, & Singer, 2013). Using a novel analysis strategy, we simultaneously decomposed self-reports of experience with descriptions of neural organisation, revealing dimensions of experience with unique phenotypical associations: positive-habitual experiences and spontaneous off-task thoughts.

Poor executive control, a well-documented association of mind wandering (McVay et al., 2009), predicted variation in positive-habitual thoughts. This pattern of thinking was linked to coupling in the mPFC, a region important for assigning value to neural signals (Roy, Shohamy, & Wager, 2012). It is

possible that deficits in executive control during mind wandering emerge because of problems in assigning value to an external task, a view supported by evidence that financial motivation limits the impact of mind wandering on performance (Mrazek et al., 2012). We found that spontaneous off-task experiences simultaneously underlie the association between mind wandering and tasks of creativity (Baird et al., 2012), as well as problems in performing tasks requiring continuous monitoring of external information. Finally, while positive-habitual experiences are linked to improved well-being, spontaneous off-task experiences are associated with increased affective disturbance, which captures the apparent contradiction that mind wandering can be associated with both negative (e.g., Killingsworth & Gilbert, 2010) and positive (e.g., Poerio et al., 2016) emotional states. Together, these data provide the most convincing evidence to date that experience during mind wandering unfolds along a set of underlying dimensions and that these explain many of the phenotypical associations that have hitherto been associated with the mind-wandering state (Smallwood, Ruby, & Singer, 2013).

Our study also demonstrates the complex contribution that the DMN makes to cognition. Strong DMN connectivity at rest was associated with an increased tendency for positive-habitual thoughts about the future, which corroborates previous research linking the DMN to mental time travel (Karapanagiotidis et al., 2017; Schacter et al., 2007). Participants also rated these experiences as habitual, a pattern that supports accounts of the role of the DMN in cognition as emphasising automatic influences during mind wandering (Christoff et al., 2016). Spontaneous off-task thoughts, in contrast, showed weaker integration between core DMN regions and were linked to poor performance in the 1-back condition, a context in which task performance depends on the DMN functioning as a coherent network (Konishi et al., 2015). More generally, we found that states of high connectivity within the DMN (positive-habitual thoughts) were associated with more functional coupling to regions outside of the core network—a key prediction of the view that activity within the DMN reflects the integration of information from across the cortex (Margulies et al., 2016). It is important to note that our analysis shows that the behavior of the DMN at rest contains information about individual variation in the type of experiences that emerge

during mind wandering. These data should not be taken as evidence that this system is exclusive in its role in mind wandering. Indeed, our whole-brain regression provides quantitative evidence that the interactions of the DMN with other regions, including those in the medial temporal lobe and the executive system (e.g., pre-SMA), are also important. In this way, our study supports recent theoretical perspectives (e.g., Christoff et al., 2016; Margulies et al., 2016), as well as prior empirical results (e.g., Ellamil et al., 2016; Golchert et al., 2017; Smallwood et al., 2016) highlighting that regions other than the DMN core are important for mind wandering.

There are a number of limitations of the current analysis. First, our study focused on describing mind wandering as a trait. Prior work has shown similarities between state and trait measures of mind wandering in terms of (a) neural processing (e.g., trait: Smallwood et al., 2016; state: Christoff et al., 2009; Stawarczyk, Majerus, Maquet, & D’Argembeau, 2011) and (b) psychological processes such as working capacity (e.g., trait: McVay et al., 2009; state: Mrazek et al., 2012) and happiness (e.g., trait: Ruby, Smallwood, Engen, & Singer, 2013; state: Killingsworth & Gilbert, 2010). Nonetheless there are certain aspects of mind wandering that can be understood only by treating it as a state, such as its temporal features (Christoff et al., 2016). Second, our study measured mind wandering in the laboratory. Although there is a correspondence between mind wandering in laboratory and naturalistic settings (e.g., McVay et al., 2009), its form and content may depend on the contexts in which the experience emerges. Consequently, our findings should be supplemented by studies examining the occurrence of different types of experience in ecologically valid settings. Finally, our study did not find evidence for links with tasks that rely on semantic memory or for links to psychological traits other than well-being. This may have been due to our selection of neural regions or from our selection of questions. Prior studies have linked regions in the temporal lobe to the contents of thought (e.g., Smallwood et al., 2016), a pattern of data that is consistent with a role of the semantic system in spontaneous thought (Binder et al., 2009). Other work has highlighted awareness of mind wandering as important in traits such as hyperactivity (Franklin et al., 2017). We anticipate that extending the selected regions of the cortex and the aspects of experience measured may extend

our understanding of the mind wandering state to encompass forms of semantic processing and additional psychological traits.

In closing, our study provides the strongest evidence to date that the mind-wandering state is heterogeneous in its content, neural basis, and functional associations. We describe two neurocognitive dimensions capturing associations with attentional lapses, creativity and well-being, confirming much of the research on mind wandering conducted over the last decade. However, we also provide an explanation for why scientific accounts of mind wandering have been dominated by controversy, such as its relationship to happiness (Killingsworth & Gilbert, 2010), creativity (Smeekens & Kane, 2016), executive control (McVay et al., 2009), and the DMN (Gilbert, Dumontheil, Simons, Frith, & Burgess, 2007). Our data suggest that these debates emerge from an erroneous assumption that mind wandering is a unitary psychological construct, when it is in fact made up of distinct states with unique neural correlates and functional associations. This ontological uncertainty has led to artificial controversies that hinder the development of a mature science of internal experience. Although our findings do not capture the full range of experiential dimensions on which the mind can wander, they convincingly demonstrate that it is untenable to characterise mind wandering as a uniform experience. As a discipline, we must embrace methodologies and analytical techniques that capture the complex nature of internal experiences, allowing researchers to accurately determine the contribution that they make to people's lives.

Chapter 4

Population Variation in the Associations Between Large-Scale Networks and Experiences at Rest

The following chapter has been adapted from: Wang, H.-T., Bzdok, D., Margulies, D., Craddock, C., Milham, M., Jefferies, E., & Smallwood, J.(2018). Patterns of thought: Population variation in the associations between large-scale network organisation and self-reported experiences at rest. *NeuroImage*, 176(1), 518–527. doi: 10.1016/j.neuroimage.2018.04.064 ¹

4.1 Abstract

Contemporary cognitive neuroscience recognises unconstrained processing varies across individuals, describing variation in meaningful attributes, such as intelligence. It may also have links to patterns of ongoing experience. This study

¹ J. Smallwood, and H.-T. Wang designed the study. The data was provided from C. Cameron and M. Milham. The analysis pipeline was constructed by H.-T. Wang under the supervision of D. Bzdok and J. Smallwood. Data were analyzed by H.-T. Wang, under the supervision of D. Bzdok, J. Smallwood, and E. Jefferies. H.-T. Wang and J. Smallwood drafted the manuscript. D. Bzdok, D. Margulies and E. Jefferies provided critical input to the interpretations. All the authors approved the final version of the manuscript prior to submission.

examined whether dimensions of population variation in different modes of unconstrained processing can be described by the associations between patterns of neural activity and self-reports of experience during the same period. We selected 258 individuals from a publicly available data set who had measures of resting-state functional magnetic resonance imaging, and self-reports of experience during the scan. We used machine learning to determine patterns of association between the neural and self-reported data, finding variation along four dimensions. ‘Purposeful’ experiences were associated with lower connectivity—in particular default mode and limbic networks were less correlated with attention and sensorimotor networks. ‘Emotional’ experiences were associated with higher connectivity, especially between limbic and ventral attention networks. Experiences focused on themes of ‘personal importance’ were associated with reduced functional connectivity within attention and control systems. Finally, visual experiences were associated with stronger connectivity between visual and other networks, in particular the limbic system. Some of these patterns had contrasting links with cognitive function as assessed in a separate laboratory session—purposeful thinking was linked to greater intelligence and better abstract reasoning, while a focus on personal importance had the opposite relationship. Together these findings are consistent with an emerging literature on unconstrained states and also underlines that these states are heterogeneous, with distinct modes of population variation reflecting the interplay of different large-scale networks.

4.2 Introduction

Unconstrained processing reflects important population level variation in measures of cognition, affect, and demographic lifestyle factors. Psychological studies show that almost a third of ongoing thought is unconstrained by events in the here-and-now (Killingsworth & Gilbert, 2010) with important links to cognitive and affective processing (Mooneyham & Schooler, 2013). In neuroscience, metrics defined from the brain during wakeful rest, describe the organisation of neural function at both the micro and macro scale (Glasser et al., 2016; Margulies et al., 2016). They also reflect individual differences in cognitive function (Finn et al., 2015), psychiatric conditions (Nooner et al., 2012) and demographic lifestyle

factors (Smith et al., 2015). These findings establish unconstrained neurocognitive processing as a core element of human cognition, highlighting the need to formally understand the underlying neural architecture, and the associated patterns of experience.

One perspective on unconstrained processing emphasises the role of memory, with contributions of conceptual and episodic representations to ongoing thought (Binder et al., 2009; Gusnard et al., 2001). Psychological studies have shown that patterns of unconstrained processing have links with memory retrieval, creativity and planning (Baird et al., 2012; Leszczynski et al., 2017; Medea et al., 2016; Poerio et al., 2017). Such evidence raises the possibility that episodic representations anchored in the medial temporal lobe (Moscovitch, Cabeza, Winocur, & Nadel, 2016) or conceptual representation anchored in anterior temporal lobe (Lambon-Ralph et al., 2017) contribute to ongoing thought (Smallwood et al., 2016). It is hypothesised that these systems' contribution to unconstrained states may be linked to the ability for these regions to become functionally decoupled from systems more directly involved in action and perception, allowing them to operate in an offline manner (Smallwood, Ruby, & Singer, 2013). This process of decoupling may also be important in neural systems closely allied to those involved in memory – the default mode network (Raichle et al., 2001). These regions of transmodal cortex are relatively distant in functional and structural space from systems involved in perception and action, potentially facilitating their role in stimulus independent aspects of cognition (Buckner & Krienen, 2013; Margulies et al., 2016; Mesulam, 1998). Together these *representational* accounts of unconstrained processing highlight default mode and limbic networks as important candidate neural systems, especially when decoupled from systems directly involved in perception and action.

Alternative perspectives on unconstrained thought emerge from links between types of ongoing experience and problems maintaining a task relevant goal in mind. This *executive-failure* view (Kane & McVay, 2012; McVay et al., 2009) takes as a starting point evidence that patterns of ongoing thought, such as the experience of mind-wandering, are linked to problems on tasks including sustained attention (McVay et al., 2009) and measures of general aptitude and executive control (Mrazek et al., 2012). Task-based neuroimaging investigations

highlight a network of regions that increase their activity across many different task situations, so called multiple demand regions (Duncan, 2010). These regions broadly correspond to three well described intrinsic networks: ventral attention, dorsal attention, and frontal-parietal networks. Since these systems are important for the effective performance of many different tasks then dysregulation within these systems could reflect the hypothesised *executive-failure* contribution to aspects of ongoing thought (McVay et al., 2009; Weissman et al., 2006).

Other aspects of unconstrained processing could reflect the importance of affective processes, or different modalities of processing. Ongoing thought is linked to mood state: Experimental inductions of mood (Smallwood, Fitzgerald, Miles, & Phillips, 2009; Smallwood & O'Connor, 2011) as well as natural fluctuations (Poerio, Totterdell, & Miles, 2013; Ruby, Smallwood, Engen, & Singer, 2013) impact on ongoing thought. Contemporary accounts of emotional processing emphasise the role of limbic regions including the amygdala (Bzdok, Laird, Zilles, Fox, & Eickhoff, 2013; Lindquist, Wager, Kober, Bliss-Moreau, & Barrett, 2012) and anterior aspects the insula (Touroutoglou, Hollenbeck, Dickerson, & Barrett, 2012), suggesting these regions may be important in determining affective aspects of ongoing thought. Psychological studies of ongoing thought also suggest that another important dimension of unconstrained processing may reflect the different modalities of processing (Konishi, Brown, Battaglini, & Smallwood, 2017; Smallwood et al., 2016). It has been shown, for example, that the visual system plays an important role in the expression of visual imagery (Ganis, Thompson, & Kosslyn, 2004; Kosslyn, Ganis, & Thompson, 2001). Recent work has extended this evidence to shown patterns of activity with visual regions are linked to the emergence of visual, non-verbal, elements of ongoing thought (Rajj & Riekk, 2017). It is also possible that sensorimotor processes may be implicated in language processing during unconstrained processing, given that a role for these regions in language processing extends beyond production (Bzdok et al., 2016; Pulvermüller & Fadiga, 2010; Pulvermüller, 2010).

Our study aimed to identify patterns of intrinsic connectivity associated with different patterns of unconstrained states and examines their neurocognitive features from the perspectives outlined above. We used a large publicly available

dataset, containing measures of resting-state functional magnetic resonance imaging (fMRI), and an accompanying self-report instrument describing cognition experienced during the resting state (Gorgolewski et al., 2014; Nooner et al., 2012). We previously explored the relationships between patterns of ongoing thought and measures of neural activity, such as the fractional amplitude of low frequency oscillations, as well as the regional homogeneity of neural activity, in a sub sample of this data set (Gorgolewski et al., 2014). In this study we focused on connectivity, we applied sparse canonical correlation analysis (SCCA) to obtain a conjoined decomposition of self-reports of experience with matrices of whole brain connectivity data. This analysis produces multivariate patterns that reflect dimensions of variation that are mutually constrained by both brain and experience. In this way we capitalise on the fact that self-reports of experience during scanning and descriptions of ongoing neural processing provide complementary descriptions of unconstrained cognition. Our analysis, therefore, helps define, at a population level, the shared links between brain patterns and different types of experience. Moreover, respecting the multivariate nature of brain and behaviour space, as our analysis does, can accommodate complex many-to-many relationships between patterns of connectivity and self-reports, and therefore is sensitive to the possibility of degeneracy in the underlying data. As a final validation step we established whether these neurocognitive dimensions are associated with performance on a battery of available cognitive tasks, including measures of executive control and intelligence.

We use the dimensions our analysis produces, and their links with cognitive function to evaluate the perspectives on unconstrained thought outlined earlier. *Representational* accounts emphasise links with neural systems involved in memory, such as the limbic system, and regions of transmodal cortex, such as the default mode network. They highlight states with lower levels of functional communication between these regions and those more directly involved in external action. In contrast, *executive-failure* accounts emphasise dysregulation in attention and control networks as contributing to patterns of ongoing thought that are linked to problems in domain general task performance. Affective accounts highlight limbic regions as important hubs in aspects of ongoing thoughts linked to emotion. Finally, modality specific influences on unconstrained thought

may depend on information codes represented in regions that specialise in that particular type of information, such as a role of visual cortex in experiences dominated by images. Notably, some views lead to dissociable predictions with respect to cognitive performance. For example, executive-failure accounts predict patterns of thoughts linked to worse performance on measures of cognitive function, while representational accounts make the opposite prediction.

4.3 Method

4.3.1 Participants

We analysed 258 participants (females = 162; age range 18 - 55, $M = 34.97$, $SD = 12.24$) obtained from the enhanced Nathan Kline Institute-Rockland sample (NKI-RS; http://fcon_1000.projects.nitrc.org/indi/enhanced/). Full details of the acquisition of this sample can be found in Nooner et al. (2012). We selected participants between 18 and 55 years old as our sample, this choice allowed us to maximise the cohesive nature of our sample. All the participants have the MRI data and less than 5 missing data points among the selected assessments.

4.3.2 Cognitive measures and questionnaires

Based on prior studies examining the links between spontaneous thought and cognitive performance (see Mooneyham & Schooler, 2013), we selected established neuropsychological measures linked to executive control, abstract reasoning and intelligence. The measures included the Delis-Kaplan Executive Function System (D-KEFS; Swanson, 2005), Wechsler Abbreviated Scale of Intelligence (WASI-II; Wechsler, 1999), and Wechsler Individual Achievement Test-Second Edition Abbreviated (WIAT-IIA; Wechsler, 2005). In D-KEFS we selected the tower test (move accuracy ratio), colour-word interference test (errors inhibition/switching), verbal fluency test (letter fluency – category fluency), design fluency test (design accuracy), trail making test (sequencing errors score + set-loss errors score + time-discontinue errors score), and the proverb test (a measure of abstract semantic reasoning). We used the rescaled score ($M = 10$, $SD = 3$) in our analysis. Tasks measures that reflected error rates (i.e. the

colour-word interference test and trail making test) were reversed, so that high rescaled scores indicated better task performance. All the scores were transformed to z-scores.

4.3.3 ongoing cognition measure

The New York Cognition Questionnaire (NYC-Q) is a self-report tool used to assess the thoughts experienced at rest (Gorgolewski et al., 2014; Sanders, Wang, Schooler, & Smallwood, 2017). It assesses thoughts and feelings experienced during the resting-state period. The first section contains 23 questions about the content of thought. These questions covers the temporal, social, emotional aspects of spontaneous thoughts that have been shown to be important by prior studies (e.g. Ruby, Smallwood, Engen, & Singer, 2013). Participants rated each question on a scale of 1 (Completely did not describe my thoughts) to 9 (Completely did describe my thoughts). The second section contains 8 questions about the forms thoughts take, capturing aspects of experience such as modality and detail associated with experience that prior studies suggest as important for spontaneous thoughts (Smallwood et al., 2016). Participants rated each question on a scale of 1 (Completely did not characterise my experience) to 9 (Completely did characterise my experience). In the current study we analysed the two sections together to provide single solutions that combined information on both the content form of experience. The full list of questions and the corresponding labels are presented in Table 4.1. The questionnaire was administrated once after the resting-state scan in order to assess experiences during the scanning session. For the full details of the NYC-Q, please refer to Gorgolewski et al. (2014). We have placed the questionnaire measure used in this study along with an example self-report collection task on GitHub at the following address: https://github.com/htwangtw/restingstate_thoughtreports.

4.3.4 MR data processing

4.3.4.1 Resting-state fMRI

We used resting-state fMRI to describe the general functional organisation of the brain. We selected resting-state multiband functional magnetic resonance imaging (R-mfMRI; TR = 1400msec; voxel size = 2mm isotropic; duration =

Table 4.1. The New York Cognition Questionnaire (NYC-Q).

Questions	Labels
I thought about things I am currently worried about	Concerns
I thought about people I have just recently met	People
I thought of people I have known for a long time (friends)	Friend
I thought about members of my family	Family
I thought about an event that took place earlier today	Today - Past
I thought about an interaction I may possibly have in the future	Social - Future
I thought about an interaction with somebody that took place in the past	Social - Past
I thought about something that happened at a place very close to me	Near Location
I thought about something that made me feel guilty	Guilt
I thought about an event that may take place later today	Today - Plan
I thought about something that happened in the recent past (last couple of days but not today)	Recent Past
I thought about something that happened a long time ago in the past	Distant Past
I thought about something that made me angry	Anger
I thought about something that made me happy	Happiness
I thought about something that made me cheerful	Cheerfulness
I thought about something that made me calm	Calm
I thought about something that made me sad	Sadness
I thought about something that is important to me	Importance
I thought about something that could still happen today	Today - Future
I thought about something that may take place in the distant future	Distant Future
I thought about something that could take place in the near future (days or weeks but not today)	Near Future
I thought about personal worries	Worries
I thought about something that happened in a place far away from where I am now	Distant Location
In the form of images:	Image
In the form of words:	Words
Like an inner monologue or audiobook:	Monologue
Like a television program or film:	Film
Had a strong and consistent personal narrative:	Narrative
Had a clear sense of purpose:	Purpose
Vague and non-specific:	Vague
Fragmented and disjointed:	Fragment

10 minutes) for our analysis. Functional and structural data were pre-processed using Configurable Pipeline for the Analysis of Connectomes (C-PAC; <https://fcp-indi.github.io/>) to interface with FMRIBs Software Library (FSL version 5.0, www.fmrib.ox.ac.uk/fsl). Individual FLAIR and T1 weighted structural brain images were extracted using Brain Extraction Tool (BET). Structural images were linearly registered to the MNI-152 template using FMRIB’s Linear Image Registration Tool (FLIRT). The resting-state functional data were pre-processed and analysed using the FMRI Expert Analysis Tool (FEAT). X, Y, Z displacement and the three axis rotations were used to calculate the mean frame displacement (FD), characterising movement of each participant during the scanning session (Power et al., 2014). Mean of the absolute values for FD were later used to account for subject specific head motion. No global signal regression was applied. The individual subject analysis involved: motion correction using MCFLIRT; slice-timing correction using Fourier space time series phase-shifting; spatial smoothing using a Gaussian kernel of FWHM 6mm; bandpass filtering ($0.1 \text{ Hz} < f < 0.009 \text{ Hz}$); six motion parameters (as estimated by MCFLIRT) regressed out; cerebrospinal fluid and white matter signal regressed out (top five

PCA components, CompCor method).

4.3.4.2 Connectivity matrices

To describe the functional architecture of the whole brain, we transformed the resting-state BOLD time series into connection strength values of the different networks for each participant. The whole brain parcellation was obtained from connectivity-based functional parcellation created by Yeo and colleagues (2011). The 7 network parcellation was used in the current study. We split the networks into two hemispheres and extracted clusters. Two voxels are considered connected only if they are adjacent within the same x, y, or z direction. This yielded 57 clusters from the Yeo 7 networks parcellation. The implementation of spatial clusters extraction was retrieved from python library Nilearn (Abraham et al., 2014, <http://nilearn.github.io/>, version 0.3.1) Next, we extracted and then averaged the time series of all voxels within each cluster to create a cluster specific time series. We used these time series to create region-to-region symmetrical correlation matrices representing the correlations of the network signal that was computed for all the individual subjects. The off-diagonal of each correlation matrix contained 1596 unique region-region connection strengths (i.e., the upper or lower triangle of the network covariance matrix). This approach provided a measure of connection strength of the whole brain for each participant. Finally, Fishers r-to-z transformation was applied to each network covariance matrix.

4.3.5 Conjoint decomposition of connectivity and experience

4.3.5.1 Decomposition methods

We performed a sparse canonical correlation analysis (SCCA; see Hastie et al., 2015) on the functional connectomes and the NYC-Q reports, to yield latent components that reflect multivariate patterns across neural organisation and experience (For similar application, see Wang, Poerio, et al., 2018). SCCA maximised the linear correlation between the low-rank projections of two sets of multivariate data sets with sparse model to regularise the decomposition solutions a process that helps maximise the interpretability of the results. The regularisation function of choice is L1 penalty, which produces 'sparse' coefficients, meaning that the canonical vectors (i.e., translating from full variables to a data matrix's low-

rank components of variation) will contain a number of exactly zero elements. L1 regularisation conducted (i) feature selection (i.e., select only relevant components) and (ii) model estimation (i.e., determine what combination of components best disentangles the neurocognitive relationship) in an identical process. This way we handle adverse behaviours of classical linear models in high-dimensional data. A reliable and robust open-source implementation of the SCCA method was retrieved as R package from CRAN (PMA, penalized multivariate analysis, version 1.0.9 Witten, Tibshirani, & Hastie, 2009). The amount of L1 penalty for the functional connectomes and the NYC-Q reports were chosen by cross-validation. The procedure is described below.

4.3.5.2 Model selection

We employed cross-validation (CV) to select the most useful model across population samples and avoid overfitting (Bzdok & Yeo, 2017). The amount of the two L_1 penalty terms for the functional connectomes and the NYC-Q reports, respectively, were chosen by a nested K-fold CV, where the coefficient for the penalty were chosen using a grid search to maximise the quality of CV objective metric. The objective metric of choice was cumulative explained variances. The explained variance of each latent component was calculated using the squared canonical correlation. High explained variance suggests a high pattern recovery rate between the two data set. The sparse assumption is fundamentally in conflict with the statistical goal of finding components with high explained variance. Therefore we decided the number of components in the model before searching for the best parameter.

We performed confound removal on functional connectomes and the NYC-Q reports as suggested by prior studies (Smith et al., 2015). We removed the effects of nuisance variables from the dataset. These confound variables were sex, age, and head motion indicated by Jenkinsons mean FD (Jenkinson et al., 2002). The removal steps was performed on the training set in each CV fold. We standardized the confound by calculating the z-score, and also squared the three confound measures to account for potentially nonlinear effects of these confounds. The 6 resulting confounds were regressed out of both data matrices. The implementation of the confound removal method (Friston et al., 1994) was




Figure 4.1. A diagram of the nested k-fold cross-validation with model selection.

retrieved from python library Nilearn (Abraham et al., 2014, <http://nilearn.github.io/>, version 0.3.1).

The number of latent components was determined by a preliminary analysis with no sparsity and calculated the explained variances for the two datasets (i.e., brain network correlations and questionnaire ratings). The explained variance increased with the number of components and growth stabilised at 10 components. We selected the number of components based on the point where the tangent stabilised. This led to a model of 4 components, and it accounted for a total of 78% of the variance in connection strength and 29% of the variance in the self-report data. Next, we determined the two coefficients for the L_1 penalty terms that was associated with the best model performance with 4 latent components. We searched for the best L_1 penalty values between 0.1 and 0.9 in 0.1 increments, which resulted in 81 set of parameters. For the nested K-Fold CV, we first separate the data into 5 consecutive folds after shuffling the data and retained one fold as the evaluation set ($N \approx 50$); the other four folds were used as the development set. The development set was further separated into 5 folds for parameter selection and each fold ($N \approx 40$) was used as the validation set once. The model was estimated on the training folds with all parameter sets, and after completion, we trained the model with the winning parameter on the whole development set and the finally tested the performance on the independent, unseen evaluation set. We selected the final parameters according to the best performance on the evaluation set across all folds of the outer CV loop (Figure 4.1). This parameter set is used to train on the full development set and tested on the evaluation set. The parameter grid search and k-fold CV was conducted by the implementation in a Python library scikit-learn (<http://scikit-learn.org/stable/>, version 0.18.2 Pedregosa et al., 2011). The detailed algorithm for selecting the penalty values are presented in Appendix B: Nested K-Fold CV.

The model with the best test performance was selected as the final model. The final models sparsity coefficient are 0.8 (functional connectivity) and 0.5 (self-reports), and the out-of-sample explained variance was 48%. We used the ensuing canonical vectors of the winning SCCA model to compute the latent component scores. There are two sets of canonical scores in a latent component, a weighted sum of variables forms the canonical vectors. For each latent component, we averaged the z-score of the canonical scores of the connection strength and NYC-Q as the combined scores. These scores described the summary of the experience with both the neural basis and the content reports.

4.3.6 Test of component robustness

After identifying the well performed components in compressing the brain-experience data, we examined the robustness of the four components in two different ways. The permutation test is a purely data-driven strategy that access the chance of discovering components in null samples. We also leveraged the brain-experience components to explain the cognitive functions, so that we can identify meaningful patterns by well-established cognitive measurements.

4.3.6.1 Permutation test

We used permutation testing to assess the robustness of the components identified through our analysis. We constructed a null distribution for each canonical component by holding the functional connectivity data in place and randomising the subject-wise order of self-reports data. This permutation scheme broke the link of individual differences in the dataset, therefore testing the robustness of the components in the hypothetical population. By calculating the false-discovery rate in the null distribution, we can conclude the possibility of discovering our components by chance with the given penalty coefficients. Hypotheses that are accepted with a 5% level of significance. In the current analyses we adopt the permutation test with the FWE-corrected p-value by Smith et al. (2015) with data argumentation to increase the size of the resampling datasets to 1000. The four components were compared to the first sparse canonical correlation of the permuted sample. The low-rank components are more relevant than the rest, therefore we yield more conservative p-value by comparing to the first canonical

correlation only. We performed 5000 permutation tests to get enough estimates for 4 decimal places.

4.3.6.2 Group analysis

To determine how patterns of unconstrained neurocognitive activity related to performance on the battery of cognitive tests, we conducted an independent statistical analysis on the identical subjects. A Type III multivariate multiple regression with Pillai's trace test was applied to 4 individual scores for each of the latent components describing experience from the SCCA were the independent variables, and the original 8 measures of cognitive performance were the dependent variables that we hoped to be described by the linear combination of the latent components. Pillai's trace test is considered to be the most powerful and robust statistic for general use (Huberty & Olejnik, 2006). The p-values reported were based on Bonferroni correction. We also performed a principal components analysis (PCA) to identify the patterns of covariance among the 8 measures of cognitive performance and compressed the data. The relation between the principle score and the 4 brain-experience dimensions identified through SCCA was examined in a linear regression model with Pillai's trace test. The analysis was conducted in R (version 3.3.1). The multivariate multiple regression was conducted in R (version 3.3.1) using function `Manova` in R package `car` (companion to applied regression, version 2.1-5).

4.3.7 Code availability

The full analysis pipeline is freely available at <https://github.com/htwangtw/patterns-of-thought>.

4.4 Results

4.4.1 Determining constituent category of experience

We used Sparse Canonical Correlation Analysis (SCCA) to determine connectome-wide dimensions that describe common variance shared by descriptions of brain and experience. This took as input individual scores for the connections between




Figure 4.2. Unique neurocognitive dimensions of population variation revealed by sparse canonical correlation analysis of measures of whole brain connectivity and self-reported descriptions of ongoing experience.

The heat map describes the canonical variate of the network-to-network connectivity between different Yeo networks. The connectivity matrices describes the coefficients from the model, separated into within and between network relationships. The word clouds reflect the coefficients on the relevant self-report items. In both cases the colour bars indicate the magnitudes of the coefficients. A detailed version of the canonical variates and alternative presentation of the self-report coefficients can be found in Online Supplementary Material Figure S1–S5

each of the regions extracted from Yeo’s 7 networks parcellation and the scores of each item of the New York Cognition Questionnaire (NYC-Q).

We applied SCCA with nested 5-fold CV as the model selection strategy. We obtained a model of 4 canonical components with penalty levels of 0.8 on the functional connectivity and 0.5 on the NYC-Q that indicated the best out-of-sample prediction on our data (see Section 4.3.5.2). The canonical correlations of the 4 latent components were 0.28, 0.19, 0.16, and 0.07. The latent components yielded by the best model are presented in Figure 4.2. For the ease of presentation and interpretation, we summarised the components as network-network connectivity instead of 57-by-57 connectivity matrices. The heat maps describe the network-to-network correlations while the word clouds describe the loadings on the self-report items. The components in full and the heat map for the self-report items can be found in Online Supplementary Materials².

Component 1, describes patterns of reduced within network connectivity

²Supplementary data related to this article can be found at <https://doi.org/10.1016/j.neuroimage.2018.04.064>

within all of the networks studied, with this pattern most prominent in the dorsal attention network. Between network connections are generally reduced, with the exception of visual to limbic. Sensorimotor was decoupled from all the other systems, and, in addition, the default and limbic were most decoupled from the attention networks. Experiential themes in Component 1 are dominated by themes related to deliberate planning with a verbal component (high loadings on ‘words’, ‘monologue’, ‘today-plan’, ‘social-future’, ‘purpose’ and ‘deliberate’). We refer to this pattern of reports as reflecting thoughts with ‘purpose’.

Component 2 is dominated by relatively higher within and between network connections. Connectivity within each network was strong with the exception of the limbic network. Between network connections were stronger, with this pattern most apparent in the connections between limbic and ventral attention. In addition, the visual network was strongly correlated with the other networks. This component is dominated by emotional responses (high loadings on ‘anger’, ‘guilt’, ‘cheerfulness’ and ‘happiness’) and social content (‘friends’ and ‘people’). We refer to this pattern of reports as reflecting ‘emotional’ experience.

Component 3 emphasises reduced connections both between and within networks. Within network connectivity is weakest for the dorsal and ventral attention networks. Edge-to-edge connections are low, with the ventral and dorsal attention and frontoparietal networks showing reduced correlations with each other as well as the visual and sensorimotor systems. This component was characterised by themes linked to personal ‘importance’ with social temporal contents (‘distant future’, ‘near future’, ‘social past’, ‘family’ and ‘recent past’). We refer to this pattern of reports as reflecting ‘personal importance’.

Component 4 has the most heterogeneous pattern of within and between network connectivity. It is associated with stronger connections within networks with the exception of the limbic system. In addition, the visual system was strongly connected to all other networks, with this pattern most apparent for the limbic network. In contrast, lower network-to-network connectivity was observed between the default mode and sensorimotor and attention networks. This component is characterised by experiential patterns reflecting a modality difference in experience, with the highest loadings on ‘images’ and lowest on ‘inner monologue’. We refer to this pattern of reports as describing ‘modality’.

4.4.2 The relationship between neuroexperiential components and cognitive functions

Having documented four neurocognitive dimensions, we next examined the robustness of the components using two complementary approaches. We first used a permutation test to identify the chance of discovering components in null samples as employed by Smith and colleagues (2015). The top three components passed the permutation test and the 4th component showed variance that was similar to that produced in a null sample (Component 1: $p = 0.0002$, Component 2: $p = 0.0010$, Component 3: $p = 0.0204$, Component 4: $p = 0.998$, $\alpha = 0.05$). This analysis suggests that Components 1–3 are unlikely to have occurred by chance. Component 4 may be a Type II error and so we discuss this component in only a limited manner moving forward. Our next test of the robustness of our components is whether they explained unique patterns of expertise in our battery of cognitive tasks. We used multiple multivariate regression model in which performance on the battery of selected tasks was the dependent variables and the individual scores for each of the canonical components describing experience from the SCCA were the independent variables. In this analysis two of the four canonical components described significant variance in our battery of tasks at multivariate level: Component 1 ($F(8, 246) = 2.21$, $p = .027$, $\eta_p^2 = .067$) and Component 3 ($F(8, 246) = 2.56$, $p = .024$, $\eta_p^2 = .068$).

In the univariate results of the significant component, Component 1 was linked to good performance in proverb test ($b = 0.48$, 95% CI = [0.1910.766], $t(251) = 3.27$, $p = .006$) and both fluid intelligent tests WASI ($b = 0.39$, 95% CI = [0.1110.677], $t(251) = 2.74$, $p = .033$) and WIAT ($b = 0.45$, 95% CI = [0.1670.724], $t(251) = 3.15$, $p = .009$). Component 3 showed a reversed pattern of the cognitive functions related to Component 1: proverb test ($b = -0.45$, 95% CI = [-0.176 - 0.727], $t(251) = -0.14$, $p = .007$); WASI ($b = -0.42$, 95% CI = [-0.151 - 0.693], $t(251) = -3.10$, $p = .012$) and WIAT ($b = -0.41$, 95% CI = [-0.148 - 0.682], $t(251) = -3.06$, $p = .012$). The relationships between the neurocognitive dimensions and the pattern of relationships on the full cognitive battery and the adjusted variable scatter plots of the significant results are summarized in the form of a heat map in Figure 4.3.

Finally, we performed a simple principle component analysis on the eight




Figure 4.3. The relationship between the different neural-cognitive components and the measures assessed in the cognitive battery.

The components 1 and 3 were significant at the multivariate level determined by multiple multivariate regression, indicated by the asterisk outside of the heat map. The cells with asterisk(s) indicates the significant results from the univariate test (bonferroni corrected) and the parameter estimates for each variable. CWI: Colour-word interference, DF: Design fluency, Pro: Proverbs, TOW: Tower of London, TMT: Trail making task, VF: Verbal Fluency, WASI: Wechsler Adult Intelligence Test, WIAT: Weschler Individual Attainment Test. P-value significant codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ .




Figure 4.4. The principle component and its relationship to the different neurocognitive components.

The heat map describes the principle component of the task battery, and the scatter plots describe the association with the components identified in our study. Component 1 and 3 passed the permutation test for component robustness significantly contributed in explaining the principle component of the task. Component 4 showed a significant contribution in the regression model, but it did not pass the permutation test. The related results should be treated cautiously.

task measures to explore the associations between experience and the structure of the laboratory data. The aim of this analysis was to see if the pattern retrieved from the univariate level in the previous multiple multivariate regression was related to the internal structure of the data. Component selection was determined based on the scree plot, and we accepted one component explaining 39% of the variance. The principle component loaded on the intelligence measures and the proverb test. We fitted a linear model to this data to understand the relationship to the four canonical components. The results are reported in Figure 4.4. The overall linear model was significant ($F(4, 253) = 5.43, p = .00003$). In the linear regression model, Component 1 ($b = 0.82, 95\% \text{ CI} = [0.361.29], t(253) = 3.5, p = .001$) showed significant contribution to explaining the task principle component. Component 3 showed a negative correlation to the task components ($b = -0.69, 95\% \text{ CI} = [-1.13 - 0.24], t(253) = -3.04, p = .003$). The relationships between tasks and the neurocognitive components here were similar to the ones uncovered by the multiple multivariate regression. In this analysis Component 4 ($b = 0.442, 95\% \text{ CI} = [0.160.72], t(253) = 3.09, p = .002$) showed a significant contribution in the regression model, but it did not pass the permutation test of robustness ($p = 0.998$). The related results should be treated cautiously. Together with our prior analysis, these results suggest that Components 1 and 3 are the most robust components identified in our study.

4.5 Discussion

We set out to describe different modes of neurocognitive patterns derived through the simultaneous decomposition of whole brain connectivity data with self-reports of ongoing experience. We used a whole brain parcellation that describes cortical function in seven independent networks (Yeo et al., 2011). We combined this data with self-reports of the experience of our participants at rest, using a multivariate approach that allows for the possibility of many-to-many mappings between neural patterns and ongoing cognition. Our analyses identified four stable canonical components, describing unique dimensions of neural-experiential variation. Permutation testing demonstrated the statistical robustness of Components 1-3. Furthermore, two components (1 and 3) described independent patterns of performance in a battery of commonly used cognitive measures. This association

with cognitive performance that establishes a source of independent validity for these neurocognitive components since they are related to independent measures of cognitive performance. We next consider the fit between the dimensions produced by our analysis and theoretical views of unconstrained neurocognitive processing.

We found evidence broadly consistent with contemporary representational accounts of unconstrained processing. The neural patterns described by Component One reflect a pattern of reduced correlation between regions with links to memory and representation (e.g. limbic, default mode) from those with links to external behaviour (e.g. visual and sensorimotor cortex and attention networks). This pattern was associated with experiences characterised by a sense of purposefulness, and with verbally mediated content that was social and temporal in nature. Participants high on this dimension were proficient at generating abstract semantic links and performed well on measures of reasoning and intelligence. Together the features of Component One support the hypothesis that the functional decoupling of systems important for memory and representation are important for aspects of unconstrained cognition (Smallwood & Andrews-Hanna, 2013). This capacity may arise from the topographical organisation of the cortex, in which neural systems that can take on more transmodal properties tend to be located in regions that are more distant in functional and structural terms (Buckner & Krienen, 2013; Margulies et al., 2016; Mesulam, 1998). This spatial location may allow neural signals in these regions to take on properties that are discrepant from the neural signal more closely tethered to inputs describing the external world (Buckner & Krienen, 2013; Friston, 2013). The pattern identified by Component One, therefore, may reflect a pattern of population variation describing the hypothesised role of functional decoupling of memory and representational systems plays in the generation of more abstract aspects of human cognition (Margulies et al., 2016; Mesulam, 1998). Importantly, in our prior work, limbic and default mode networks were the most distant in functional connectivity terms from unimodal systems (Margulies et al., 2016). Our data also highlights neural patterns that capture the hypothesised influence of attention and control on ongoing thought (McVay et al., 2009). Component 3 highlights links between reduced connectivity within attention and control systems and pat-

terns of thought that emphasise personal importance. This is associated with worse performance on measures of intelligence and reasoning. The combination of a focus on personally important themes linked to poor performance on measures of general aptitude, captures the hallmark psychological features of the ‘current concerns–executive-failure’ accounts of ongoing thought (McVay et al., 2009). This view suggests that failures in attentional control lead to highly personally relevant cognition to intrude into ongoing thought, leading to lapses in task performance. Importantly, the neural pattern described by this component emphasises dysregulated connectivity both within and between networks implicated in attention and control by task-based studies (Duncan, 2010). Our prior work established that spontaneous mind-wandering is linked to cortical thinning within regions linked to attention and control, such as the intra-parietal sulcus (Golchert et al., 2017). Spontaneous mind-wandering has been linked to worse cognitive control (Robison & Unsworth, 2018), as well as showing stronger links with attention related problems, including ADHD (Seli, Smallwood, et al., 2015). Together with these prior studies, our data suggests that population variation in the intrinsic neural functioning within networks with an established role in external task performance captures the hypothesised contribution of executive-failure to patterns of ongoing thought.

The method of decomposition used in the current study also highlighted patterns related to affective processing and the modality of the experience that are similar to those seen in our prior work that applied principal components analysis (PCA) to self-reported data only. Component Four places experiences with visual features (‘images’) in opposition to experiences with verbal features (‘monologue’), capturing dissociations between visual and verbal thinking observed in our prior studies (Konishi et al., 2015; Medea et al., 2016; Smallwood et al., 2016). The accompanying neural pattern were associated with higher connectivity between the visual network with other networks, in particular the limbic system. It is important to note that our permutation analysis failed to validate this component, so despite its association with task performance using the PCA analysis it should be treated with relative caution. Component Two loads on emotional experiences (‘cheerfulness’, ‘anger’, ‘guilt’ and ‘happiness’) with the exception of those that are unhappy (‘sad’). In neural terms this component was

characterised by high levels of connectivity, however, unlike Component Four, this was highest between limbic and ventral attention networks. This pattern of coupling is consistent with accounts that emphasise interactions between saliency and limbic systems in affective processing (Touroutoglou et al., 2012). In the case of Component Two permutation testing indicated this component was likely to be robust in statistical terms, however, we did not observe associations with task performance. As with Component Four, interpretations of Component Two should be made with caution in lieu of more empirical work.

Before closing it is worth considering several important limiting factors of our study. We focused on patterns of population variance in unconstrained neuro-cognitive processing that were measured once in each individual. Our study, therefore, cannot separate the influences of state and traits on our observed components. Treating patterns of unconstrained processing as a trait is common in both the psychological (McVay et al., 2009; Smallwood, Ruby, & Singer, 2013) and neural domains (Smith et al., 2015). Nonetheless, it remains an open question how consistent these components will be across individuals over time, as well as which aspects may be better described as traits. Importantly, by its very nature there are dimensions of experience that our study cannot adequately address. We cannot, for example, identify brain-experience associations that are highly dynamic in nature and in particular those that change rapidly within an individual. Insight into this issue could be achieved by a focus on dynamic rather than static connectivity (Kucyi, Hove, Esterman, Hutchison, & Valera, 2017). For example, the application of techniques such as sliding window analysis (Chang & Glover, 2010) or Hidden Markov models (Vidaurre, Smith, & Woolrich, 2017) to fMRI could provide information that would complement our analyses. However, it may also be more important to examine these across multiple sessions within the same individuals, as this would also make it most possible to dissociate state from trait related influences on neural activity (Mueller et al., 2013). There are also types of experience that may be difficult to assess using the measure of retrospective experience sampling we have employed (Smallwood & Schooler, 2015). For important features of experience, such as whether it has evolving features (Mills, Raffaelli, Irving, Stan, & Christoff, 2018), or when the participant is unaware of the content of their experience (Schooler, 2002) these

experiential features may be best assessed using experience sampling techniques that capture momentary elements of experience (Smallwood, 2013).

There are a number of methodological improvements that could enhance future studies of brain-experience association. A recent benchmark study by Ciric et al. (2017) shows that scrubbing can improve the performance of resting state analyses. Regarding to the analysis pipeline, we gained hyper-parameters and best model with nested-CV an approach that can help prevent overfitting (Bzdok & Yeo, 2017). There are also alternative ways that could provide better tests of the robustness of the components we identified. We assessed the validity of the components in three different ways; (i) with a data-driven, non-parametric permutation test (Smith et al., 2015) that establishes the statistical validity of the identified components and (ii) by establishing the relationship between the laboratory cognitive measures and (iii) by consideration of their links with contemporary theoretical accounts of ongoing cognition. In our study, Components 1 and 3 were statistically significant in both cases and fitted well with contemporary accounts of ongoing cognition. Accordingly we place encourage readers to focus on these patterns from our data. There are alternative strategies that could help validate the robustness of patterns of brain-experience association. One approach could be to compare the relationship between multiple sessions within the same individual (Poldrack et al., 2015) and to have a larger sample that would allow the reproducibility of these results through a formal split-half validation procedure. To achieve this latter aim for future studies, we have placed the questionnaire measure used in this study along with an example self-report collection task on GitHub at the following address: https://github.com/htwangtw/restingstate_thoughtreports. We encourage interested investigators to apply these measures in their resting-state investigation and to also upload the resultant data onto open fMRI. These studies could be used in conjunction with the openly access data used in this study to enable future investigations the opportunity to cross validate experiential analyses in a more sophisticated manner than we have been able to achieve in this study. The analysis pipeline of the current study can be further unified into one frame work that benefits from both validation strategies. We can include the number of components along with penalty coefficients in the hyper-parameters

determined in the CV process, or determine the best penalty terms with the first component. The permutation test will then identify the reliable components occurring above chance level. After all the data-driven component selection, we can examine the survived components through their relations with well-documented cognitive measures and conclude the meaningful patterns. Finally, it is likely that our measure of ongoing thought lacks important questions regarding the content of experience. It will be important, therefore, in the future to examine the relationships of the type described in this study with a more exhaustive description of ongoing experience. We hope that by publishing our questionnaire collection task in a GitHub repository we will be able to harness the power of the broader community to help generate and test plausible questions for use in future studies.

Chapter 5

Inhibition of Prior Information Contributes to Internal Content Representation

5.1 Abstract

Although the patterns of ongoing thought that make up our day to day lives are important, we know relatively little about how these experiences are constrained by an individuals' neurocognitive architecture. In the current study we used machine learning to identify stable patterns describing shared variance between performance on a battery of cognitive tasks in the laboratory, and intrinsic neural architecture observed at rest. Next we explored whether these dimensions explained variance in measures of ongoing thought recorded in the laboratory. We identified five neurocognitive dimensions characterised at the cognitive level as creative association, fluid intelligence, temporally specified cognition, and, separate dimensions of episodic memory linked to visual and verbal codes of representation. Variation in *temporally specified cognition*—the ability to inhibit a prior mental set during task switching—predicted substantial variance in ongoing experience recorded in the lab, accounting for reduced variance in two principle dimensions identified by principal component analysis—less off-task thought and less immersive detail. In neural terms temporally specified cognition was characterised by patterns of high within-network limbic connectivity coupled with

relative isolation between this system and other regions of cortex. Together this analysis suggests that whether an individual's thoughts are a pristine representation of the moment, or an immersive experience generated via imagination, depends cognitively on the ability to inhibit prior mental sets, and neurally on the balance of segregation and integration between the limbic system and the rest of the cortex.

5.2 Introduction

Human cognition is not always tethered to the events in the here now. Phenomena such as mind-wandering highlight that we can become immersed in experiences generated from memories, rather than information in the immediate environment (Smallwood & Schooler, 2015). Although understanding self-generated experiences may ultimately inform theoretical accounts of both normal cognition and disease states, we currently lack a comprehensive understanding of how these experiences are constrained by an individual's neurocognitive architecture.

Contemporary studies have shown that ongoing thought shares important links with both brain and behaviour. For example, the occurrence of off-task thought can jeopardise the integrity of tasks depending on executive control (McVay & Kane, 2009; Mrazek et al., 2012). On the other hand, states of off-task thought and daydreaming can be associated with better performance on tasks of memory and creativity (Ruby, Smallwood, Engen, & Singer, 2013; Poerio et al., 2017; Wang, Poerio, et al., 2018; Baird et al., 2012). Neuroimaging studies have highlighted important roles for a number of large-scale brain networks (see meta-analysis from K. C. R. Fox et al., 2015; Stawarczyk & D'Argembeau, 2015). These networks include the default mode network (Mason et al., 2007; Christoff et al., 2009), the frontoparietal and attention networks (Wang, Bzdok, et al., 2018; Golchert et al., 2017) and the limbic system (Ellamil et al., 2016; Smallwood et al., 2016; Golchert et al., 2017).

Associations between patterns of ongoing thought with objective metrics defined from brain and behaviour, raise the possibility that these metrics of individual difference could be used to gain traction on the architecture that underlies patterns of ongoing thought. In the current study, a large cohort of participants ($n = 197$) performed a battery of neurocognitive tasks in the laboratory, and, on

a separate session, we measured their intrinsic brain activity using resting-state functional magnetic resonance imaging (fMRI). These individuals also provided descriptions of their experience while they performed a simple laboratory task across several days.

Using these data we build on our prior work that used sparse canonical correlation analysis to reveal neurocognitive dimensions that relate to patterns of ongoing experience (Wang, Poerio, et al., 2018; Wang, Bzdok, et al., 2018). In this study, we used sparse canonical correlation analysis to identify dimensions that linked brain organisation to behaviour and used these as explanatory variables in analyses predicting patterns of ongoing thought in the laboratory. This allows us to test the view that ongoing thought is an emergent property of an individuals' neural architecture (i.e. Gratton et al., 2018).

5.3 Method

5.3.1 Participants

Two hundred and seven healthy participants were recruited from the University of York (132 females, 65 males; age range = 18-31 years, $M = 20.21$, $SD = 2.36$). This analysis included two data sets with some shared measurements and the same MRI protocol as Chapter 3. Participants were right-handed native English speakers with normal or corrected-to-normal vision and no history of psychiatric or neurological illness. Participants underwent MRI scanning, completed a 1-hr online questionnaire. The first cohort is identical to the sample in Chapter 3. Participant attended three (165 participants; 99 females, 66 males; age range = 18-31 years, $M = 20.43$, $SD = 2.63$) 2-hr behavioural testing sessions to complete a battery of cognitive tasks. The second cohort (42 participants; 33 females, 9 males; age range = 18-23 years, $M = 19.79$, $SD = 1.37$) underwent two 2-hr behavioural testing sessions to complete a battery of cognitive tasks. The behavioural sessions took place within a week of the scan. Ten participants were excluded from the multivariate pattern analysis because they failed to complete all of the behavioural testing sessions. In total, 197 participants (126 females, 71 males; age range = 18-31 years, $M = 20.11$, $SD = 2.24$) were included in the multivariate pattern analysis and the comparison with cognitive performance.

Participants were rewarded with either a payment of £10 per hour or a commensurate amount of course credit. All participants provided written consent prior to the fMRI session and the first behavioural testing session. Approval for the study was obtained from the ethics committee of the University of York Department of Psychology and the University of York Neuroimaging Centre.

5.3.2 MRI acquisition

The MRI acquisition protocol was identical to the study documented in Chapter 3. Please refer to Section 3.3.2 for details.

5.3.3 Resting state data preprocessing

All preprocessing and denoising steps for the MRI data were carried out using the SPM software package (Version 12.0) and Conn functional connectivity toolbox (Version 17.f), based on the MATLAB platform (Version 17.a). The first three functional volumes were removed in order to achieve steady-state magnetisation. The remaining data were first corrected for motion using six degrees of freedom (x, y, z translations and rotations), and adjusted for differences in slice-time. Subsequently, the high-resolution structural images were co-registered to the mean functional image via rigid-body transformation, segmented into grey/white matter and cerebrospinal fluid probability maps and all functional volumes were spatially normalized to Montreal Neurological Institute (MNI) space using the segmented images and a priori templates. This indirect procedure utilises the unified segmentationnormalization framework, which combines tissue segmentation, bias correction, and spatial normalization in a single unified model. No smoothing was employed, complying with recent studies that report the negative influence of this procedure on the construction of connectivity matrices analysis.

Moreover, a growing body of literature indicates the potential influence of participant motion inside the scanner on the subsequent estimates of functional connectivity. To ensure that motion and other artefacts did not confound our data, we have employed an extensive motion-correction procedure and denoising steps, comparable to those reported in the literature. In addition to the removal of six realignment parameters and their second-order derivatives using the general linear model (GLM), a linear detrending term was applied as well

as the CompCor method that removed five principal components of the signal from white matter and cerebrospinal fluid. Moreover, the volumes affected by motion were identified and scrubbed based on the conservative settings of motion greater than 0.5 mm and global signal changes larger than $z = 3$. Though recent reports suggest the ability of global signal regression to account for head motion, it is also known to introduce spurious anti-correlations and was thus not utilised in our analysis. Finally, a band-pass filter between 0.009 Hz and 0.08 Hz was employed in order to focus on low-frequency fluctuations.

5.3.4 ROI-ROI functional connectivity.

We adopted a set of 57 regions based on the Yeo 7 networks. We split the networks into two hemispheres and extracted clusters. Two voxels are considered connected only if they are adjacent within the same x, y, or z-direction. This yielded 57 clusters from the Yeo 7 networks parcellation. The implementation of spatial clusters extraction was retrieved from python library Nilearn (Abraham et al., 2014, <http://nilearn.github.io/>, version 0.3.1). Fully-connected, undirected and weighted matrices of bivariate correlation coefficients (Pearson's r) were constructed for each participant using the average BOLD signal time series obtained from all the 57 ROIs described above. The off-diagonal of each correlation matrix contained 1596 unique region-region connection strengths (i.e., the upper or lower triangle of the network covariance matrix). This approach provided a measure of the connection strength of the whole brain for each participant. Finally, Fisher's r -to- z transformation was applied to each network covariance matrix.

5.3.5 Behavioural data

5.3.5.1 Cognitive tasks

We selected 9 cognitive tasks that are common across the two cohorts. The selected tasks measures cognitive functions that have been examined in previous mind-wandering literature, using the same or similar tasks adopted by previous mind-wandering research, including executive control (digit span: (from Wechsler, 1999), task switching task: Handy and Kam (2015)), generation of information (unusual uses task: Mrazek et al. (2012), verbal fluency task: Adlam,

Patterson, Bozeat, and Hodges (2010) and Balota and Coane (2008)), semantic memory (semantics relatedness judgement tasks and feature matching task, both developed by Krieger-redwood (2012)), episodic memory (paired-associate task: Cairney, Lindsay, Paller, and Gaskell (2016), four mountains task: Hartley et al. (2007)), and fluid intelligence (Raven Advanced Progressive Matrices; RAPM: Baird et al. (2012)). Please refer to Appendix A.2 for the detailed descriptions and references of the tasks.

Thirteen cognitive scores were calculated from the selected tasks. Performance of the digit span task was represented as the average of digit span in the forward and backward recall conditions. The verbal fluency score is the contrast of the letter condition and the category condition (letter–category). Picture naming tasks, the four mountains tasks, RAMP were summarised with accuracy scores. The task switching tasks provided two scores (a) flexibility¹ as the ability to switch from a different condition and (b) inhibition as the ability to suppress information from the previous trial. The calculation of the task switching contrast can be found in Appendix A.2.7. All the semantics related judgement tasks, feature mating task and the paired associate task were summarised using efficiency scores. The efficiency scores were calculated as reaction time divided by accuracy. A smaller score indicates better performance, thus the scores were reversed to ease the interpretation. For the semantics tasks, we calculated three reaction contrasts based on the semantics modules tested: (a) strength (strong–weak), modality (picture–word), and (c) specificity (specific–general). All the scores were standardised as z-scores in the subsequent analysis. We defined outliers as scores greater than 3. The identified outliers were then imputed with medians of each variable.

5.3.5.2 Experience sampling

Multi-dimensional experience sampling (MDES) was used to collect thoughts during a 0-back/1-back working memory task. Please refer to Chapter 3 Section 3.3.4 for the MDES data collection and Table 3.1 for the detailed questions.

¹ In the original contrast ‘switch cost’, smaller values indicates a better ability to switch away from the previous condition. For the ease of interpretation, we reversed the scores and re-named the contrast as flexibility.

The MDES questions were aggregated (a) across all conditions, (b) within the 0-back condition, and (c) within the 1-back condition. All the scores were standardised as z-scores in the subsequent analysis. We defined outliers as scores greater than 3. The identified outliers were then imputed with medians of each variable.

5.3.5.3 Dimensions of ongoing thought

For the purpose of analyses, the scores on the 13 MDES questions were entered into a PCA to describe the underlying structure of the participants' responses. Following prior studies (Konishi et al., 2017; Medea et al., 2016; Ruby, Smallwood, Engen, & Singer, 2013) we concatenated the responses of each participant in each task into a single matrix and employed a principal components reduction with varimax rotation (see the top panel of Figure 5.1). We selected the number of components based on the elbow in the scree plot.

5.3.6 Multivariate pattern analysis

5.3.6.1 SCCA

We performed a sparse canonical correlation analysis (SCCA; see Hastie et al., 2015) on the functional connectomes and the cognitive tasks, to yield latent components that reflect multivariate patterns across neural organisation and cognition (For similar application, see Wang, Poerio, et al., 2018). SCCA maximised the linear correlation between the low-rank projections of two sets of multivariate data sets with a sparse model to regularise the decomposition solutions a process that helps maximise the interpretability of the results. The regularisation function of choice is the L1 penalty, which produces 'sparse' coefficients, meaning that the canonical vectors (i.e., translating from full variables to a data matrix's low-rank components of variation) will contain a number of exactly zero elements. L1 regularisation conducted (a) feature selection (i.e., select only relevant components) and (b) model estimation (i.e., determine what combination of components best disentangles the neurocognitive relationship) in an identical process. This way we handle adverse behaviours of classical linear models in high-dimensional data. A reliable and robust open-source implementation of the SCCA method was retrieved as R package from CRAN (PMA, penalized mul-

tivariate analysis, version 1.0.9 Witten et al., 2009). The amount of L1 penalty for the functional connectomes and cognitive task performance were chosen by cross-validation. The procedure is described below.

5.3.6.2 Model Selection

The model selection process was conducted with two parts: the L1 penalty coefficient selection and component selection. For the L1 penalty coefficient selection, we performed a grid search combined with cross-validation (CV) to avoid over-fitting (Bzdok & Yeo, 2017). Of each penalty pair on the search grid, a 10-fold CV was performed to search for the best out-of-sample the rank-1 canonical correlation (see the middle panel of Figure 5.1). We then decomposed the full dataset with the selected L1 penalty coefficients. The K-Fold CV was conducted by the implementation in Python library scikit-learn (<http://scikit-learn.org/stable/>, version 0.18.2 Pedregosa et al., 2011).

Confound removal was performed on the functional connectomes and the cognitive scores as suggested by prior study (Smith et al., 2015). The confound variables were sex, age, and head motion indicated by mean frame-wise displacement (Jenkinson et al., 2002). The removal steps were performed on the training set in each CV fold. The z-scores of the confound variables were calculated, and also squared the three confound measures to account for potentially nonlinear effects of these confounds. The 6 resulting confounds were regressed out of both data matrices. The implementation of the confound removal method (Friston et al., 1994) was retrieved from Python library Nilearn (Abraham et al., 2014, <http://nilearn.github.io/>, version 0.3.1).

After finding the optimal hyper-parameters, 1000 permutation tests with family-wise error (FWE) correction was applied to access the component(s) that occur above chance (see the bottom panel of Figure 5.1). We constructed a null distribution for each canonical component by holding the functional connectivity data in place and randomising the subject-wise order of self-report data. This permutation scheme broke the link of individual differences in the dataset, therefore testing the robustness of the components in the hypothetical population. By calculating the false-discovery rate in the null distribution, we can conclude the possibility of discovering our components by chance with the given

penalty coefficients. Hypotheses that are accepted with a 5% level of significance. In the current analyses, we adopt the permutation test with the FWE-corrected p-value by Smith and colleagues (2015). All components were compared to the first sparse canonical correlation of the permuted sample. The low-rank components are more relevant than the rest, therefore we yield more conservative p-value by comparing to the first canonical correlation only.




Figure 5.1. Analysis pipeline.

Top: PCA on MDES (Section 5.3.5.2). Middle: sparsity parameter selection (Section 5.3.6.2). Bottom: permutation test procedure (Section 5.3.6.2).

5.3.7 Group analysis

To determine how patterns of unconstrained neurocognitive activity related to performance on the self-report experience summarised in three different ways (see section 5.3.5.2), we conducted three independent statistical analysis on the

identical subjects. A Type III multivariate multiple regression with Pillai’s trace test was applied to the data. Each of the latent components describing the neurocognitive mechanism from the SCCA was the independent variables, and the 13 measures from MDES were the dependent variables. We hoped to described the neurocognitive components by the linear combination of the self-report questions collected via MDES. The p-values reported were based on Bonferroni correction. The analysis was conducted in R (version 3.3.1). The multivariate multiple regression was conducted in R (version 3.3.1) using function `Manova` in R package `car` (companion to applied regression, version 2.15).

5.4 Results

5.4.1 Dimensions of ongoing thought

Experience sampling probes revealed four unique dimensions of ongoing thought during the 0-back/1-back task. PCA of the 13 experience sampling questions resulted in four principle components of thought presented in Figure 5.2. Consistent with our prior study (Poerio et al., 2017), the components were characterised as: detailed thought, off-task thought, modality of thought, and emotional thought.

5.4.2 Neurocognitive component selection

Sparse Canonical Correlation Analysis (SCCA) was used to determine connectome-wide dimensions that describe common variance shared by descriptions of brain and experience.




Figure 5.2. Dimensions of ongoing thought.

The result from the PCA is presented as a heatmap. The colour bar represents the value of the principal component loading.

This took as input individual scores for the connections between each of the regions extracted from Yeo’s 7 networks parcellation and the 13 cognitive task scores.

We applied SCCA with 10-fold CV and permutation tests as the model selection strategy. We obtained penalty levels of 0.6 on both the functional connectivity and cognitive tasks indicated the best out-of-sample prediction on our data through the grid search (Figure 5.3), obtaining 0.70 on the rank-1 canonical correlation. Five significant canonical components were identified through FWE-corrected p-value through permutation tests. The p-values of the 5 components are 0.028, 0.042, 0.041, 0.012, 0.033. The canonical correlations of the 5 significant latent components were 0.68, 0.68, 0.68, 0.70, and 0.68. Since the sparsity turns CCA into a convex optimisation problem, the modes didn’t come out in the descending order of their canonical correlations.




Figure 5.3. Grid search result.

This heatmap represents the rank-1 canonical correlations of each sparsity coefficient pairs determined by the CV. The red square indicates the best result.

5.4.3 Determining constituent category of cognitive functions

The latent components yielded by the best model are presented in Figure 5.4. Using SCCA we identified five neurocognitive dimensions characterised at the cognitive level as the creative association (Component 1), fluid intelligence (Component 9), temporally specified cognition (Component 7), and, separate dimensions of episodic memory linked to visual (Component 8) and verbal (Component 12) codes of representation. For the ease of presentation and interpretation, we summarised the components as network-network connectivity instead of 57-by-57 connectivity matrices. The heat maps describe the network-to-network correlations and the cognitive tasks.




Figure 5.4. Significant components from SCCA.

The colour in the heatmap indicates the value of the canonical coefficients of each component. VIS: visual network, S-M: somatomotor network, VAN: ventral attention network, DAN: dorsal attention network, LIM: limbic network, FPN: frontoparietal network, DMN: default mode network.

Component 1 emphasis semantic control with the picture-naming task and the unusual uses task, along with intelligence in the cognition component. The negative coefficient in semantic strength indicated the ability to detect weaker semantics relationships. The functional connectivity pattern shows a general dissociation among all networks, except between the unimodal systems. Network segregation was especially pronounced between the sensorimotor network and the limbic system, and a general dissociation between the unimodal sensory system and the attention and transmodal regions. Component 1 demonstrated semantic control ability to generate mental representation and semantic associations.

Component 7 reflects better performance on task switching tasks, both in terms of a reduced switch cost, and the ability to suppress prior mental sets. The performance was better for both visual semantic decisions and specific semantic imagery. In addition, participants performed better on the four mountains task and had a larger digit span. This pattern of performance was linked to the better

performance of tasks that requires mental representations to be controlled across time (task switching, four mountains task and digit span) particularly with a particular emphasis on visual processing. The connectivity pattern shows strong connectivity between sensory systems. In addition, the visual system showed reduced connectivity between the attention systems, while the limbic system was less correlated with all systems other than the FPN. Within network connectivity was strong for the limbic, frontoparietal and default mode networks. Overall, this component demonstrates the ability to control representations over time and is linked to integrity within limbic and transmodal systems and separation between visual and attention systems.

Component 8 pits executive tasks against verbal episodic memory systems since it is linked to better paired-associate memory, verbal semantic memory and feature matching and worse task switching, digit span and verbal fluency. The functional connectivity pattern shows a general pattern of reduced network connectivity. Exceptions to this include the unimodal networks the visual network shows stronger connections to attention and default mode networks, while the sensorimotor network shows stronger coupling to the ventral attention and limbic networks. Within network connectivity is higher in the ventral attention, limbic and default mode networks.

Component 9 is composed of tasks that rely on controlled processing (fluid intelligence, task switching, fluency and controlled semantic retrieval). It is also linked to worse picture naming. In connectivity terms, the default mode network shows stronger connectivity with unimodal and attention networks, and the dorsal attention network is linked to stronger connectivity with the visual network. Within network connectivity is low within the visual, dorsal attention and default mode networks and high in the ventral attention network.

Component 12 highlights the between episodic memory (better paired-associate memory) and worse visual semantic memory (worse verbal semantics and poor figure matching). Fluency was better when organised alphabetically rather than by categories. The default mode network and the visual system showed a general coupling pattern with the other networks, while a strong dissociation of the limbic with the attention systems and the sensorimotor system. The component demonstrates a strong ability to retain and recall information that does not

benefit from the semantic organisation.

5.4.4 The relationship between neurocognitive components and self-reports on thoughts

Three regression models were performed to evaluate the relationship between neurocognitive components and self-reports on thoughts: average scores of all thought probes thought probes in 0-back and 1-back conditions (Figure 5.5).

We first examined the relations of the neurocognitive components and the overall thought reports. In the model of all thought probes, we got multivariate main effect in component 7 ($F(13, 160) = 2.239, p = .010, \eta_p^2 = .154$) and 12 ($F(13, 160) = 1.946, p = .029, \eta_p^2 = .137$). There was only one significant univariate effect after Bonferroni correction, which is the negative correlation on self-question under component 7 ($b = 0.25, 95\% \text{ CI} = [0.40, 0.10], t(172) = -3.369, p = .005$). The results revealed two types of thought patterns. Component 7 focuses on information maintenance in cognitive tasks and the integration within the separated transmodal systems. The related thought patterns shows a low tendency in reporting personal issues. Although there were no significant contributing univariate pattern in component 12, we see a trend of deliberate, focused thought pattern with low imagery information related to retrieval of semantically novel associations.

Two models examined the average scores of thought probes during the 0-back and 1-back condition separately. The aim is to uncover potential differences in thought reports under the two conditions. In the 0-back condition, only component 7 ($F(13, 160) = 1.924, p = .031, \eta_p^2 = .135$) showed the main effect in multivariate level. There was only one significant univariate effect after Bonferroni correction, which is the self-question under the model for component 7 ($b = 0.23, 95\% \text{ CI} = [0.37, 0.79], t(172) = -3.03, p = .017$). In the 0-back condition, the participants perform a visual matching task. The focused state during the 0-back condition is associated with more cognitive mechanism that sustains ongoing thought.

The 1-back condition model showed significant results in component 7 ($F(13, 160) = 2.192, p = .012, \eta_p^2 = .151$) and component 12 ($F(13, 160) = 2.312, p = .008, \eta_p^2 = .158$). There was only significant univariate effect in the model

for component 7 after Bonferroni correction. The significant variable is the self question ($b = 0.26$, 95% CI = [0.40, 0.11], $t(172) = -3.51$, $p = .003$) and the habit question ($b = 0.22$, 95% CI = [0.37, 0.75], $t(172) = -2.97$, $p = .020$). The 1-back condition required participants to maintain meaningless associations of the two shapes presented on the screen.




Figure 5.5. Group level analysis on neurocognitive components and self-report on thoughts.

The codes next to the component number indicate the significant level of the multivariate results, and those in the coloured cells are for the univariate results. The colour represent the univariate b value. P-value significant codes: ‘***’: < 0.001 ; ‘**’: < 0.05 ; ‘.’: < 0.01

5.4.5 The relationship between neurocognitive components and ongoing thought patterns

Pearson’s correlations were calculated to explore the relationships between the neurocognitive components and the dimensions of ongoing thought. The temporally specific cognition component (Component 7) is negatively correlated to details and off-task components (Figure 5.6).




Figure 5.6. The relationships between the neurocognitive components and the ongoing thought.

This heatmap represents correlation between the principal components and the univariate level predictions in Section 5.4.4.

5.5 Discussion

We set out to identify patterns that described the association between different aspects of cognition and the intrinsic organisation of the cortex and to explore whether these accounted for variations in patterns of ongoing thought. Using SCCA we identified five neurocognitive dimensions characterised at the cognitive level as the creative association (Component 1), fluid intelligence (Component 9), temporally specified cognition (Component 7), and, separate dimensions of episodic memory linked to visual (Component 8) and verbal (Component 12) codes of representation. In our subsequent analysis, we identified that variation in temporally specified cognition was associated with substantial variance in patterns of ongoing thought recorded in the laboratory. In particular, we found that this neurocognitive dimension was associated with variation in both the task relatedness of cognition and its level of immersive details. In the discussion, we consider the implications of these data for theoretical accounts of ongoing thought.

In neural terms, our CCA analysis suggests that the relative degree of integration/segregation of the limbic system is at the core of whether an individual's experience is a pristine reflection of their current external goal, or instead, they are immersed in thoughts generated through imagination. We found that indi-

viduals who maintained attention on the task in hand, tended to show a pattern of brain activity dominated by whether the limbic system was coupled to itself, but decoupled from other cortical areas, while individuals who reported off-task experiences with immersive qualities showed the reverse pattern (low within network coupling, and high between network coupling for the limbic system). These data add to a growing body of evidence that highlights limbic regions as critical for patterns of spontaneous thought. For example, lesions to the hippocampus are associated with reductions in off-task thinking (McCormick, Rosenthal, Miller, & Maguire, 2018) and episodic future thinking as part of a task (Maguire & Hassabis, 2011; Race, Keane, & Verfaellie, 2011). Likewise, semantic dementia, which targets the lateral and medial aspects of the temporal pole, reduced the capacity to imagine the future (Irish, Addis, Hodges, & Piguet, 2012; Viard et al., 2014). Furthermore, hippocampus activity has been shown to be important for spontaneous thought during the occurrence of spontaneous thought (Ellamil et al., 2016), while its connectivity with regions of the default mode network is important for both episodic features of spontaneous experience (Karapanagiotidis et al., 2017) as well as its immersive features (Smallwood et al., 2016). Based on our results, the contribution of limbic structures to spontaneous experience may depend on their coupling with other regions, allowing these hub regions to integrate information from across the cortex to create a mental scene (Hassabis & Maguire, 2009). This account is broadly consistent with views of limbic structures, such as the hippocampus (Moscovitch et al., 2016) and the anterior temporal lobe (Lambon-Ralph et al., 2017) which are both thought to share a hub and spoke architecture in which their contribution to cognition arise from their capacity to integrate information from across the cortical mantle.

In our analysis, participants with whom the limbic system was relatively isolated within the cortical mantle, performed well in a task switching context that required them to suppress representations of a previous task set. Prior uses of this task paradigm have documented that this ability is linked to the tendency to ruminate. Whitmer and Banich (2007) found that individuals who were high on trait rumination were better than non-ruminators when switching back to a prior task. This pattern is broadly consistent with our data which shows that people who show the smallest cost from inhibiting a prior mental set, were

more likely to reports patterns of ongoing thought that were characterised by immersive experiences characterised by self-relevant, social and episodic content. Building on our study a promising area for future study would be to examine how limbic connectivity supports the recurrence of negative self-relevant experiences that are thought to be important in rumination (Kleckner et al., 2017; Peters, Burkhouse, Feldhaus, Langenecker, & Jacobs, 2016; Cooney, Joormann, Eugène, Dennis, & Gotlib, 2010).

Our data suggest that the default mode network shows a similar, albeit less pronounced pattern, to the limbic system. Given evidence of a role for the DMN in both immersive experiences in task states (Sormaz et al., 2018; Richter, Cooper, Bays, & Simons, 2016) as well as in off-task states (Mason et al., 2007; Christoff et al., 2009; Stawarczyk et al., 2011). It is possible that these two networks work in tandem when cognition is focused on self-relevant information with the limbic systems providing the episodic and conceptual content, and the default mode network allowing this content to be represented at a relatively abstract level. This interpretation is consistent with the observation that the default mode network is spatially located at the top of a hierarchy and most distant from unimodal inputs, while limbic regions occupy an intermediate position (Margulies et al., 2016).

Finally, it is worth considering a number of limitations of this study. First, we did not measure patterns of ongoing thought while individuals performed the battery of cognitive tasks. It is, therefore, possible that part of the shared variance that our analyses capture emerges because of the patterns of ongoing thought occur during the cognitive tasks (see Mrazek et al., 2012, for evidence of a similar point in the context of executive control or intelligence tasks). However, this interpretation of our data is unlikely since individuals who were off-task tended to perform better on the tasks when they returned to a mental set that had recently been active. Second, our neural data was only measured on a single occasion, raising the possibility that this measure of brain function reflects a state rather than a trait. While this remains a possibility, recent studies have shown that individual patterns of functional connectivity remain relatively consistent across tasks and time (Gratton et al., 2018). Nonetheless, future studies could benefit from measuring an individuals architecture across multiple

points to provide a more robust indication of its trait like features.

Chapter 6

General discussion

Contemporary research on patterns of self-generated thought, such as those occurring during states of mind-wandering, is riddled with contradictions. The content of ongoing thought varies from future-orientated planning thoughts that may help refine personal goals (Medea et al., 2016) to negative past concerns that can maintain unpleasant affective states (Killingsworth & Gilbert, 2010). Likewise, research has highlighted the disadvantage of off-task thought during tasks that demand continuous external attention (McVay & Kane, 2009, 2012a), whereas research on creativity and problem solving suggests evidence of beneficial influence from off-task thought (Smeekens & Kane, 2016; Baird et al., 2012).

The conflict presented above is thought to emerge because of variability in the nature of ongoing thought. Heterogeneous ongoing thought may be composed of a set of experiences with overlapping features—the so-called family resemblance account of mind-wandering (Smallwood, 2013; Seli et al., 2018). One aim of this thesis was to develop an empirical approach sensitive to possible similarities among the heterogeneous patterns of thought. In particular, the goal was to identify multiple patterns of ongoing experience that share common and distinct features that can be empirically measured. To implement this goal, this thesis examined the intersection between subjective reports and objective measures—in this case patterns of resting-state functional connectivity and performance on cognitive tests. The current thesis employed SCCA (Witten & Tibshirani, 2009)—a multivariate approach that measures similarities of linear patterns in two domains of data. SCCA identifies various patterns of thought while reflecting the covariance between both brain and cognition. The overarching aim of this

thesis is to provide evidence in support of the family resemblance view of ongoing thought, and facilitate a more constrained theoretical account of how different patterns of ongoing thought emerge.

6.1 Empirical findings

The current thesis comprised three studies focusing on resolving the heterogeneous features of ongoing thought by considering its intersection with measures of neural function—in this case intrinsic functional connectivity at rest. Chapter 3 revealed that mind-wandering is a collection of different ongoing thoughts that are derived from the connectivity patterns in DMN. Chapter 4 found that the population variance in intelligence is related to different whole-brain neural hierarchies and ongoing thoughts. Chapter 5 showed that either momentary or longer mental representation is associated with the ability to inhibit prior mental sets and the balance of segregation and integration between the limbic system and the rest of the cortex. Describes in detail two themes that emerged from this work—the heterogeneity of patterns of ongoing thought and the integration and segregation of transmodal networks.

6.1.1 Heterogeneity

In mind-wandering literature, converging evidence highlights heterogeneity in the variety of functional outcome linked to off-task thought (Smallwood & Andrews-Hanna, 2013), the definitions that researchers have used to study different types of ongoing thought (Seli, Risko, Smilek, & Schacter, 2016), and the number of competing theoretical accounts (e.g. Smallwood, 2010; McVay & Kane, 2010). The current thesis aimed to systematically explore whether this conflicted literature is a consequence of the emergence of distinct patterns of ongoing thought with different experiential features and associated outcomes. This question was explored by focusing on a single candidate neural system—the DMN (Chapter 3) and at the whole-brain level (Chapter 4).

It is a widely held view that the DMN is important in certain types of ongoing thought, such as mind-wandering (see a review from Smallwood & Schooler, 2015). The aim of Chapter 3 was to explore whether associations between pat-

terns of DMN connectivity and measures of experience yielded unique patterns with distinctive patterns of functional outcomes. In this study we found two reliable neuroexperiential patterns, each with distinct functional outcomes (see Figure 3.4). Internal connectivity in the mPFC was related to positive habitual experience and was predicted by poor executive control (see Figure 3.2), suggesting that this may correspond to patterns of executive failure linked to the mind-wandering state (McVay et al., 2009). The relationship between ongoing thought and deficits in executive control could be a result of failure to allocate cognitive resource to an external task. In addition, patterns of PCC-TPJ-mPFC decoupling (see Figure 3.2), associated with off-task experience, provided a link to better performance on tasks requiring the generation of information. The continuous content generation associated with an external task is similar to the association between patterns of off-task thought and creativity (Baird et al., 2012). Together, Chapter 3 provided evidence that ongoing thoughts unfold along a set of heterogeneous dimensions. Critically, Chapter 3 explained how the conflicts between the representational and executive failure accounts could be a consequence of different configurations in the DMN.

Since the brain works as a distributed system when engaging in tasks, a natural question following Chapter 3 is how regions other than DMN contribute to the heterogeneity of ongoing thought. Recent works on the hierarchical functional cortical organisation suggest variability in relationships among large-scale networks (e.g. Margulies et al., 2016). In Chapter 4, the focus shifted from the DMN to the whole brain in order to explore the contribution of other brain networks to the heterogeneity of ongoing experience. The identified neuroexperiential components suggested that different patterns of ongoing thought may be linked to distinctive neural hierarchies. For example, experiences characterised as purposeful monologue (see component 1 in Figure 4.2) were linked to sensory segregation—dissociation between the DMN and unimodal networks—at the whole-brain level. This pattern of decoupling is a well-documented element of ongoing thought (see Smallwood, 2013; Smallwood & Andrews-Hanna, 2013) and is thought to reduce the interference between patterns of self-generated thought and events in the external environment (Murphy et al., 2018). Consistent with the adaptive view on the sensory segregation process, this pattern

of experience was linked to better performance on measures of cognition and intelligence (Figure 4.3). The study also highlighted that dysfunctions within a second hierarchy—the MDN (Duncan, 2010)—may also make an important contribution to patterns of ongoing thought (see component 3 in Figure 4.2). Individuals whose thoughts were directed towards their personal concerns had low levels of connectivity within both the ventral and dorsal attention networks, as well as the FPN. Critically, these individuals tended to show poor performance at measures of intelligence and control. This phenotypical pattern provides evidence for the hallmarks of executive failure (i.e. McVay & Kane, 2010). In contrast to data present in Chapter 3, where a singular system gives rise to different patterns of thoughts and behaviour, the data presented in Chapter 4 shows that heterogeneous patterns of ongoing thought emerge from functional connectivity that reflects previously documented neurocognitive associations.

Multiple overlapping patterns of ongoing thought can be realised by combining experiential data with objective neurocognitive measures. The demonstrations in this thesis suggest that some of the theoretical controversy surrounding the nature of ongoing thought can be explained as relating to distinct patterns of ongoing thought. For example, in both Chapters 3 and 4, we found certain patterns of ongoing experience that have beneficial features (such as better creativity or intelligence) and others with less beneficial correlates (such as lower intelligence or worse executive control). Based on these findings, some of the controversies regarding whether mind-wandering should be considered a failure of executive control (e.g. McVay & Kane, 2010; Smallwood, 2010) result from prior studies lumping together experiential states with different features into a single category. In other words, one important contribution of this thesis is providing a synthesis to resolve the competing theoretical positions of the same phenomenon. The competing elements highlighted in the theoretical accounts can be seen to reflect independent aspects of ongoing thought. The capacity to identify these overlapping patterns of experience is made possible in part because of the use of a biological marker (i.e. functional connectivity) as well as assessments on multiple aspects of ongoing thought and task performance. Moving forward, studies of ongoing thought would, therefore, benefit from measuring multiple dimensions of experience, as well as measures of covert function such as

neuroimaging measures, pupillometry (Konishi et al., 2017), or other biological measures (Engert, Smallwood, & Singer, 2014).

6.1.2 Integration and segregation in transmodal networks

A second theme emerging from this thesis lies in the function of transmodal networks. Patterns of heterogeneity emerge through differential patterns of integration and segregation between and within large-scale neural systems. Integration and segregation have been both assumed to be an important principle in brain organisation. For example, hierarchical integration of sensory information is thought to support more abstract aspects of cognition (Mesulam, 1998). In contrast, segregating neural systems are thought to provide flexibility in the operations that can be performed (Buckner & Krienen, 2013). A hierarchical organisation implicating both integration and segregation is captured by the primary gradient which stretches from the unimodal to the transmodal networks (Margulies et al., 2016). Other examples of cognitive hierarchies that depend on integration and segregation include the MDN (Duncan, 2010). The integrated activity of large-scale network concerned with integration and segregation is important whenever individuals perform complex goal-directed tasks. Notably, the principal gradient and the MDN are differentiable in terms of the degree of separation between the DMN and FPN. The differences in functional distance indicate how patterns of integration within transmodal cortex is a defining feature of the neurocognitive hierarchies. The other hierarchy emerges from the limbic system. This limbic system hierarchy is composed of visceromotor regions that connect with DMN, and salience network (Kleckner et al., 2017). These authors suggest that this forms an allostatic-interoceptive system, segregated from the unimodal and attention systems. This hierarchy is assumed to emerge because the limbic system can selectively integrate information from systems involved in attention and cognition, as well as those important for emotion and affect. These past studies illustrate that at the core of different neural hierarchies are patterns of integration and segregation between distributed neural regions.

This thesis underscores the implication of the importance of integration and segregation in the neural hierarchies that support patterns of ongoing thought. Chapter 3 demonstrated that the role of the DMN in distinct patterns of ongoing

thought emerges because of differences in the integration and segregation within DMN, with spontaneous off-task thought linked to lower connectivity within this system, while the positive habitual thought was linked to integration within the same system. Importantly, this latter pattern of experience was also linked to stronger coupling with a number of regions outside of the DMN including left temporoparietal cortex, left hippocampus/entorhinal cortex, left lateral middle temporal gyrus, and the left pre-SMA (see Section 3.4.2 and Figure 3.4). We also found evidence for integration and segregation in Chapter 4 (see Section 4.4.1 and Figure 4.2). The pattern of purposeful future planning thought was linked to segregation between DMN and the primary sensory systems. The second pattern of thoughts reflected ongoing thoughts of personal importance and was linked to reduced connectivity (i.e. lower integration) within many regions of attention and control systems. Within both Chapters 3 and 4 patterns of ongoing thought were differentiable based on the patterns of integration and segregation between neural systems.

The most powerful evidence for integration and segregation within this thesis is provided by Chapter 5. This analysis highlighted the integration and segregation in the limbic system as a core determinant of patterns of ongoing thought. The limbic system has been previously argued to form a hierarchy integrating the transmodal system while simultaneously segregating the unimodal system to facilitate attention inhibition (Kleckner et al., 2017). In Chapter 5 we found a pattern of population variation anchored at one end by a highly inter-connected limbic system, integrating with the other transmodal area and segregated from the sensorimotor system (component 7 in Section 5.4.3). Such segregation pattern was predictive of behaviour that entailed the flexibility of retaining mental content. At the other extreme, the limbic system was highly coupled with neural system and individuals were unable to inhibit their prior mental set. Importantly, we found that this pattern of population variation was linked to patterns of ongoing thought that varied from task-focused thought at one end to personal, habitual content at the other. This analysis not only suggests that the degree of integration and segregation between the limbic system is important for attentional control (i.e. Kleckner et al., 2017). Moreover, it also suggests that the degree to which this system is coupled or decoupled from other aspects of

the cortex is a primary determinant of whether patterns of ongoing thought are focused on the task, or are instead focused on personally relevant matters in a detailed manner.

Together the three studies presented in this thesis show that at the core of different patterns of ongoing thought are the integration and segregation between neural systems. Moving forward studies should formally consider how patterns of integration and segregation between neural systems can give rise to the heterogeneity of patterns of ongoing thought that make up our daily lives.

6.2 Limitations

The primary limitation of this work is how it dealt with the temporal elements of cognition. For example, the studies focused exclusively on individual differences within a population rather than a state level of ongoing thought. Studies exploring the associations between static functional connectivity and psychological traits have brought fruitful results to ongoing thought (Smallwood et al., 2016; McVay et al., 2009; Ruby, Smallwood, Engen, & Singer, 2013). However, it is important to bear in mind that these studies confound traits with states since a defining feature of patterns of ongoing thought is their intermittent nature. While individual traits allow a way to understand links between cognition and the brain, it remains to be seen whether the patterns discovered at the population level will be applicable in the momentary state.

There are two ways that future studies could provide a more valid temporal perspective on patterns of ongoing thought. One approach would be to measure experience and neural activity on multiple occasions. One potential strategy is collecting multiple examples of experience using online probes while the neural function is recorded. Recently using the same 0-back/1-back task in conjunction with measures of neural function provided by fMRI we found that different dimensions of experience can have unique neural representations (Sormaz et al., 2018). It would be possible to apply CCA to data collected in this manner which would allow neurocognitive patterns to emerge that describe momentary states rather than population variation. A second approach would be to explore the association between patterns of dynamic neural function and ongoing experience. The recent discovery of temporal dynamic using hidden Markov models (HMM;

Vidaurre et al., 2017) demonstrated that time spent in brain states predicts behavioural traits including measures of inhibition control and attention. HMM allow the identification of temporally re-occurring states that are defined by a similarity between neural data across time. It is possible that the application of HMM to neural data would resolve covert patterns of neural function that reflect momentary changes in patterns of ongoing thought.

Another limitation emerges from the statistical technique that was applied in this thesis in terms of both model selection strategy employed. The three studies in the current thesis explored and improved the model selection strategy of SCCA. The analysis in Chapter 3 did not select the hyper-parameters in a data-driven manner. With formal hyper-parameter selection, Chapter 4 is more transparent and data-driven in the model selection. A nested cross-validation scheme was adopted to simultaneously select the hyper-parameter and the final model. With the motivation to construct a pipeline that can be generalised to the basic version of linear CCA, Chapter 5 separates the hyper-parameter selection step and the mode selection. The final canonical correlations of the principle mode improved from 0.28 in Chapter 4 to 0.70 in Chapter 5 with a simpler pipeline. The scope of the current thesis focuses on the psychological question of ongoing thought, hence the two pipelines presented in Chapters 4 and 5 are not formally bench-marked on the same data set. Future works on a structurally simple and well-performed pipeline would be important for the application of CCA and its variation on biomedical data.

The other concern is the choice of optimisation target for model selection. The current thesis uses out-of-sample explained variance as the learning target. The rationale is to maximise the potential of predictability in a wider unknown sample with the limited sample size. The alternative choice would be the out-of-sample prediction error, which minimises the mistake when applied to an unknown sample. This thesis did not explore the second option, hence the performance is unknown. These two optimisation targets are asking two fundamentally different questions—explained variance provides a more optimistic view of the model, while prediction error is more conservative. It is uncertain whether the choice of learning target should be question-driven or performance-driven. Again, a bench-marking study would be helpful to clarify the potential of the

options.

6.3 Future directions

Before concluding it is worth considering the implications for two specific areas of the study of ongoing thought. Much debate has been around the intermittent disruption caused by experience sampling methods when intending to measure the train of thought in a concurrent task (Smallwood & Schooler, 2006). This problem is especially concerning with MDES, where participants spend around a minute to report the thought rather than one or two questions that can be done in seconds. A covert marker would allow studying the ongoing thought while not interrupting the natural flow of thought.

This thesis shows that the variation in whole-brain functional hierarchy potentially supports different types of ongoing thought. If, as this thesis suggests, patterns of integration and segregation in neural activity are important aspects of different features of ongoing thought, then the covert marker could be based on patterns of functional connectivity. However, the calculation of connectivity depends on the processing of time-series data making the determination of rapid temporal changes problematic. This is compounded by the low temporal resolutions of fMRI. The application of magnetoencephalography (MEG) is possibly helpful for the determination of an online marker, given its ability to reveal neural processes at the level of milliseconds superior to fMRI.

In conclusion, the current thesis provided a proof of principle on the utility of whole-brain functional connectivity in exploring ongoing thought. It has the potential to be the covert online marker for spontaneous thought. However, with the current limitation in fMRI temporal resolution, functional connectivity calculation would be the main challenge of such application. MEG is the possible candidate method for understanding the dynamic of ongoing thought and underlying neural pattern.

6.4 Concluding remarks

This thesis set out to examine the neurocognitive mechanism of ongoing thought and establish the basic component processes to incorporate the heterogeneity of

ongoing thought. Three major questions were posed at the start of the thesis. These will now be revisited in light of the work performed.

Why does ongoing thought show both costs and benefits? Ongoing thought is a collective phenomenon with multiple types of experience each with their own associated functional outcomes at the trait level. This thesis suggests that pattern of costs and benefits related to mind-wandering may be usefully conceptualised as characterising overlapping but distinct aspects of the ongoing experience. Further work will be important to understand the underlying mechanisms that explain why these different states emerge.

Can functional neural hierarchy explain the heterogeneity? This thesis demonstrates that ongoing thoughts with different experimental profiles are associated with different neural hierarchy. Further work is suggested to incorporate the neural basis with the ongoing thought profiles at the state level to understand the dynamic of ongoing thought.

Is the family resemblance view viable for ongoing thought? Overall this thesis supports the contention that ongoing thought can be conceived of as a family of experiences with similar and overlapping features. The current thesis finds common component processes that determine population variation. Further work is necessary on the application of these findings at a state level to improve the understanding of the architecture of the component processes.

Appendix A

Chapter 3 Supplemental Materials

Adapted from the online supplemental material of:

Wang, H.-T., Poerio, G. L., Murphy, C. E., Bzdok, D., Jefferies, E., & Smallwood, J. (2018). Dimensions of Experience: Exploring the Ontology of the Wandering Mind. *Psychological Science*, 29 (1), 56-71. doi: 10.1177/0956797617728727

A.1 Questionnaires

A.1.1 Health Organization Adult ADHD Self-Report Scale

This is a self-report screening scale of adult ADHD, developed by the world health organisation (Kessler et al., 2005). This questionnaire comprises 18 questions to assess the frequency of DSM-IV Criterion A symptoms of adult ADHD. We take the average scores across all 18 questions to assess the participants ADHD tendency.

A.1.2 Autism Spectrum Quotient

The Autism Spectrum Quotient (Baron-Cohen et al., 2001) comprises 50 questions, included 10 questions measuring 5 different dimensions: social skills, attention switching, attention to detail, communication, and imagination. For each question the participant has four options: definitely agree, slightly agree, definitely disagree, and slightly disagree. Definitely agree or slightly agree responses

scored 1 point on half of the designated questions. Definitely disagree or slightly disagree responses scored 1 point on the other half of the questions. The scores of each dimension is calculated with the sum of the scores of designated the questions.

A.1.3 British Dyslexia Association Dyslexia checklist

The British Dyslexia Association Dyslexia checklist (Smythe & Everatt, 2001) comprises 15 questions to assess the tendency of dyslexic. Each answer of the questions have a designated scores. Individuals scoring less than 45 are probably non-dyslexic; individual scoring 45 – 60 shows a mild level of dyslexia; scoring above 60 suggests moderate or severe dyslexia. The score in the current study is the sum of the scores.

A.1.4 World Health Organization Quality of Life

The World Health Organization Quality of Life (WHO, 2002) assessment measures the quality of life cross-culturally. In the current study, a shorter version of the original instrument, WHOQOL-BREF, was used, as it is recommended for large research studies. WHOQOL-BREF comprises 26 questions. The assessment measure the following broad domains: physical health, psychological health, social relationships, and environment. The official scoring system can be obtained on request from the official website.

A.1.5 CES-Depression scale

The CES-Depression scale (Radloff, 1977) is a self-report scale designed to measure the symptoms of depression in the general population. The scale contains 20 questions accessing the frequency of depressive symptoms in the past one week. In the current study we used the sum of the scores as an indicator of depression.

A.1.6 State-Trait Anxiety Inventory

The State-Trait Anxiety Inventory (Spielberger, 1983) is a measure of trait and state anxiety, composing with 20 state anxiety questions and 20 trait anxiety questions. The state anxiety questions measure the level of anxiety when taking the questionnaire; the trait anxiety questions measure the general level of anxiety.

The questions are rated on a 4-point scale. Mean scores of the state and trait questions were taken as our measurement, where higher scores indicates a greater level of anxiety.

A.1.7 Ruminative Response Scale

The Ruminative Response Scale (Treyner, Gonzalez, & Nolen-Hoeksema, 2003) is a 22-question self-report measure of rumination. Rumination involves introverted focus on negative mood and was found associating with depressive symptoms and stress (Moberly & Watkins, 2008). The questions are rated on a 4-point scale. Mean scores of the questions were taken as our measurement, where higher scores indicates a greater level of rumination.

A.2 Cognitive tasks

The behavioural tasks were allocated into three sessions based on apparatus needed. Visual attention and generative semantic tasks were in session A, and semantic and episodic memory tasks were in session B and C. In each session, the first and second tasks were the mind-wandering task and the flanker task. In session B and C, the third task was the encoding and the delayed-recall phases of the word pair memory task respectively. The rest of the tasks were performed based on a pre-allocated order.

A.2.1 General apparatus of the laboratory session

In session B and C, the participants were in a sound proofed booth with a big glass window for the testers to monitor them. There were four testing spaces separated by office screen dividers. The tasks were delivered on Windows 7 computers and presented on 21 inches LCD monitors. Headsets were given to participants to deliver audio stimulus and blocking distracting noises. Participants were instructed to view the screen from a distance of 65 cm. The participants raised their hand to inform the experimenter to start each task. In session A, the visual attention tasks were delivered on a Windows 7 computer and presented on a 21 inches CRT monitor in a small room with light switch. The generative semantic tasks were delivered on a Windows 7 computer and presented on a

21 inches LCD monitor and a headset with microphone attached were used to recording verbal responses.

A.2.2 Semantic tasks

The tasks employed a 3 alternative force choice (3AFC) paradigm with the probe presented alongside the three choices among which the target was selected. There are four tasks: Relatedness Task (Word-to-Picture Matching; Word-to-Word Matching; Picture-to-Picture Matching), Identity Matching Task (Word-to-Picture Matching), Feature Matching Task, and Scrambled Picture Matching as the control task.

The unrelated distracters of each trial were selected among the targets from other trials ensuring that they were not linked to the probe. Except for the Feature Matching Task, all the tasks contain the same trial structure. Each trial started with 500ms blank screen. The three choices were subsequently presented on the bottom of the screen for 900ms. Finally the probe was presented on the top middle section of the screen. Probe and choices remained visible until participants response or for a maximum of 3 seconds. In the Feature Matching Task, the 500ms blank screen was replaced by the probe with, in bracket, the feature (cue) as criterion for the matching (colour, size, shape or texture). Probe and cue were presented for 1000ms. The three choices were subsequently presented on the bottom of the screen. Probe, cue and choices were presented as written words and remained visible until participants response or for a maximum of 3 seconds.

The stimuli employed in the tasks were selected from a larger dataset of words and photographs used in previous experiment (Davey et al., 2015; Krieger-redwood, 2012; Krieger-redwood, Teige, Davey, Hymers, & Jefferies, 2015; Whitney, Kirk, O’Sullivan, Lambon-Ralph, & Jefferies, 2012). The pictures were coloured photographs collected on internet and re-sized to fit the trial structure (200pixel, 72 dpi). All the coloured pictures and words were rated for familiarity and imageability using 7-point Likert scales. Lexical frequency count for the words was obtained by the SUBTLEX-UK database (van Heuven, Mandera, Keuleers, & Brysbaert, 2014).

For the details of the design, please refer to the online supplementary material

of (Wang, Poerio, et al., 2018).

A.2.3 Fluency Task

During Verbal Fluency (Adlam et al., 2010; Balota & Coane, 2008), participants had 1 minute to generate as many unique words as possible belonging to a semantic category (category fluency) or starting with a specific letter (letter fluency). Semantic fluency was assessed for six categories split in two blocks (Block A: fruits, vehicles, type of dogs; Block B: animals, tools, type of boats). Letter fluency was assessed for three letter cues (Block C: A, F, S). Block order was counterbalanced across participants and the order of cues within each block was randomized. Participants verbal responses were collected and the audio recordings were transcribed and scored off-line.

A.2.4 Word pair memory task

Participants also undertook a Word Pair Memory Task (WPMT) to assess episodic memory (Cairney et al., 2016). 80 words were selected from an adapted version of The University of South Florida (USF) word association, rhyme, and word fragment norms (Nelson, McEvoy, & Schreiber, 2004) to create 40 semantically unrelated cue and target word pairs (e.g. owl frame). Both the cue and target words were singular and they were matched for concreteness ($t(39) = 0.39$; $p = .696$), lexical frequency ($t(39) = -4.71$; $p = .640$), word length ($t(39) = 0.09$; $p = .933$) and number of syllables ($t(39) = -0.73$; $p = .472$). There were no pre-existing forward or backward associated relationships between any of the words, reducing the likelihood of erroneous associations between words in separate pairs.

During an initial learning phase, participants were presented with the unrelated words pairs, one at a time for 5 seconds each. This encoding phase was followed by a recall phase during which they attempted to recall the second word from the first word in the pair, they had 12 seconds for each trial and received a feedback after each response. In case of no response or error the feedback included the correct match. Participants were required to reach a performance criterion of 60% correct responses, with a maximum of three repetitions of the recall phase for the entire list of word pairs. In the subsequent behavioral testing session

that took place at least one day apart, participants attempted to recall the pairs immediately (without feedback) and provided a confidence rate about each of their responses using a 7-point Likert scale.

A.2.5 Digit span

For the Forward and Backward Digit Span Test we used the stimuli and the score procedure described in the WASI battery. For each trial, audio files of each digit were played in the sequential order reported in the WASI battery. The Forward and Backward Digit Span versions were tested in separate blocks and instructions were presented at the beginning of each block asking participants to listen to the sequence of numbers and type them in the same order, for the Forward block, or in reverse order for the Backward block.

A.2.6 Flanker task

We used the flanker task paradigm developed by (Eriksen & Eriksen, 1974) as a baseline executive measure in this study. This task is conducted at the beginning of each laboratory sessions. The target was an arrowhead at the centre, pointing to the left or right direction. This target was flanked on either side by two to four items. The items were arrows in the same direction (congruent condition), or in the opposite direction (incongruent condition), or lines (neutral condition). The participants task was to identify the direction of the centrally presented arrow by pressing the left arrow key for the left direction and the right arrow key for the right direction. The stimuli were white and displayed on a black background. Each trial lasted for 4000 msec. A trial started with a fixation period of 900 - 2100 msec and then the target and the flankers appeared simultaneously. The target and the flankers were presented until the participant responded but for no longer than 1700 msec. After the participant made a response, the target and flankers disappeared immediately and then a post-target fixation cross was presented. The duration of the post-target fixation period was based on the duration of the first fixation and RT (4000 ms minus duration of the first fixation minus RT). After this interval, the next trial began.

A.2.7 Task-switching task

We used the task-switching paradigm developed by (Mayr & Keele, 2000) and the design and task materials were constructed based on (Whitmer & Banich, 2007). This task measured executive control on inhibiting previously relevant information. In this task, the participant identified the spatial location of a deviant object with a verbal instruction cue. The participant used a number pad to respond. Number 1,2,4, and 5 were used. Each of them responded to the spatial location of the designated rectangle. In each trial, four blue rectangles arranged into a two-by-two matrix were displayed on screen. The rectangles can vary from each other on one of three dimensions: size, motion, or orientation. Before a set time interval of 100ms or 900ms, a verbal cue on dimensions appeared on the centre of the screen. There were one practice block and two experiment blocks. The cue-stimuli interval in the practice is 500 msec, and 900 msec and 100 msec respectively in the two experiment blocks. The trials are categorised into four: control, inhibitory, uncategorised switch and repeat, see (Whitmer & Banich, 2007) for details.

A.2.8 Four mountains task

We used the four mountains task developed by (Hartley et al., 2007) as a measure of spatial scene construction memory. In this task, the participant identified the target image that match the topography of the sample image across 30 trials. The participant was presented with a sample image of four mountains for 10 seconds, and then a four-choice of landscapes arranged in a two-by-two grid shown on the screen. The participant had no limit on thinking time for each trial, and they pressed number 1 to 4 to select the image. The target image is the same landscape as the sample image, but the perspective and environment (lighting, weather and vegetation) is changed.

A.2.9 Ravens advanced progressive matrices

The Ravens Advanced Progressive Matrices (RAPM Raven et al., 1998) measured educative ability that is the ability to make sense and meaning out of complex non-verbal stimuli. In order to complete the task participants were tasked with finding new patterns and relationships between the stimuli. The

RAPM used in the current study contained two tests: (i) practice test - containing 2 problems and (ii) the full test containing 36 problems. For each problem a set of 9 boxes (ordered in a 3x3 design) were shown on the screen. All but one box contained a pattern. At the bottom of the screen were 4 additional boxes, each containing a unique pattern. Participants were required to select out of these 4 potential boxes which pattern should go in the empty box. During the practice phase participants were given online feedback outlining whether their response was correct and, if not, how they should decide which box was the correct answer. If participants had any further questions, then they were instructed to ask the experimenter before starting the main experiment. During the full test no feedback was given. Participants were given 20 minutes to complete as many problems as they could, the problems got progressively more difficult.

A.2.10 Unusual uses task

The Unusual Uses Task (Guilford, 1967) assessed divergent thinking and creativity. Participants were instructed to list as many unusual uses as they can for a familiar object. Three objects were selected (newspaper, brick and shoe). Uses were considered unusual if they were not the original use of the item. For example, saying crosswords for newspaper would not be considered unusual, however saying animal bedding would. For each object, the object name appeared on screen for two minutes and participants were required to type as many unusual uses as they could. The total number of unique uses they listed for each item was calculated. Repetition of uses was not included (e.g., saying animal bedding and bedding for animal cage would only count as one unusual use). The participants creativity score was based upon the mean number of unusual uses across the three objects.

A.3 Supplementary analysis and figures

Table A.1. Correlation between motion parameter (Mean FD Jenkinson) and variable of interests

	1-back	0-back	SEM	EXE	GEN	AD	SOC	DYSL	ATT	Positive/Habit	Spontaneous off task
r	-0.140	0.161	-0.186	-0.234	-0.046	0.002	-0.007	0.043	-0.007	0.057	-0.096
p	0.080	0.044	0.020	0.003	0.566	0.981	0.931	0.597	0.928	0.514	0.271
N	157	157	157	157	157	157	157	157	157	134	134

We selected participants for whom movement greater than .2mm occurred on less than 5% of the resting state data ($N = 134$) and re-ran the SCCA.




Figure A.1. Restricted temporal sampling and bootstrapping resampling distribution with 1000 iteration.




Figure A.2. Scree plots of the principle component analysis.

See Online Supplemental Material <http://journals.sagepub.com/doi/suppl/10.1177/0956797617728727>

Figure A.3. Full set of components.




Figure A.4. Decomposition with motion outlier subjects excluded.

Appendix B

Nested K-Fold Cross-Validation

Algorithm 1 : Nested k-fold cross-validation

```
1: for Each outer fold k do
2:   for Each parameter set do
3:     Separate the development set into j folds.
4:     for Each inner fold j do
5:       Train the model on the training set
6:       Calculate test error in the validation set j
7:     end for
8:     Compute the average inner cross-validation test error
9:   end for
10:  Choose the best parameter set with minimum average test error.
11:  Use this parameter set to train on the development set.
12:  Calculate test error in the test set
13: end for
14: Determine the optimal model based on the outer fold test error
15: Train the full dataset on the optimal model
```

Bibliography

- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Muller, A., Kossaifi, J., ... Varoquaux, G. (2014). Machine Learning for Neuroimaging with Scikit-Learn. *Frontiers in Neuroinformatics*, 8, 14. doi:10.3389/fninf.2014.00014
- Adlam, A.-L. R., Patterson, K., Bozeat, S., & Hodges, J. R. (2010). The Cambridge Semantic Memory Test Battery: Detection of semantic deficits in semantic dementia and Alzheimer's disease. *Neurocase*, 16(3), 193–207. doi:10.1080/13554790903405693
- Allen, N., Sudlow, C., Downey, P., Peakman, T., Danesh, J., Elliott, P., ... Collins, R. (2012). UK Biobank: Current status and what it means for epidemiology. *Health Policy and Technology*, 1(3), 123–126. doi:10.1016/j.hlpt.2012.07.003
- Andrew, G., Arora, R., Bilmes, J., & Livescu, K. (2013). Deep Canonical Correlation Analysis. In S. Dasgupta & D. McAllester (Eds.), *Proceedings of the 30th international conference on machine learning* (Vol. 28, pp. 1247–1255). PMLR. Retrieved from <http://proceedings.mlr.press/v28/andrew13.html>
- Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., & Buckner, R. L. (2010). Functional-Anatomic Fractionation of the Brain's Default Network. *Neuron*, 65(4), 550–562. doi:10.1016/j.neuron.2010.02.005
- Andrews-Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self-generated thought: Component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, 1316(1), 29–52. doi:10.1111/nyas.12360
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8(4), 170 - 177. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1364661304000531> doi:https://doi.org/10.1016/j.tics.2004.02.010
- Bailey, K., West, R., & Kuffel, J. (2013). What would my avatar do? Gaming, pathology, and risky decision making. *Frontiers in Psychology*, 4(SEP), 609. doi:10.3389/fpsyg.2013.00609
- Baird, B., Smallwood, J., Mrazek, M. D., Kam, J. W. Y., Franklin, M. S., & Schooler, J. W. (2012). Inspired by distraction: Mind wandering facilitates creative incubation. *Psychological Science*, 23(10), 1117–1122. doi:10.1177/0956797612446024
- Baird, B., Smallwood, J., & Schooler, J. W. (2011). Back to the future: Autobiographical planning and the functionality of mind-wandering. *Consciousness and Cognition*, 20(4), 1604–1611. doi:10.1016/j.concog.2011.08.007
- Balota, D. A., & Coane, J. H. (2008). 2.28 Semantic Memory. In J. H. Byrne (Ed.), *Learning and memory: A comprehensive reference* (Vol. 2, pp. 512–

- 531). Academic Press. doi:http://dx.doi.org/10.1016/B978-012370509-9.00170-4
- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The autism-spectrum quotient (AQ): Evidence from asperger syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, *31*(1), 5–17. doi:10.1023/A:1005653411471
- Barrick, m. R., & Mount, M. K. (1991). The Big Five Personality Dimensions and Job Performance : A Meta-Analysis. *Personnel psychology*, *44*(1), 1–26. doi:10.1111/j.1744-6570.1991.tb00688.x
- Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London - Series B: Biological Sciences*, *360*(1457), 1001–1013. doi:10.1098/rstb.2005.1634
- Beckmann, C. F., Mackay, C. E., Filippini, N., & Smith, S. M. (2009). Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Neuroimage*, *47*(Suppl 1), S148. doi:10.1016/S1053-8119(09)71511-3
- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, *19*(12), 2767–2796. doi:10.1093/cercor/bhp055
- Biswal, B., Zerrin Yetkin, F., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*(4), 537–541. doi:10.1002/mrm.1910340409
- Broca, P. P. (1861). Perte de la Parole , Ramollissement Chronique et Destruction Partielle du Lobe Antérieur Gauche du Cerveau. *Bulletin de la Société Anthropologique*, *2*, 235–238. doi:10.2215/CJN.06440616
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*(1), 1–38. doi:10.1196/annals.1440.011
- Buckner, R. L., & Krienen, F. M. (2013). The evolution of distributed association networks in the human brain. *Trends in Cognitive Sciences*, *17*(12), 648–665. doi:10.1016/j.tics.2013.09.017
- Bzdok, D., Hartwigsen, G., Reid, A., Laird, A. R., Fox, P. T., & Eickhoff, S. B. (2016). Left inferior parietal lobe engagement in social cognition and language. *Neuroscience & Biobehavioral Reviews*, *68*(September 2016), 319–334. doi:10.1016/j.neubiorev.2016.02.024
- Bzdok, D., Heeger, A., Langner, R., Laird, A. R., Fox, P. T., Palomero-Gallagher, N., ... Eickhoff, S. B. (2015). Subspecialization in the human posterior medial cortex. *NeuroImage*, *106*(1), 55–71. doi:10.1016/j.neuroimage.2014.11.009
- Bzdok, D., Laird, A. R., Zilles, K., Fox, P. T., & Eickhoff, S. B. (2013). An investigation of the structural, connectional, and functional subspecialization in the human amygdala. *Human Brain Mapping*, *34*(12), 3247–3266. doi:10.1002/hbm.22138
- Bzdok, D., Langner, R., Schilbach, L., Jakobs, O., Roski, C., Caspers, S.,

- ... Eickhoff, S. B. (2013). Characterization of the temporo-parietal junction by combining data-driven parcellation, complementary connectivity analyses, and functional decoding. *NeuroImage*, *81*(1), 381–392. doi:10.1016/j.neuroimage.2013.05.046
- Bzdok, D., & Yeo, B. T. T. (2017). Inference in the age of big data: Future perspectives on neuroscience. *NeuroImage*, *155*(1), 549–564. doi:10.1016/j.neuroimage.2017.04.061
- Cairney, S., Lindsay, S., Paller, K., & Gaskell, M. (2016). Sleep preserves original and distorted memories following retrieval-induced distortion. (*under revision*).
- Chang, C., & Glover, G. H. (2010). Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *NeuroImage*, *50*(1), 81–98. doi:10.1016/j.neuroimage.2009.12.011
- Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, *106*(21), 8719–8724. doi:10.1073/pnas.0900234106
- Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: a dynamic framework. *Nature Reviews Neuroscience*, *17*(11), 718–731. doi:10.1038/nrn.2016.113
- Ciric, R., Wolf, D. H., Power, J. D., Roalf, D. R., Baum, G. L., Ruparel, K., ... Satterthwaite, T. D. (2017). Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *NeuroImage*, *154*(1), 174–187. doi:10.1016/j.neuroimage.2017.03.020
- Cooney, R. E., Joormann, J., Eugène, F., Dennis, E. L., & Gotlib, I. H. (2010, Dec 01). Neural correlates of rumination in depression. *Cognitive, Affective, & Behavioral Neuroscience*, *10*(4), 470–478. doi:10.3758/CABN.10.4.470
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215. doi:10.1038/nrn755
- Crittenden, B. M., Mitchell, D. J., & Duncan, J. (2015). Recruitment of the default mode network during a demanding act of executive control. *eLife*, *4*, e06481. doi:10.7554/eLife.06481
- Crittenden, B. M., Mitchell, D. J., & Duncan, J. (2016). Task Encoding across the Multiple Demand Cortex Is Consistent with a Frontoparietal and Cingulo-Opercular Dual Networks Distinction. *Journal of Neuroscience*, *36*(23), 6147–6155. doi:10.1523/JNEUROSCI.4590-15.2016
- D’Argembeau, A., Jeunehomme, O., Majerus, S., Bastin, C., & Salmon, E. (2015). The Neural Basis of Temporal Order Processing in Past and Future Thought. *Journal of Cognitive Neuroscience*, *27*(1), 185–197. doi:10.1162/jocn_a.00680
- D’Argembeau, A., & Van der Linden, M. (2006). Individual differences in the phenomenology of mental time travel: The effect of vivid visual imagery and emotion regulation strategies. *Consciousness and Cognition*, *15*(2), 342–350. doi:10.1016/j.concog.2005.09.001
- Davey, J., Cornelissen, P. L., Thompson, H. E., Sonkusare, S., Hallam, G., Smallwood, J., & Jefferies, E. (2015). Automatic and Controlled Semantic Retrieval: TMS Reveals Distinct Contributions of Posterior Middle Temporal

- Gyrus and Angular Gyrus. *Journal of Neuroscience*, *35*(46), 15230–15239. doi:10.1523/JNEUROSCI.4705-14.2015
- De Bie, T., Cristianini, N., & Rosipal, R. (2005). Eigenproblems in Pattern Recognition. In *Handbook of geometric computing* (pp. 129–167). Springer. doi:10.1007/3-540-28247-5_5
- Delamillieure, P., Doucet, G., Mazoyer, B., Turbelin, M. R., Delcroix, N., Mellet, E., ... Joliot, M. (2010). The resting state questionnaire: An introspective questionnaire for evaluation of inner experience during the conscious resting state. *Brain Research Bulletin*, *81*(6), 565–573. doi:10.1016/j.brainresbull.2009.11.014
- Diaz, B. A., Van Der Sluis, S., Moens, S., Benjamins, J. S., Migliorati, F., Stoffers, D., ... Linkenkaer-Hansen, K. (2013). The Amsterdam Resting-State Questionnaire reveals multiple phenotypes of resting-state cognition. *Frontiers in Human Neuroscience*, *7*, 446. doi:10.3389/fnhum.2013.00446
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, *55*(10), 78. doi:10.1145/2347736.2347755
- Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, *14*(4), 172–179. doi:10.1016/j.tics.2010.01.004
- Efron, B. (2010). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511761362
- Eickhoff, S. B., Laird, A. R., Fox, P. T., Bzdok, D., & Hensel, L. (2016). Functional Segregation of the Human Dorsomedial Prefrontal Cortex. *Cerebral Cortex*, *26*(1), 304–321. doi:10.1093/cercor/bhu250
- Ellamil, M., Fox, K. C. R., Dixon, M. L., Pritchard, S., Todd, R. M., Thompson, E., & Christoff, K. (2016). Dynamics of neural recruitment surrounding the spontaneous arising of thoughts in experienced mindfulness practitioners. *NeuroImage*, *136*(1), 186–196. doi:10.1016/j.neuroimage.2016.04.034
- Engert, V., Smallwood, J., & Singer, T. (2014). Mind your thoughts: Associations between self-generated thoughts and stress-induced and baseline levels of cortisol and alpha-amylase. *Biological psychology*, *103*, 283–291. doi:10.1016/j.biopsycho.2014.10.004
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143–149. doi:10.3758/BF03203267
- Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., ... Constable, R. T. (2015). Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, *18*(October), 1–11. doi:10.1038/nn.4135
- Fox, K. C. R., Spreng, R. N., Ellamil, M., Andrews-Hanna, J. R., & Christoff, K. (2015). The wandering brain: Meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes. *NeuroImage*, *111*(1), 611–621. doi:10.1016/j.neuroimage.2015.02.039
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678. doi:10.1073/pnas.0504136102
- Franklin, M. S., Mrazek, M. D., Anderson, C. L., Johnston, C., Small-

- wood, J., Kingstone, A., & Schooler, J. W. (2017). Tracking Distraction: The Relationship Between Mind-Wandering, Meta-Awareness, and ADHD Symptomatology. *Journal of Attention Disorders*, *21*(6), 475–486. doi:10.1177/1087054714543494
- Franklin, M. S., Smallwood, J., & Schooler, J. W. (2011). Catching the mind in flight: Using behavioral indices to detect mindless reading in real time. *Psychonomic Bulletin and Review*, *18*(5), 992–997. doi:10.3758/s13423-011-0109-6
- Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, *2*(1-2), 56–78. doi:10.1002/hbm.460020107
- Friston, K. J. (2013). Life as we know it. *Journal of The Royal Society Interface*, *10*(86), 20130475–20130475. doi:10.1098/rsif.2013.0475
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. ., Frith, C. D., & Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*(4), 189–210. doi:10.1002/hbm.460020402
- Ganis, G., Thompson, W. L., & Kosslyn, S. M. (2004). Brain areas underlying visual mental imagery and visual perception: An fMRI study. *Cognitive Brain Research*, *20*(2), 226–241. doi:10.1016/j.cogbrainres.2004.02.012
- Gilbert, S. J., Dumontheil, I., Simons, J. S., Frith, C. D., & Burgess, P. W. (2007). Comment on 'Wandering minds: the default network and stimulus-independent thought'. *Science*, *317*(5834), 43. doi:10.1126/science.1140801
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ... Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. doi:10.1038/nature18933
- Golchert, J., Smallwood, J., Jefferies, E., Seli, P., Huntenburg, J. M., Liem, F., ... Margulies, D. S. (2017). Individual variation in intentionality in the mind-wandering state is reflected in the integration of the default-mode, fronto-parietal, and limbic networks. *NeuroImage*, *146*(1), 226–235. doi:10.1016/j.neuroimage.2016.11.025
- Gorgolewski, K. J., Lurie, D., Urchs, S., Kipping, J. A., Craddock, R. C., Milham, M. P., ... Smallwood, J. (2014). 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. doi:10.1371/journal.pone.0097176
- Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., ... Petersen, S. E. (2018). Functional Brain Networks Are Dominated by Stable Group and Individual Factors, Not Cognitive or Daily Variation. *Neuron*, *98*(2), 439–452. doi:10.1016/j.neuron.2018.03.035
- Guilford, J. P. (1967). *The nature of human intelligence*. New York: McGraw-Hill.
- Gusnard, D. A., Akbudak, E., Shulman, G. L., & Raichle, M. E. (2001). Medial prefrontal cortex and self-referential mental activity: Relation to a default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(7), 4259–4264. doi:10.1073/pnas.071043098
- Handy, T. C., & Kam, J. W. Y. (2015). Mind wandering and selective attention

- to the external world. *Canadian journal of experimental psychology*, 69(2), 183–189. doi:10.1037/cep0000051
- Hardoon, D. R., Szedmak, S., & Shawe-Taylor, J. (2004). Canonical correlation analysis: an overview with application to learning methods. *Neural computation*, 16(12), 2639–2664. doi:10.1162/0899766042321814
- Hartley, T., Bird, C. M., Chan, D., Cipolotti, L., Husain, M., Vargha-Khadem, F., & Burgess, N. (2007). The hippocampus is required for short-term topographical memory in humans. *Hippocampus*, 17(1), 34–48. doi:10.1002/hipo.20240
- Hassabis, D., & Maguire, E. A. (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1521), 1263–1271. doi:10.1098/rstb.2008.0296
- Hastie, T., Tibshirani, R., & Wainwright, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*.
- Hemphill, J. F. (2003). Interpreting the Magnitudes of Correlation Coefficients. *American Psychologist*, 58(1), 78–79. doi:10.1037/0003-066X.58.1.78
- Hotelling, H. (1936). Relations Between Two Sets of Variates. *Biometrika*, 28(3/4), 321. doi:10.2307/2333955
- Huberty, C. J., & Olejnik, S. (2006). *Applied MANOVA and Discriminant Analysis*. Hoboken, NJ, USA: John Wiley & Sons, Inc. doi:10.1002/047178947X
- Iaria, G., Fox, C., Waite, C., Aharon, I., & Barton, J. (2008). The contribution of the fusiform gyrus and superior temporal sulcus in processing facial attractiveness: Neuropsychological and neuroimaging evidence. *Neuroscience*, 155(2), 409–422. doi:10.1016/j.neuroscience.2008.05.046
- Irish, M., Addis, D. R., Hodges, J. R., & Piguet, O. (2012). Considering the role of semantic memory in episodic future thinking: evidence from semantic dementia. *Brain*, 135(7), 2178–2191. doi:10.1093/brain/aws119
- Jenkinson, M., Bannister, P. R., Brady, M., & Smith, S. M. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841. doi:10.1016/S1053-8119(02)91132-8
- Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306(5702), 1776–1780. doi:10.1126/science.1103572
- Kane, M. J., Conway, A. R. A., Miura, T. K., & Colflesh, G. J. H. (2007). Working Memory, Attention Control, and the N-Back Task: A Question of Construct Validity. *Journal of Experimental Psychology: Learning Memory and Cognition*, 33(3), 615–622. doi:10.1037/0278-7393.33.3.615
- Kane, M. J., & McVay, J. C. (2012). What Mind Wandering Reveals About Executive-Control Abilities and Failures. *Current Directions in Psychological Science*, 21(5), 348–354. doi:10.1177/0963721412454875
- Kanwisher, N. (2010). Functional specificity in the human brain: A window into the functional architecture of the mind. *Proceedings of the National Academy of Sciences*, 107(25), 11163–11170. doi:10.1073/pnas.1005062107
- Karapanagiotidis, T., Bernhardt, B. C., Jefferies, E., & Smallwood, J. (2017). Tracking thoughts: Exploring the neural architecture of men-

- tal time travel during mind-wandering. *NeuroImage*, *147*(1), 272–281. doi:10.1016/j.neuroimage.2016.12.031
- Kessler, R. C., Adler, L., Ames, M., Demler, O., Faraone, S., Hiripi, E., ... Walters, E. E. (2005). The World Health Organization Adult ADHD Self-Report Scale (ASRS): a short screening scale for use in the general population. *Psychological medicine*, *35*(2), 245–256. doi:10.1017/S0033291704002892
- Killingsworth, M. A., & Gilbert, D. T. (2010). A Wandering Mind Is an Unhappy Mind. *Science*, *330*(6006), 932. doi:10.1126/science.1192439
- Kleckner, I. R., Zhang, J., Touroutoglou, A., Chanes, L., Xia, C., Simmons, W. K., ... Barrett, L. F. (2017). Evidence for a large-scale brain system supporting allostasis and interoception in humans. *Nature Human Behaviour*, *1*(5), 69. doi:10.1038/s41562-017-0069
- Knapp, T. R. (1978). Canonical correlation analysis: A general parametric significance-testing system. *Psychological Bulletin*, *85*(2), 410–416. doi:10.1037/0033-2909.85.2.410
- Konishi, M., Brown, K. S., Battaglini, L., & Smallwood, J. (2017). When attention wanders: Pupillometric signatures of fluctuations in external attention. *Cognition*, *168*, 16–26. doi:10.1016/j.cognition.2017.06.006
- Konishi, M., McLaren, D. G., Engen, H., & Smallwood, J. (2015). Shaped by the past: the default mode network supports cognition that is independent of immediate perceptual input. *PloS one*, *10*(6), e0132209. doi:10.1371/journal.pone.0132209
- Kosslyn, S. M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, *2*(9), 635–642. doi:10.1038/35090055
- Krieger-Redwood, K., Jefferies, E., Karapanagiotidis, T., Seymour, R., Nunes, A., Ang, J. W. A., ... Smallwood, J. (2016). Down but not out in posterior cingulate cortex: Deactivation yet functional coupling with prefrontal cortex during demanding semantic cognition. *NeuroImage*, *141*, 366 - 377. doi:doi.org/10.1016/j.neuroimage.2016.07.060
- Krieger-redwood, K. M. (2012). *An Investigation of Phonological and Semantic Control Using TMS and fMRI* (Unpublished doctoral dissertation). University of York.
- Krieger-redwood, K. M., Teige, C., Davey, J., Hymers, M., & Jefferies, E. (2015). Conceptual control across modalities: Graded specialisation for pictures and words in inferior frontal and posterior temporal cortex. *Neuropsychologia*, *76*, 92–107. doi:10.1016/j.neuropsychologia.2015.02.030
- Kucyi, A., Hove, M. J., Esterman, M., Hutchison, R. M., & Valera, E. M. (2017). Dynamic Brain Network Correlates of Spontaneous Fluctuations in Attention. *Cerebral Cortex*, *27*(3), 1831–1840. doi:10.1093/cercor/bhw029
- Lambon-Ralph, M. A., Jefferies, E., Patterson, K., Rogers, T. T., Lambon-Ralph, M. A., Jefferies, E., ... Rogers, T. T. (2017). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, *18*(1), 42–55. doi:10.1038/nrn.2016.150
- Leszczynski, M., Chaieb, L., Reber, T. P., Derner, M., Axmacher, N., & Fell, J. (2017). Mind wandering simultaneously prolongs reactions and promotes creative incubation. *Scientific Reports*, *7*(1), 10197. doi:10.1038/s41598-017-10616-3

- Lindquist, K. A., Wager, T. D., Kober, H., Bliss-Moreau, E., & Barrett, L. F. (2012). The brain basis of emotion: A meta-analytic review. *Behavioral and Brain Sciences*, *35*(3), 121–143. doi:10.1017/S0140525X11000446
- Maguire, E. A., & Hassabis, D. (2011). Role of the hippocampus in imagination and future thinking. *Proceedings of the National Academy of Sciences*, *108*(11), E39–E39. doi:10.1073/pnas.1018876108
- Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., ... Smallwood, J. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, *113*(44), 12574–12579. doi:10.1073/pnas.1608282113
- Marquand, A. F., Haak, K. V., & Beckmann, C. F. (2017). Functional corticostriatal connection topographies predict goal-directed behaviour in humans. *Nature Human Behaviour*, *1*(8), 0146. doi:10.1038/s41562-017-0146
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. *Science*, *315*(5810), 393–5. doi:10.1126/science.1131295
- Matthews, G., Joyner, L., Gilliland, K., Campbell, S., Falconer, S., & Huggins, J. (1999). Validation of a comprehensive stress state questionnaire- Towards a state 'big three'? In *Personality psychology in europe* (Vol. 7, pp. 335–350). Ghent, Belgium: Tilburg University.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General*, *129*(1), 4–26. doi:10.1037/0096-3445.129.1.4
- McCormick, C., Rosenthal, C. R., Miller, T. D., & Maguire, E. A. (2018). Mind-wandering in people with hippocampal damage. *Journal of Neuroscience*, *38*(11), 2745–2754. doi:10.1523/JNEUROSCI.1812-17.2018
- McVay, J. C., & Kane, M. J. (2009). Conducting the train of thought: Working memory capacity, goal neglect, and mind wandering in an executive-control task. *Journal of experimental psychology. Learning, memory, and cognition*, *35*(1), 196–204. doi:10.1037/a0014104
- McVay, J. C., & Kane, M. J. (2010). Does Mind Wandering Reflect Executive Function or Executive Failure? Comment on Smallwood and Schooler (2006) and Watkins (2008). *Psychological Bulletin*, *136*(2), 188–197. doi:10.1037/a0018298
- McVay, J. C., & Kane, M. J. (2012a). Drifting from slow to d'oh!: Working memory capacity and mind wandering predict extreme reaction times and executive control errors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(3), 525–549. doi:10.1037/a0025896
- McVay, J. C., & Kane, M. J. (2012b). Why does working memory capacity predict variation in reading comprehension? On the influence of mind wandering and executive attention. *Journal of Experimental Psychology: General*, *141*(2), 302–320. doi:10.1037/a0025250
- McVay, J. C., Kane, M. J., & Kwapil, T. R. (2009). Tracking the train of thought from the laboratory into everyday life: An experience-sampling study of mind wandering across controlled and ecological contexts. *Psychonomic bulletin & review*, *16*(5), 857–63. doi:10.3758/PBR.16.5.857
- Medea, B., Karapanagiotidis, T., Konishi, M., Ottaviani, C., Margulies, D.,

- Bernasconi, A., ... Smallwood, J. (2016). 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*. doi:10.1007/s00221-016-4729-y
- Mesulam, M. M. (1998). From sensation to cognition. *Brain*, *121*(6), 1013–1052. doi:10.1093/brain/121.6.1013
- Miller, K. L., Alfaro-Almagro, F., Bangerter, N. K., Thomas, D. L., Yacoub, E., Xu, J., ... Smith, S. M. (2016). Multimodal population brain imaging in the UK Biobank prospective epidemiological study. *Nature Neuroscience*, *19*(11), 1523–1536. doi:10.1038/nn.4393
- Miller, R. L., Yaesoubi, M., Turner, J. A., Mathalon, D., Preda, A., Pearlson, G., ... Calhoun, V. D. (2016). Higher Dimensional Meta-State Analysis Reveals Reduced Resting fMRI Connectivity Dynamism in Schizophrenia Patients. *PloS one*, *11*(3), e0149849. doi:10.1371/journal.pone.0149849
- Mills, C., Raffaelli, Q., Irving, Z. C., Stan, D., & Christoff, K. (2018). Is an off-task mind a freely-moving mind? Examining the relationship between different dimensions of thought. *Consciousness and Cognition*, *58*(2018), 20–33. doi:10.1016/j.concog.2017.10.003
- Mittner, M., Hawkins, G. E., Boekel, W., & Forstmann, B. U. (2016). A neural model of mind wandering. *Trends in Cognitive Sciences*, *20*(8), 570–578. doi:10.1016/j.tics.2016.06.004
- Moberly, N. J., & Watkins, E. R. (2008). Ruminative self-focus, negative life events, and negative affect. *Behaviour Research and Therapy*, *46*(9), 1034–1039. doi:10.1016/j.brat.2008.06.004
- Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-wandering: A review. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *67*(1), 11–18. doi:10.1037/a0031569
- Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic Memory and Beyond: The Hippocampus and Neocortex in Transformation. *Annual Review of Psychology*, *67*(1), 105–134. doi:10.1146/annurev-psych-113011-143733
- Mrazek, M. D., Smallwood, J., Franklin, M. S., Chin, J. M., Baird, B., & Schooler, J. W. (2012). The role of mind-wandering in measurements of general aptitude. *Journal of Experimental Psychology: General*, *141*(4), 788–798. doi:10.1037/a0027968
- Mueller, S., Wang, D., Fox, M. D., Yeo, B. T. T., Sepulcre, J., Sabuncu, M. R., ... Liu, H. (2013). Individual Variability in Functional Connectivity Architecture of the Human Brain. *Neuron*, *77*(3), 586–595. doi:10.1016/j.neuron.2012.12.028
- Murphy, C., Jefferies, E., Rueschemeyer, S.-A., Sormaz, M., Wang, H.-t., Margulies, D. S., & Smallwood, J. (2018). Distant from input: Evidence of regions within the default mode network supporting perceptually-decoupled and conceptually-guided cognition. *NeuroImage*, *171*(2018), 393–401. doi:10.1016/j.neuroimage.2018.01.017
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, *36*(3), 402–407. doi:10.3758/BF03195588

- Nooner, K. B., Colcombe, S. J., Tobe, R. H., Mennes, M., Benedek, M., Moreno, A. L., ... Milham, M. P. (2012). The NKI-Rockland sample: A model for accelerating the pace of discovery science in psychiatry. *Frontiers in Neuroscience*, *6*(OCT), 152. doi:10.3389/fnins.2012.00152
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.
- Peters, A. T., Burkhouse, K., Feldhaus, C. C., Langenecker, S. A., & Jacobs, R. H. (2016). Aberrant resting-state functional connectivity in limbic and cognitive control networks relates to depressive rumination and mindfulness: A pilot study among adolescents with a history of depression. *Journal of Affective Disorders*, *200*, 178 - 181. doi:10.1016/j.jad.2016.03.059
- Poerio, G. L., Sormaz, M., Wang, H.-T., Margulies, D. S., Jefferies, E., & Smallwood, J. (2017). The role of the default mode network in component processes underlying the wandering mind. *Social Cognitive and Affective Neuroscience*, *104*(7), 6430–5. doi:10.1093/scan/nsx041
- Poerio, G. L., Totterdell, P., Emerson, L. M., & Miles, E. (2016). Social daydreaming and adjustment: An experience-sampling study of socio-emotional adaptation during a life transition. *Frontiers in Psychology*, *7*, 13. doi:10.3389/fpsyg.2016.00013
- Poerio, G. L., Totterdell, P., & Miles, E. (2013). Mind-wandering and negative mood: Does one thing really lead to another? *Consciousness and Cognition*, *22*(4), 1412–1421. doi:10.1016/j.concog.2013.09.012
- Poldrack, R. A., Laumann, T. O., Koyejo, O., Gregory, B., Hover, A., Chen, M. Y., ... Mumford, J. A. (2015). Long-term neural and physiological phenotyping of a single human. *Nature Communications*, *6*(1), 8885. doi:10.1038/ncomms9885
- Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, *84*, 320–341. doi:10.1016/j.neuroimage.2013.08.048
- Pulvermüller, F. (2010). Brain embodiment of syntax and grammar: Discrete combinatorial mechanisms spelt out in neuronal circuits. *Brain and Language*, *112*(3), 167–179. doi:10.1016/j.bandl.2009.08.002
- Pulvermüller, F., & Fadiga, L. (2010). Active perception: Sensorimotor circuits as a cortical basis for language. *Nature Reviews Neuroscience*, *11*(5), 351–360. doi:10.1038/nrn2811
- Race, E., Keane, M. M., & Verfaellie, M. (2011). Medial temporal lobe damage causes deficits in episodic memory and episodic future thinking not attributable to deficits in narrative construction. *Journal of Neuroscience*, *31*(28), 10262–10269. doi:10.1523/JNEUROSCI.1145-11.2011
- Radloff, L. S. (1977). The CES-D Scale: A Self-Report Depression Scale for Research in the General Population. *Applied Psychological Measurement*, *1*(3), 385–401. doi:10.1177/014662167700100306
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682. doi:10.1073/pnas.98.2.676
- Raij, T. T., & Riekkki, T. J. (2017). Dorsomedial prefrontal cortex supports

- spontaneous thinking per se. *Human Brain Mapping*, 38(6), 3277–3288. doi:10.1002/hbm.23589
- Raven, J., Raven, J., & Court, J. (1998). *Manual for Raven's progressive matrices and vocabulary scales*.
- Richter, F. R., Cooper, R. A., Bays, P. M., & Simons, J. S. (2016). Distinct neural mechanisms underlie the success, precision, and vividness of episodic memory. *eLife*, 5, e18260. doi:10.7554/eLife.18260
- Robison, M. K., & Unsworth, N. (2018). Cognitive and Contextual Correlates of Spontaneous and Deliberate Mind-Wandering. *Journal of Experimental Psychology: Learning Memory and Cognition*, 44(1), 85–98. doi:10.1037/xlm0000444
- Roy, M., Shohamy, D., & Wager, T. D. (2012). Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends in Cognitive Sciences*, 16(3), 147–156. doi:10.1016/j.tics.2012.01.005
- Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, 52(3), 1059–1069. doi:10.1016/j.neuroimage.2009.10.003
- Ruby, F. J. M., Smallwood, J., Engen, H., & Singer, T. (2013). 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. doi:10.1371/journal.pone.0077554
- Ruby, F. J. M., Smallwood, J., Sackur, J., & Singer, T. (2013). Is self-generated thought a means of social problem solving? *Frontiers in Psychology*, 4(December), 1–10. doi:10.3389/fpsyg.2013.00962
- Rugg, M. D., & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology*, 23(2), 255 - 260. (Macrocircuits) doi:10.1016/j.conb.2012.11.005
- Sanders, J. G., Wang, H.-T., Schooler, J. W., & Smallwood, J. (2017). Can I get me out of my head? Exploring strategies for controlling the self-referential aspects of the mind-wandering state during reading. *The Quarterly Journal of Experimental Psychology*, 70(6), 1053–1062. doi:10.1080/17470218.2016.1216573
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nature Reviews Neuroscience*, 8(9), 657–661. doi:10.1038/nrn2213
- Schooler, J. W. (2002). Re-representing consciousness: Dissociations between experience and meta-consciousness. *Trends in Cognitive Sciences*, 6(8), 339–344. doi:10.1016/S1364-6613(02)01949-6
- Seli, P., Cheyne, J. A., Xu, M., Purdon, C., & Smilek, D. (2015). Motivation, intentionality, and mind wandering: Implications for assessments of task-unrelated thought. *Journal of Experimental Psychology: Learning Memory and Cognition*, 41(5), 1417–1425. doi:10.1037/xlm0000116
- Seli, P., Kane, M. J., Smallwood, J., Schacter, D. L., Maillet, D., Schooler, J. W., & Smilek, D. (2018). Mind-Wandering as a Natural Kind: A Family-Resemblances View. *Trends in Cognitive Sciences*, 22(6), 479–490. doi:10.1016/j.tics.2018.03.010
- Seli, P., Risko, E. F., Purdon, C., & Smilek, D. (2017). Intrusive thoughts: linking spontaneous mind wandering and OCD symptomatology. *Psychological Research*, 81(2), 392–398. doi:10.1007/s00426-016-0756-3

- Seli, P., Risko, E. F., & Smilek, D. (2016). On the Necessity of Distinguishing Between Unintentional and Intentional Mind Wandering. *Psychological Science*, *27*(March), 0956797616634068-. doi:10.1177/0956797616634068
- Seli, P., Risko, E. F., Smilek, D., & Schacter, D. L. (2016). Mind-wandering with and without intention. *Trends in Cognitive Sciences*, *20*(8), 605–617. doi:10.1016/j.tics.2016.05.010
- Seli, P., Smallwood, J., Cheyne, J. A., & Smilek, D. (2015). On the relation of mind wandering and ADHD symptomatology. *Psychonomic bulletin & review*, *22*(3), 629–36. doi:10.3758/s13423-014-0793-0
- Sepulcre, J., Sabuncu, M. R., Yeo, T. B., Liu, H., & Johnson, K. A. (2012). Step-wise Connectivity of the Modal Cortex Reveals the Multimodal Organization of the Human Brain. *Journal of Neuroscience*, *32*(31), 10649–10661. doi:10.1523/JNEUROSCI.0759-12.2012
- Shafiq, M. A., Tyler, L. K., Dixon, M., Taylor, J. R., Rowe, J. B., Cusack, R., ... Matthews, F. E. (2014). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) study protocol: A cross-sectional, lifespan, multidisciplinary examination of healthy cognitive ageing. *BMC Neurology*, *14*(1), 204. doi:10.1186/s12883-014-0204-1
- Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common blood flow changes across visual tasks: II. decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, *9*(5), 648–663. doi:10.1162/jocn.1997.9.5.648
- Smallwood, J. (2010). Why the global availability of mind wandering necessitates resource competition: Reply to McVay and Kane (2010). *Psychological Bulletin*, *136*(2), 202–207. doi:10.1037/a0018673
- Smallwood, J. (2013). Distinguishing how from why the mind wanders: a process-occurrence framework for self-generated mental activity. *Psychological Bulletin*, *139*(3), 519–35. doi:10.1037/a0030010
- Smallwood, J., & Andrews-Hanna, J. R. (2013). Not all minds that wander are lost: The importance of a balanced perspective on the mind-wandering state. *Frontiers in Psychology*, *4*, 441. doi:10.3389/fpsyg.2013.00441
- Smallwood, J., Brown, K. S., Tipper, C., Giesbrecht, B., Franklin, M. S., Mrazek, M. D., ... Schooler, J. W. (2011). Pupillometric evidence for the decoupling of attention from perceptual input during offline thought. *PLoS ONE*, *6*(3), 1–8. doi:10.1371/journal.pone.0018298
- Smallwood, J., Fitzgerald, A., Miles, L. K., & Phillips, L. H. (2009). Shifting moods, wandering minds: negative moods lead the mind to wander. *Emotion*, *9*(2), 271–6. doi:10.1037/a0014855
- Smallwood, J., Karapanagiotidis, T., Ruby, F., Medea, B., de Caso, I., Konishi, M., ... Jefferies, E. (2016). Representing Representation: Integration between the Temporal Lobe and the Posterior Cingulate Influences the Content and Form of Spontaneous Thought. *PLOS ONE*, *11*(4), e0152272. doi:10.1371/journal.pone.0152272
- Smallwood, J., McSpadden, M., & Schooler, J. W. (2008). When attention matters: The curious incident of the wandering mind. *Memory and Cognition*, *36*(6), 1144–1150. doi:10.3758/MC.36.6.1144
- Smallwood, J., & O'Connor, R. C. (2011). Imprisoned by the past: Unhappy moods lead to a retrospective bias to mind wandering. *Cognition & emotion*, *25*(930884466), 1481–1490. doi:10.1080/02699931.2010.545263

- Smallwood, J., O'Connor, R. C., Sudbery, M. V., & Obonsawin, M. (2007). Mind-wandering and dysphoria. *Cognition & Emotion*, *21*(4), 816–842. doi:10.1080/02699930600911531
- Smallwood, J., Ruby, F. J. M., & Singer, T. (2013). Letting go of the present: Mind-wandering is associated with reduced delay discounting. *Consciousness and Cognition*, *22*(1), 1–7. doi:10.1016/j.concog.2012.10.007
- Smallwood, J., & Schooler, J. W. (2006). The restless mind. *Psychological Bulletin*, *132*(6), 946–958. doi:10.1037/0033-2909.132.6.946
- Smallwood, J., & Schooler, J. W. (2015). The Science of Mind Wandering: Empirically Navigating the Stream of Consciousness. *Annual Review of Psychology*, *66*(1), 487–518. doi:10.1146/annurev-psych-010814-015331
- Smallwood, J., Tipper, C., Brown, K., Baird, B., Engen, H., Michaels, J. R., . . . Schooler, J. W. (2013). Escaping the here and now: Evidence for a role of the default mode network in perceptually decoupled thought. *NeuroImage*, *69*(1), 120–125. doi:10.1016/j.neuroimage.2012.12.012
- Smeeckens, B. A., & Kane, M. J. (2016). Working Memory Capacity, Mind Wandering, and Creative Cognition: An Individual-Differences Investigation Into the Benefits of Controlled Versus Spontaneous Thought. *Psychology of Aesthetics, Creativity, and the Arts*, *10*(4), 389–415. doi:10.1037/aca0000046
- Smith, S. M., Nichols, T. E., Vidaurre, D., Winkler, A. M., Behrens, T. E. J., Glasser, M. F., . . . Miller, K. L. (2015). A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature Neuroscience*, *18*(11), 1565–1567. doi:10.1038/nn.4125
- Smythe, I., & Everatt, J. (2001). *British Dyslexia Association Dyslexia checklist*.
- Sormaz, M., Murphy, C., Wang, H.-T., Hymers, M., Karapanagiotidis, T., Porerio, G., . . . Smallwood, J. (2018). The default mode network can support the level of detail in experience during active task states. *Proceedings of the National Academy of Sciences*. doi:10.1073/pnas.1721259115
- Spielberger, C. D. (1983). State-Trait Anxiety Inventory (Vol. 19) [Computer software manual]. doi:10.1037/t06496-000
- Sporns, O. (2014). Contributions and challenges for network models in cognitive neuroscience. *Nature neuroscience*, *17*(5), 652–60. doi:10.1038/nn.3690
- Stawarczyk, D., & D'Argembeau, A. (2015). Neural correlates of personal goal processing during episodic future thinking and mind-wandering: An ALE meta-analysis. *Human Brain Mapping*, *36*(8), 2928–2947. doi:10.1002/hbm.22818
- Stawarczyk, D., Majerus, S., Maquet, P., & D'Argembeau, A. (2011). Neural correlates of ongoing conscious experience: Both task-unrelatedness and stimulus-independence are related to default network activity. *PLoS ONE*, *6*(2), e16997. doi:10.1371/journal.pone.0016997
- Sui, J., Adali, T., Pearlson, G., Yang, H., Sponheim, S. R., White, T., & Calhoun, V. D. (2010). A CCA+ICA based model for multi-task brain imaging data fusion and its application to schizophrenia. *NeuroImage*, *51*(1), 123–134. doi:10.1016/j.neuroimage.2010.01.069
- Swanson, J. (2005). The Delis-Kaplan Executive Function System: A Review. *Canadian Journal of School Psychology*, *20*(1-2), 117–128. doi:10.1177/0829573506295469
- Taylor, J. R., Williams, N., Cusack, R., Auer, T., Shafto, M. A., Dixon, M., . . .

- Henson, R. N. A. (2017). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage*, *144*, 262–269. doi:10.1016/j.neuroimage.2015.09.018
- Thompson, B. (2015). The Case for Using the General Linear Model as a Unifying Conceptual Framework for Teaching Statistics and Psychometric Theory. *Journal of Methods and Measurement in the Social Sciences*, *6*(2), 30–41. doi:10.2458/azu_jmmss.v6i1.blair
- Touroutoglou, A., Hollenbeck, M., Dickerson, B. C., & Barrett, L. F. (2012). Dissociable large-scale networks anchored in the right anterior insula subserved affective experience and attention. *NeuroImage*, *60*(4), 1947–1958. doi:10.1016/j.neuroimage.2012.02.012
- Treynor, W., Gonzalez, R., & Nolen-Hoeksema, S. (2003). Rumination reconsidered: A psychometric analysis. *Cognitive Therapy and Research*, *27*(3), 247–259. doi:10.1023/A:1023910315561
- Tsvetanov, K. A., Henson, R. N. A., Tyler, L. K., Razi, A., Geerligs, L., Ham, T. E., . . . Cam-CAN (2016). Extrinsic and Intrinsic Brain Network Connectivity Maintains Cognition across the Lifespan Despite Accelerated Decay of Regional Brain Activation. *The Journal of Neuroscience*, *36*(11), 3115–26. doi:10.1523/JNEUROSCI.2733-15.2016
- Unsworth, N., & McMillan, B. D. (2013). Mind wandering and reading comprehension: Examining the roles of working memory capacity, interest, motivation, and topic experience. *Journal of Experimental Psychology: Learning Memory and Cognition*, *39*(3), 832–42. doi:10.1037/a0029669
- van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, *80*, 62–79. doi:10.1016/j.neuroimage.2013.05.041
- van Heuven, W., Mandera, P., Keuleers, E., & Brysbaert, M. (2014). SUBTLEX-UK: a new and improved word frequency database for British English. *Quarterly Journal of Experimental Psychology*, *67*(6), 11761190. doi:10.1080/17470218.2013.850521
- Vatansever, D., Bzdok, D., Wang, H.-T., Mollo, G., Sormaz, M., Murphy, C. E., . . . Jefferies, E. (2017). Varieties of semantic cognition revealed through simultaneous decomposition of intrinsic brain connectivity and behaviour. *NeuroImage*, *158*(1), 1–11. doi:10.1016/j.neuroimage.2017.06.067
- Vatansever, D., Menon, D. K., & Stamatakis, E. A. (2017). Default mode contributions to automated information processing. *Proceedings of the National Academy of Sciences*, *114*(48), 201710521. doi:10.1073/pnas.1710521114
- Viard, A., Piolino, P., Belliard, S., de La Sayette, V., Desgranges, B., & Eustache, F. (2014). Episodic future thinking in semantic dementia: A cognitive and fmri study. *PLOS ONE*, *9*(10), 1–22. doi:10.1371/journal.pone.0111046
- Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, *114*(48), 201705120. doi:10.1073/pnas.1705120114
- Villena-Gonzalez, M., Wang, H.-t., Sormaz, M., Mollo, G., Margulies, D. S., Jefferies, E. A., & Smallwood, J. (2018). Individual variation in the propensity for prospective thought is associated with functional integration between visual and retrosplenial cortex. *Cortex*, *99*(2018), 224–234. doi:10.1016/j.cortex.2017.11.015

- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature.*, *428*(6984), 748–751. doi:10.1038/nature02447
- Wang, H.-T., Bzdok, D., Margulies, D. S., Craddock, R. C., Milham, M. P., Jefferies, E., & Smallwood, J. (2018). Patterns of thought: Population variation in the associations between large-scale network organisation and self-reported experiences at rest. *NeuroImage*, *176*(1), 518–527. doi:10.1016/j.neuroimage.2018.04.064
- Wang, H.-T., Poerio, G. L., Murphy, C. E., Bzdok, D., Jefferies, E., & Smallwood, J. (2018). Dimensions of Experience: Exploring the Ontology of the Wandering Mind. *Psychological Science*, *29*(1), 56–71. doi:10.1177/0956797617728727
- Wechsler, D. (1999). *Wechsler Abbreviated Scale of Intelligence*. San Antonio, TX: The Psychological Corporation.
- Wechsler, D. (2005). *Wechsler Individual Achievement Test 2nd Edition (WIAT II)*. London: The Psychological Corporation.
- Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, *9*(7), 971–978. doi:10.1038/nm1727
- Whitmer, A. J., & Banich, M. T. (2007). Inhibition Versus Switching in Different Deficits Forms of Rumination. *Psychological Science*, *18*(6), 546–553.
- Whitney, C., Kirk, M., O’Sullivan, J., Lambon-Ralph, M. A., & Jefferies, E. (2012). Executive Semantic Processing Is Underpinned by a Large-scale Neural Network: Revealing the Contribution of Left Prefrontal, Posterior Temporal, and Parietal Cortex to Controlled Retrieval and Selection Using TMS. *Journal of Cognitive Neuroscience*, *24*(1), 133–147. doi:10.1162/jocn_a.00123
- WHO. (2002). *The World Health Organization Quality of Life*.
- Witten, D. M., Tibshirani, R., & Hastie, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, *10*(3), 515–534. doi:10.1093/biostatistics/kxp008
- Witten, D. M., & Tibshirani, R. J. (2009). Extensions of Sparse Canonical Correlation Analysis with Applications to Genomic Data. *Statistical Applications in Genetics and Molecular Biology*, *8*(1), 29. doi:10.2202/1544-6115.1470
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165. doi:10.1152/jn.00338.2011
- Zeki, S. M. (1978). Functional specialization in the visual cortex of the rhesus monkey. *Nature*, *274*, 423–428. doi:10.1038/274423a0