Mapping the Landscape of Brain Activity using Experience Sampling

By

Ian Goodall-Halliwell

A thesis submitted to the Graduate Program in Neuroscience

in conformity with the requirements for the

Degree of Master of Science

Queen's University

Kingston, Ontario, Canada

November 2023

Abstract

The thoughts we have and the activity in our brain are influenced by what we are doing. Studies within the lab and in the field suggest that tasks, whether lab-coordinated or organic, consistently push our thoughts and neural activity into “states”. Drawing on this research, this study had participants perform a testing battery made up of a diverse range of cognitive tasks which had been previously scanned in an fMRI. During the testing, participants reported their thoughts during each task. Following previous work, common patterns of thought were extracted using principal component analysis. These patterns of thought were used to predict neural activity and vice versa using simple machine learning techniques. This study found that neural activity and patterns of thought could be predicted better than chance. This study also suggests that certain brain networks, notably the default mode network, were essential for the successful prediction of thought patterns. These findings may be useful to researchers and clinicians who wish to more robustly sample cognition.

Co-Authorship

This study was conducted at Queen's University in the ThinC Lab by Ian Goodall-Halliwell, under the supervision of Dr. Jonathan Smallwood. Ian Goodall-Halliwell, Raven Wallace, Louis Chitiz, Bridget Mulholland, Dr. Brontë Mckeown, and Dr. Jonathan Smallwood contributed to study design, with consultations from Dr. Robert Leech, Dr. Adam Turnbull, Dr. Hao-Ting Wang, Dr. Brontë Mckeown, Louis Chitiz, Raven Wallace, and Bridget Mulholland. Data collection was organized by Raven Wallace and Ian Goodall-Halliwell. Formal analysis was completed by Ian Goodall-Halliwell. Figure development was completed by Ian Goodall-Halliwell, with assistance from Dr. Brontë Mckeown. The final manuscript was written by Ian Goodall-Halliwell, with valuable writing and editing contributions from Dr. Jonathan Smallwood, Dr. Jeffrey D. Wammes, and Dr. Jason Gallivan.

Acknowledgements

I would like to extend my sincere thanks to my supervisor, Dr. Jonathan Smallwood. Jonny, without your guidance through my Masters, and especially this thesis, I would have certainly found myself scrubbing the bilge at the nearest shipyard by now instead of looking at my completed thesis. I certainly couldn't have done this without you, and I specifically appreciate your patience through the arduous process of me relearning how to write good [*sic*] after so much time away from writing. I'd also like to thank my lab mates, Dr. Brontë Mckeown, Bridget Mulholland, Raven Wallace, Louis Chitiz, and Aryanna Rastan. Your help in running so many participants has been invaluable, and a special thanks to Raven for organizing so much of the testing logistics. I may have made the battery, but it was thanks to all of you that it evolved into a study. I am also grateful for everyone at ThinC lab for imbuing the past two years with some of the most fun times of my life. I would also like to personally thank Dr. Robert Leech and Dr. Brontë Mckeown for their many contributions to my analysis, without which this paper would likely be half the size. I would like to extend a final thanks to my parents, Dr. Kathleen Goodall and Brad Halliwell for their support, as well as my partner, Hannah Bannister for her thoughtful comments and care.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

mDES	Multidimensional Experience Sampling
`MDN	Multiple Demand Network
DMN	Default Mode Network
fMRI	Functional Magnetic Resonance Imaging
PCA	Principal Component Analysis
KMO	Kaiser-Meyer-Olkin test
MAE	Mean Absolute Error

Chapter 1: Introduction

The goal of my Master's is to determine whether there is a common mapping for task states that can predict across dimensions of experience and brain activity, by using machine learning to determine a relationship between task-associated mDES responses and fMRI data. mDES response data was recorded in a lab setting while participants completed tasks, each of which are associated with fMRI scans from a prior study. First, the resulting mDES response data was decomposed into components, which formed our experience sampling space. Concurrently, the task-associated brain maps were correlated to standard gradient maps to create our gradient space. Second, we then performed a permuted linear regression using the experience sampling space to predict positions of a held-out task in the gradient space, and *vice versa*. This method is known as a cross-validated approach. Third, this relationship was further explored in two ways: (i) by understanding the similarity across the two spaces we could determine which individual gradients and principal components were predictable, and (ii) by removing a brain network from the gradients, a process referred to as "lesioning", we determined which regions of the brain contained the information which contributed to the predictions.

The long-term goal of the research program is to establish this comparative framework by using brain dimensions as an organizing framework for contexts and examine the relationship between patterns of mDES results and brain activity. If we can map experience sampling onto brain states reliably, then we can infer neural activity from self-reported cognitive states. Additionally, this will expand on existing work (Konu, Turnbull et al. 2020) with a more robust task battery. Further, formalization of this mDES-gradient framework will also enable complex improvements to both the task battery and the mDES questionnaire.

The exploratory nature of this investigation does not enable any specific predictions of the outcome, but there are two general hypotheses for this analysis. The first hypothesis is that our decomposed mDES responses will be able to predict the position of brain maps in the gradient space better than chance, and the second is that our gradient space can predict the position of decomposed mDES responses in the experience sampling space.

Chapter 2: Literature Review

A core goal of cognitive neuroscience is to understand the link between neural activity and thoughts and actions. To formalize this relationship, recent work has highlighted the importance of context in the generation of thought (Shine, Breakspear et al. 2019, Smallwood, Turnbull et al. 2021) . The results of these experiments show that context, including tasks (Konu, Mckeown et al. 2021), introspective states like mind-wandering (Smallwood, Beach et al. 2008, Trautwein, Singer et al. 2016), and activities of daily life (Mckeown, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022) contribute to people's reported patterns of ongoing thought . These studies show that patterns of thought vary across situations, suggesting that measurement of cognition may be more complete if contextual variation is considered. To establish these contextual relationships, cognition needs to be sampled in a consistent way across multiple contexts and the contexts used for sampling must be strongly differentiable. Accounting for context-dependent changes in cognition may also be important for studies that test mechanistic principles underpinning cognition. The first goal of consistent sampling can be addressed through an experimental task battery, but the second goal of differentiating across contexts requires a formalized framework in which objective measures of task similarity, such as observed patterns of neural activity on fMRI, can act as objective descriptions of similarities and differences in context.

2.1 Linking Task Contexts to Ongoing Thought through Experience Sampling

A common approach to assessing the influence of context on cognition is experience sampling (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013). Previous work has linked task contexts to recognizable patterns of ongoing thought (see for reviews,

(Smallwood and Schooler 2015, Smallwood, Turnbull et al. 2021). Experience sampling is a method of assessing mental states that asks participants to describe their thoughts along a set of dimensions. The difference between this form of sampling and other self-report methods is primarily in the immediacy of the sampling; experience sampling generally occurs immediately, and is less dependent on memory compared to other common self-reporting methods (Barrett and Barrett 2001). One advantage of experience sampling is the variety of situations that can be sampled (Smallwood, Turnbull et al. 2021).There are many different methods of experience sampling, but one technique which is especially well suited for sampling ongoing thought is multi-dimensional experience sampling (mDES) (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013, Engert, Smallwood et al. 2014, Ho, Poerio et al. 2020, Konu, Mckeown et al. 2021, Mckeown, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022). mDES is a technique that measures multiple different elements of ongoing experience during tasks. It differs from conventional self-report surveys in the brevity and immediacy of sampling (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013, Engert, Smallwood et al. 2014). The questions within the mDES aim to sample a broad range of characteristics of ongoing thought, such as whether thought was directed towards the future, if thought was focused on a task, or if the emotion of the thoughts was positive or negative. mDES has been applied to studies in lab environments (Konu, Mckeown et al. 2021, Martinon, Smallwood et al. 2022), fMRI tasks (Christoff, Gordon et al. 2009, Konu, Turnbull et al. 2020), and in the real world (Mckeown, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022).

While the relationship between experience sampling and context has not been fully mapped across many cognitive states, prior studies have shown that mDES is sensitive to changes in cognition across contexts (Ho, Poerio et al. 2020, Konu, Mckeown et al. 2021,

Mckeown, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022). Previous work has found context-related differences in mDES responses across task contexts within a lab environment, such as episodic social cognition being positively associated with social cognition and self-reference tasks, and negatively associated with suspenseful movies and working memory tasks (Konu, Mckeown et al. 2021). Others have explored this thought-context relationship in the world outside of the lab, finding that episodic social cognition is positively associated with activities like texting and phone calls (Mulholland, Goodall-Halliwell et al. 2022). Based on these findings, my Master's will show that differences in mDES responses across states can be used to infer the relationship between neural patterns and thought patterns.

2.2 Understanding Task Contexts through Neural Activity

A complimentary approach to experience sampling that can be used to understand the influence of context on cognition is to scan the brains of people while they do tasks to identify the patterns of neural activity that differentiate situations. Some examples of this include the use of patterns of activity to differentiate language tasks (Bansal, Medaglia et al. 2018), or classifying whether participants were in a task condition or a control condition (Mourao-Miranda, Bokde et al. 2005). Many studies have highlighted two broad brain patterns that correlate with a wide variety of different types of situations. These two patterns are the Multiple Demand Network (MDN) and the Default Mode Network (DMN) (for reviews, see (Duncan 2010, Raichle 2015)). Many of the situations in which the MDN is active are also situations when the DMN shows lower levels of activity.

The MDN is composed of many different brain regions, including the lateral prefrontal cortex, inferior frontal sulcus, the anterior insula/frontal operculum, the dorsal anterior cingulate,

and the intraparietal sulcus. Studies focusing on the MDN have found that MDN activity generally increases as task demands increase (Camilleri, Müller et al. 2018) . It has also been shown that some regions within the MDN are modulated context-independently as task difficulty increases, while other regions are modulated depending on the specific task context (Wen and Egner 2023). Similar network-level changes have been found in single task conditions, such as task difficulty modulating activity in the MDN during image identification tasks (Duncan 2010) and visual searching tasks (Jung, Min et al. 2021). The MDN has even been implicated in some vibrotactile tasks (Woolgar and Zopf 2017). These examples indicate that the MDN is implicated in response to a broad set of tasks and contexts, with a key role in accommodating increases in task difficulty. The common association between MDN and task difficulty has led to activation of the MDN being regarded as a "task-positive" state (Cole, Reynolds et al. 2013).

The DMN, on the other hand, has considerably more nuanced roles in cognition. It is typically negatively correlated with the MDN, showing reduced activation when tasks demand more attention (Raichle, MacLeod et al. 2001). The DMN includes the anterior medial prefrontal cortex, posterior cingulate cortex and angular gyrus. The DMN has been linked to many psychological domains but is particularly linked to memory (Andreasen, O'Leary et al. 1995, Binder, Frost et al. 1999, Buckner and Wheeler 2001). Other work has expanded this linkage to involve conceptual processing (Andreasen, O'Leary et al. 1995), our sense of self (Kelley, Macrae et al. 2002), and self-awareness (Andreasen, O'Leary et al. 1995). With regard to task-associated contexts, the DMN has also been shown to correlate to off-task states in sustained attention tasks (Trautwein, Singer et al. 2016). Although the DMN is implicated in memory and conceptual processing, the opposing activity of the DMN in relation to the MDN has led to DMN activation being seen as a "task-negative" state (Fox, Snyder et al. 2005). These changes in

neural activity seen on fMRI maps across task-positive and task-negative states indicate that task states have a broad influence on activity in both the MDN and the DMN.

2.3 The Need for Low-Dimensional Representations

The influence of context on cognition is seen on multiple levels, both in subjective experience and neural activity. The degree to which context constrains measures of subjective experience and neural activity is not known, but we should expect that experience should be more similar in two tasks where the brain is broadly similar, and quite different if the brain activation is different. It should, therefore, be possible to use context in order to better map the relationship between subjective experience and neural activity. If both subjective experience and neural activity are constrained by context, it should also be possible to predict across these two aspects of cognition while they are constrained by context.

However, mapping between these high-dimensional datasets derived from experience sampling and neuroimaging techniques presents challenges. For example, subjective experience as measured by the mDES, and fMRI scans might have patterns so specific to their unique context that the raw data would hamper generalizability. Further, with so many specific associations between subjective experience and neural activity, any analysis involving multiple comparisons dramatically increases the chance of spurious results. One way to address these challenges is to reduce the dimensionality of the data and plot it into comparable lower-dimensional space. Within this low-dimensional space, each datapoint will have a location on axes defined by differentiating characteristics of the input data. In this space, the distance between points becomes a measure of similarity. Points with smaller distances between them are similar, and points with larger distances between them are dissimilar. Once datasets are

represented in this low-dimensional space, mapping between them is simplified. To evaluate the correlation between subjective experience and neural activity, it is only necessary to determine if the location of points in one space can predict the location of points in another space. This approach is similar to representational similarity analysis (Kriegeskorte, Mur et al. 2008) in that it projects data into lower-dimensional spaces to map across domains, but is extended to apply to the experience sampling data we can collect.

Dimensionality reduction techniques have emerged as valuable tools for interpreting experience sampling data (Christoff, Gordon et al. 2009, Franklin, Smallwood et al. 2016, Martinon, Smallwood et al. 2019, Konu, Turnbull et al. 2020, Turnbull, Karapanagiotidis et al. 2020, Konu, Mckeown et al. 2021, Mckeown, Poerio et al. 2021, Turnbull, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022). The reduction of complex subjective experience to principal components of experience allows for easier interpretability, as shown in decomposed experience sampling data from previous work (Fig. 1), and for the creation of experience spaces. These experience spaces use components of experience as dimensions and enable the comparison of experience samples using distance within the space.

In the same way that subjective experiences can be broken into components, fMRI data can be broken into patterns of activity. In fMRI analysis, dimension reduction is a common practice due to the extremely high dimensionality of scans. One approach that can be used to form a common space is to use axes defined by functional gradients, such as the gradient of Unimodal-Transmodal activity found by Dr. Margulies (Margulies, Ghosh et al. 2016), to describe patterns of activations and deactivations evoked by specific tasks (for reviews, see (Bernhardt, Smallwood et al. 2022). For example, this approach demonstrated that across multiple tasks, regions associated with unimodal processing exhibit high levels of task-related

activity, whereas those associated with transmodal processing show low levels of task-related activity (Ito, Hearne et al. 2020). Similarly, Lanzoni et al. investigated semantic cue integration using task-based fMRI and found maximal activation for their convergent cue condition at the heteromodal default mode end of the principal gradient(Lanzoni, Ravasio et al. 2020). These findings suggest that the default mode network, situated at the top of the unimodal-transmodal gradient, supports information integration processes that influence ongoing cognition. An additional study investigating semantics contextualized task-based fMRI patterns during a semantic judgment task and demonstrated the strongest response when input overlaps with long-term memory, particularly at the transmodal apex (Wang, Margulies et al. 2020). Similar to these studies, we use this principal gradient to facilitate mapping fMRI data to behaviour, thereby shedding light on the associations between task contexts, patterns of ongoing thought, and neural activity.

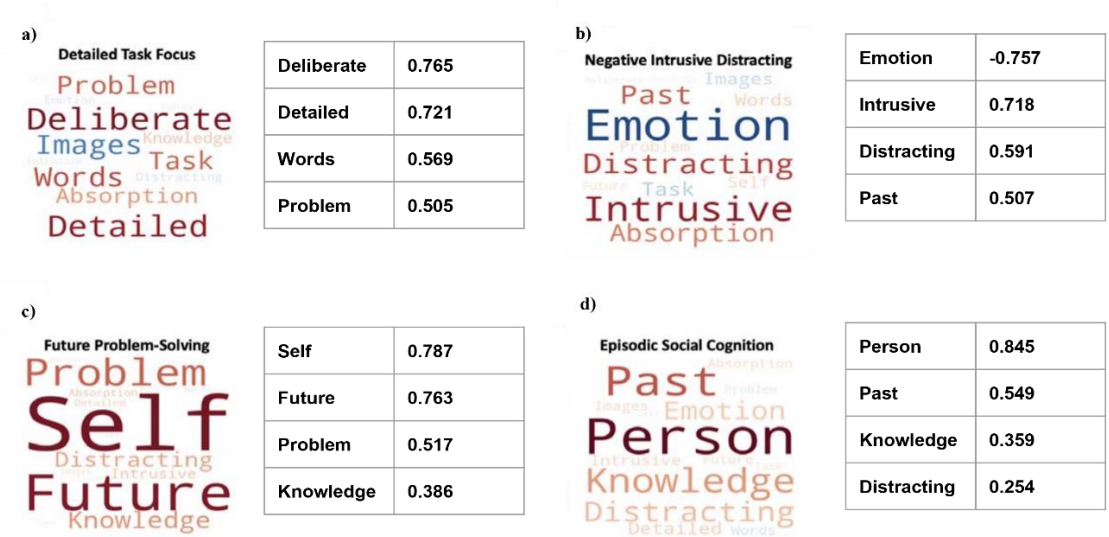

Deliberate	0.765
Detailed	0.721
Words	0.569
Problem	0.505

Emotion	-0.757
Intrusive	0.718
Distracting	0.591
Past	0.507

Self	0.787
Future	0.763
Problem	0.517
Knowledge	0.386

Person	0.845
Past	0.549
Knowledge	0.359
Distracting	0.254

Figure 1. Example experience sampling decomposition. Data was taken from previous work that used experience sampling (Mulholland, Goodall-Halliwell et al. 2022). Results originate from a PCA decomposition of mDES data. Named wordclouds are shown next to the 4 largest loadings from each component. Words with larger text in wordclouds have larger loadings, and the colour of the text indicates the direction of the loadings. High loadings are red, and low loadings are blue.

2.4 Dimension Reduction and Stable Patterns of Thought

As described previously, the dimensionality of mDES data can easily be decomposed using techniques like principal component analysis (PCA). Reducing the dimensions of mDES data using PCA organizes cognition into more common and generalizable patterns of thought. Using mDES, stable patterns of thought have been identified in several different task contexts (Fig. 2;(Smallwood, Turnbull et al. 2021)). mDES during an n-back task has generated 4 component patterns of ongoing thought, including task-relatedness, detail, modality, and emotion (Turnbull, Karapanagiotidis et al. 2020), which were previously found to be consistent across both scanning and laboratory conditions (Turnbull, Wang et al. 2019). Further work with a wider array of tasks has shown similar component patterns of ongoing thought (Konu, Mckeown et al. 2021). Key among these examples is that alterations in the task resulted in changes in individual component loadings (Turnbull, Wang et al. 2019, Turnbull, Karapanagiotidis et al. 2020, Konu, Mckeown et al. 2021). The magnitude of differences in mDES patterns across different contexts are also generally related to the magnitude of differences between the cognitive demands of the contexts. This relationship is reflected in the ranking of tasks based on the similarity of their mDES results at a component level (Konu, Mckeown et al. 2021). Within a common space created by PCA decomposition, tasks can be differentiated by their relative Euclidian distance from other tasks. For example, tasks involving detailed thinking, such as a two-back task, would more strongly load on the PCA component describing detailed thought, while tasks with less requirement for detailed thinking, such as the zero-back task, would load negatively on that PCA component. Creating a common space of task-related mDES components enables a simpler comparison of the impact of context on subjective experience. For example, mDES collected during similar tasks would have more similar PCA scores and would therefore be closer in the

space. The reduction of mDES dimensions into a space allows for these representations of context similarity.

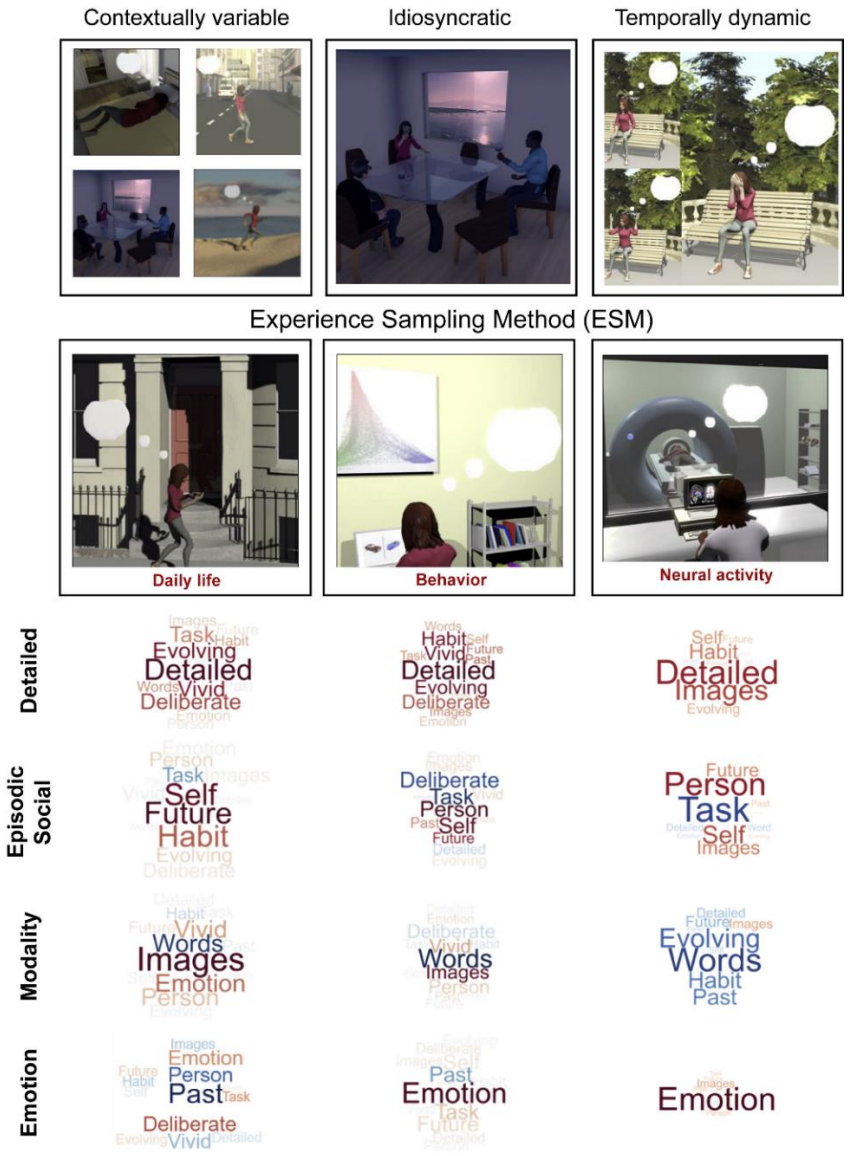

Figure 2. Heterogeneity in the patterns of ongoing thought from previous work (Smallwood, Turnbull et al. 2021). Upper panel. Using experience sampling to understand patterns of experience. Ongoing experience varies across people, places, and time. Studies use experience sampling to gain better insight into the patterns of experience that are common in different individuals and in different situations, including in daily life. These can be related to objective behavioral and neural signatures recorded under laboratory conditions to help understand the neural contributions to different patterns of ongoing thought. Lower panel. Consistency in the latent structure of thought patterns as revealed through the decomposition of mDES data. The left-hand panel shows the results of the application of the same decomposition algorithm (principal components analysis, PCA) to sets of experience sampling questions recorded in different situations. It can be seen that this approach reveals a similar structure in reports across situations and that this structure is correlated across individuals. In this figure, the word clouds show the importance of the item through the font size (bigger = more important) and the direction of the association through the color (warmer colors = positive, cooler colors = negative).

2.5 Dimension Reduction as a Tool to Compare Cognition Across Domains.

Dimension reduction is also commonly used to simplify spatial maps of brain activity (Cohen, Daw et al. 2017, Tang, Chen et al. 2021). There are many different techniques, ranging from simple linear approaches like PCA (Hong, Xu et al. 2020), to more complex nonlinear methods like Laplacian Eigenmaps (Belkin and Niyogi 2003) and diffusion embedding (Coifman and Lafon 2006). Of these techniques, diffusion embedding was noted to explain greater variance in higher order gradients (Hong, Xu et al. 2020). This technique creates a graph to represent similarities between brain regions and uses a diffusion process to identify key variations in brain activity. The resulting simplified representations, called diffusion embeddings, help visualize the data more easily. This technique was used to identify the principal gradients of macroscale cortical organization (Margulies, Ghosh et al. 2016). The principal gradient maps were extracted from resting state data and broadly represent common patterns of brain activity. These common patterns are not unlike the components that are extracted from a PCA in that they maximize the explained variance of the original data.

The gradients found by Margulies et al. have been used as standard maps to differentiate on-task versus off-task states (Turnbull, Wang et al. 2019). The gradients have also been used to organize both task (Turnbull, Karapanagiotidis et al. 2020) and resting state (Karapanagiotidis, Vidaurre et al. 2020) brain maps into a state space (Fig. 3a and 3b, respectively). These examples illustrate that the gradients are effective standard maps to compare task-associated context maps in order to visualize the differentiation of individual tasks. This approach allows for a more generalizable representation of the differences between context-associated brain maps. In this way, dimension reduction creates a common space, where maps can be compared to each other

based on distance within the space. The projection of brain maps into a standard space allows for contextual variation to be quantified.

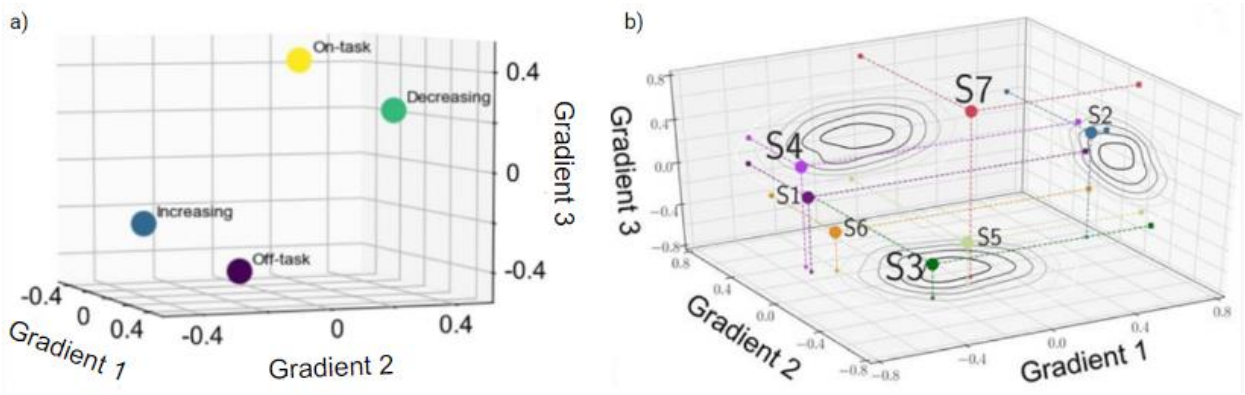

Figure 3. Gradients organizing on-task and off-task thought. a) This figure from Turnbull (Turnbull, Karapanagiotidis et al. 2020) shows that gradients 1, 2, and 3 (referred to here as hierarchies) can differentiate on-task from off-task thought, and that these align with regions that decrease after a task, and increase after a task, respectively. These maps can be represented in a three-dimensional space based on each hierarchy (top right). In this plot, neural activity patterns that decline over time (green), and patterns seen during on task states (yellow), are in the top right quadrant. In contrast, patterns that increase over time (blue) and those associated with off-task thought (purple) are in the bottom right quadrant. b) This figure from Karapanagiotidis (Karapanagiotidis, Vidaurre et al. 2020) shows the spatial similarity between the spatial maps describing gradients 1-3 and the set of spatial maps generated from hidden Markov modelling decomposition of the resting state data. In these figures, the contour plots projected onto each plane describe the distribution of the synthetic states generated through permutation.

When comparing brain maps, one approach is to correlate brain maps to each other and use that correlation as an indication of their relative difference. When examining many brain maps however, this approach can be limited by the maps used during the initial construction of the space. Specifically, the space could lack generalizability since its dimensions will fit the specific data upon which it was built. This often presents no issue, and can lead to great representations of the differences in input data, as demonstrated by prior work (Shine, Breakspear et al. 2019), which found that PCAs performed on timeseries of neural activity could effectively differentiate task states. An alternative approach is to correlate brain maps of interest to standardized templates which are known to differentiate many brain states. The limitation of

using methods relying on the structure of the input data, which comparison to standard maps overcomes, is that the structure of the low-dimensional representation is entirely dependent on the input data. By using external standard maps to construct our space, and placing input data into this constructed space, we can be sure that the axes of the space will always remain consistent across different experiments. In other words, a brain map's location in this "standard space" will always be the same regardless of what other maps are being placed into the space. Using maps external to the dataset can help determine how well a selection of maps represents broader cognition. In other words, without using external maps, the degree to which the sampled maps represent all of cognition, or all possible maps, is undeterminable. The choice of standard maps can be challenging, because they need to account for activity which may not be seen in the sampled maps and need to be descriptive of a wide number of cognitive states. These challenges can be met using the principal gradients extracted by Margulies (Margulies, Ghosh et al. 2016).

Chapter 3: Materials and Methods

3.1 Participants

194 participants were recruited to take part in the study. The sample size is guided by the sample sizes of prior studies in the literature that have investigated differences in ongoing thought across easy and hard task contexts (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013, Turnbull, Karapanagiotidis et al. 2020). This sample was drawn from the undergraduate student population at Queens' University. The study was approved by the institutional ethics committee at the Queens' University Psychology Department. All volunteers provided informed written consent and received course credits for their participation. Of these participants, 164 identified as women, 24 as men, and two as non-binary. Mean age of participants was 18.56 years (SD = 1.09). Age range was from 17 to 24. Four participants declined to provide demographic information.

3.2 Multidimensional Experience Sampling (mDES)

Participants' ongoing thought was measured using mDES. Each mDES probe was made up of sixteen items (see Table 1), presented in a random order. All items were rated by participants on a continuous scale of one to ten, using the left and right arrow keys to move a marker on a screen. The marker start location was random for each item. When participants were satisfied with the marker position on the scale, they pressed "Enter" to confirm their answer and move on to the next item. mDES probes were presented at the end of each task block. There were three blocks per task, apart from the two-back and movie tasks which each had two. Overall, participants completed a total of 38 probes. Every block of each task was followed by mDES probes, except for the movie task, which also had a probe during the task.

Table 1.

mDES items with scale and label.

Item	Label	Scale Low	Scale High
My thoughts were focused on an external task or activity:	Task	Not at all	Completely
My thoughts involved future events:	Future	Not at all	Completely
My thoughts involved past events:	Past	Not at all	Completely
My thoughts involved myself:	Self	Not at all	Completely
My thoughts involved other people:	People	Not at all	Completely
The emotion of my thoughts was:	Emotion	Negative	Positive
My thoughts involved images	Images	Not at all	Completely
My thoughts were detailed and specific:	Detailed	Not at all	Completely
My thoughts were:	Deliberate	Spontaneous	Deliberate
I was thinking about solutions to problems (or goals):	Problem	Not at all	Completely
My thoughts were intrusive:	Intrusive	Not at all	Completely
My thoughts contained information I already knew (e.g., knowledge or memories):	Knowledge	Not at all	Completely
I was absorbed in the contents of my thoughts	Absorption	Not at all	Completely
My thoughts were distracting me from what I am doing	Distracting	Not at all	Completely
My thoughts involved words	Words	Not at all	Completely
My thoughts involved sounds	Sounds	Not at all	Completely

3.3 Task Battery Procedure

All 14 tasks were presented using PsychoPy3 (Peirce, Gray et al. 2019). Each task included three task blocks for every task except for the movie and two-back tasks, which had two blocks. Each task lasted 90s per block, with a task runtime which was jittered (offset) by +/- 15s. Task blocks were also grouped into meta-blocks based on task similarity (Table 2). Task blocks were randomized within meta-blocks, and the meta-blocks were randomized based on a unique seed generated for each participant. Written instructions were presented at the start of each block (Fig. 4).

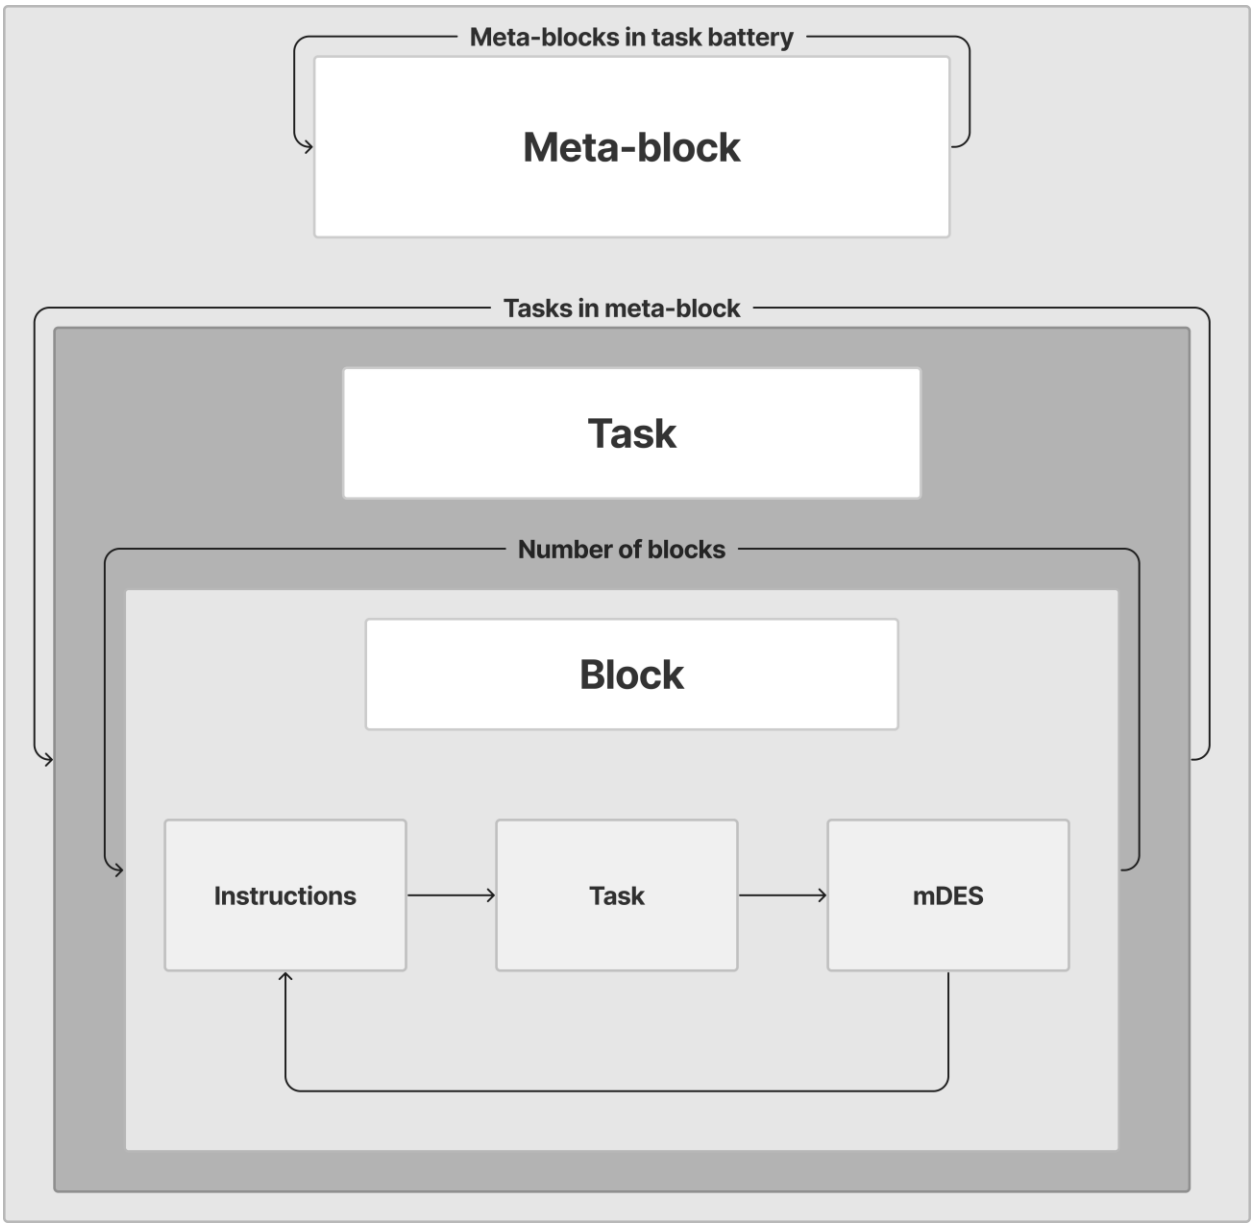

Figure 4: Task battery schematic. A diagram of the task battery workflow is shown, where each block is shown to include instructions, the task, and mDES sampling. Each block is shown to repeat, and tasks are cycled for each meta-block.

Table 2.

Tasks and their associated meta-blocks.

Meta-block name	Task name
Self-Other Task	Self Task (Konu et al., 2021)
	Other Task (Konu et al., 2021)
Go-NoGo/Finger Tapping Task	Go-NoGo Task (Konu et al., 2021)
	Finger Tapping Task (Livesey et al., 2007)
Reading/Memory Task	Reading Task (Zhang et al., 2022)
	Memory Task (Zhang et al., 2022)
n-Back Task	Zero-Back Task (Turnbull et al., 2019a)
	One-Back Task (Turnbull et al., 2019a)
Math Task	Easy Math Task (Wang et al., 2020)
	Hard Math Task (Wang et al., 2020)
Movie Task	Movie Task - Bridgeville (Konu et al., 2021)
	Movie Task - Inception (Konu et al., 2021)
Two-Back Task	Two-Back Task – Faces (Barch et al., 2013)
	Two-Back Task – Scenes (Barch et al., 2013)

3.4 PsychoPy Tasks

3.4.1 Self-Other Tasks

This task was based on a task used by Dr. Konu and colleagues (Konu, Mckeown et al. 2021) and was similar to other self-reference effect paradigms (Craik, Moroz et al. 1999, Vanderwal, Hunyadi et al. 2008). Participants were instructed to attend to the centre of the screen, where adjectives were presented one word at a time. Each trial consisted of a probe signalling the participant to judge the text stimuli in accordance with the referent (self or other).

Participants had to indicate whether they would associate the word presented with a referent or not using the left and right arrow keys. For the “other” condition, participants were instructed to think of a single friend. Each block consisted of participants making judgements about themselves (self condition) or a friend (other condition). The words used in this task paradigm were selected from a list of normalised personality trait adjectives with the highest meaningfulness ratings from Anderson and colleagues (Anderson 1968) and de Caso and colleagues (de Caso, Karapanagiotidis et al. 2017, de Caso, Poerio et al. 2017). Participants saw a unique list of words in each block.

3.4.2 Go-NoGo Task

This task was based on a similar task used by Dr. Konu (Konu, Mckeown et al. 2021) and consisted of a series of presentations of a rhombus interspersed with fixation crosses. Each rhombus was presented for 1.5s, jittered by 0.5s. The participants were instructed to press the space bar if the rhombus was not skewed (Go), and not press the space bar if the rhombus was skewed (NoGo). A random 15% of rhombus presentations were skewed for each participant.

3.4.3 Finger Tapping Task

This task was adapted from a similar task used by Dr. Livesey (Livesey, Wall et al. 2007). Participants were instructed to press the space bar when a square was shown on the screen. A fixation cross appeared between square presentations. Squares were presented for two seconds, and failure to press the space bar on time was recorded as a failure. The delay between presentations was jittered, with a minimum delay being two seconds and a maximum delay being five seconds.

3.4.4 Reading and Memory Tasks

Both the reading and memory portions were adapted from tasks used by Dr. Zhang (Zhang, Bernhardt et al. 2022). In the reading task, participants were asked to read the words presented on the screen. The words came from a series of full sentences, provided by Zhang, broken down into individual words and shown sequentially. The memory task consisted of asking participants to first think of an event to remember that was related to a word shown on the screen, and then to remember the event they had selected for 40s, jittered by five seconds.

3.4.5 n-Back Tasks

The one-back and zero-back tasks were provided by Dr. Turnbull (Turnbull, Wang et al. 2019). Both consisted of a presentation of a series of shapes on the left and right side of the screen. Participants were asked to pay attention to the shapes presented, and when given a prompt consisting of a shape in the center of the screen, to press either the left or right arrow key to select which side either currently has the correct shape on it (zero-back) or had the correct shape on it one iteration ago (one-back). If the participant did not respond in 5s, then the trial was counted as an incorrect response and the task continued.

3.4.6 Maths Task

The maths task was largely unaltered from Dr. Wang's original code (Wang, Margulies et al. 2020). The prompts were split into easy/hard versions. The maths task consisted of several different addition equations with either two terms less than ten (easy) or two terms larger than ten (hard). Each equation was presented for three seconds, followed by a fixation cross, and then two numbers were shown. The participant was instructed to press the left arrow key if the left

answer was correct and the right arrow key if the right answer was correct. If the participant did not respond within 5s, the task continued, and the response was counted as incorrect.

3.4.7 Movie Task

The movie task consisted of subjects being presented a short (~90s) clip from either a suspenseful movie (Inception), or a documentary (Welcome to Bridgeville). The presentation was counterbalanced within subject. The videos were adapted from clips outlined in Dr. Barch's experiment (Barch, Burgess et al. 2013). The volume of the clips was balanced to reduce large spikes in noise, and the volume of the testing machines was kept consistent. Each movie presentation was interrupted at a random point after the first 15s and before the last 15s with an mDES prompt, with another shown at the conclusion of the clip.

3.4.8 Two-Back Task

In this task, participants were presented with a series of stimuli, either human faces or environmental scenes, one at a time. They were instructed to indicate when the current stimulus matched the one presented two stimuli back. The images and task paradigm were pulled from the design used in the Human Connectome Project (Barch, Burgess et al. 2013).

3.5 Brain map sourcing

3.5.1 Gradient Maps

Gradient brain maps were received from Dr. Margulies' extraction (Margulies, Ghosh et al. 2016), available at the following link. These maps were extracted from the resting data from the Human Connectome Project (Van Essen, Smith et al. 2013) using diffusion embedding. The gradient maps are formally labelled numerically, but they are given names here based on the

regions that dominate either end of the dimension. Gradient 1 separates unimodal activation from transmodal activation. Gradient 2 separates the visual system from the sensorimotor system. Gradient 3 separates activity from task-negative (DMN) from task-positive (MDN). Gradient 4 separates parietal and lateral frontal regions of the DMN from temporal and medial regions of the same system. Gradient 5 separates limbic and dorsal attention networks. These can be seen in Figure 5.

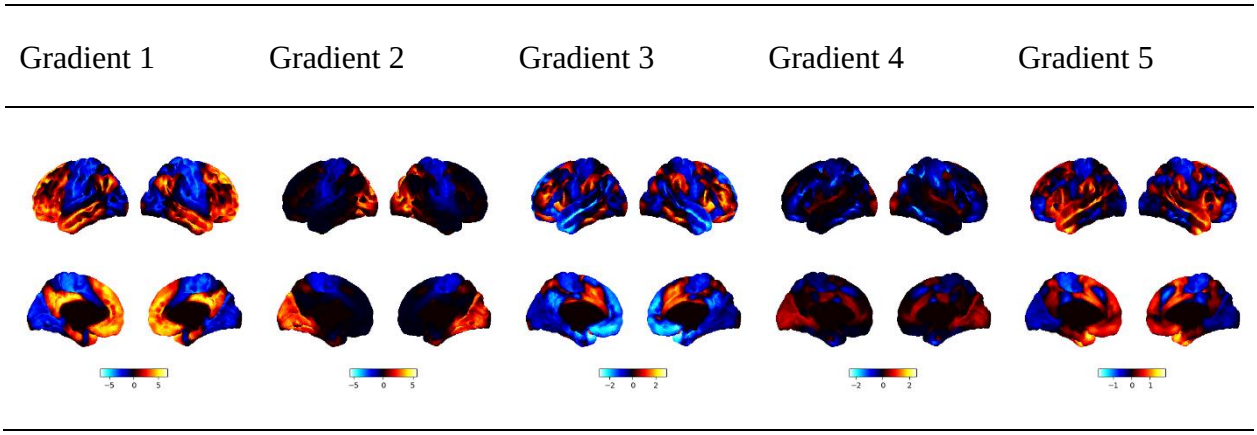

Figure 5: fMRI maps from previous gradient analyses (Margulies, Ghosh et al. 2016). First five functional gradients are listed.

3.5.2 Task Maps

The maps we selected were chosen to maximize differentiation across our gradient axes. This allowed us to more robustly sample possible cognitive states. Task maps were selected from numerous sources, listed in Table 3. All maps were z-stat maps, group-averaged, and contrasted the task condition against rest. Maps were not thresholded due to the nature of our analysis focusing on whole brain activity. Since the coverage of task maps differed, we created a binarized mask which removed any regions in which there were missing data in any of the task maps and applied this mask to all maps prior to analysis.

Table 3

Task maps and sources

Task name	Paper of origin	Task Description	Surface map
Self Task	(Murphy, Poerio et al. 2019)	Participant is shown an adjective and answers whether the adjective applies to themselves or not.	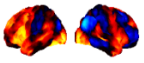
Other Task	(Murphy, Poerio et al. 2019)	Participant is shown an adjective and answers whether the adjective applies to a chosen friend or not.	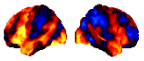
Go-NoGo Task	(Alam, Murphy et al. 2018)	Participant is shown a target stimulus and responds with a button press when the stimulus is presented. However, if the stimulus is deformed the participant does not respond with a button press.	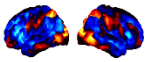
Finger Tapping Task	(Barch, Burgess et al. 2013)	Participant responds to a stimulus presentation with button press as soon as it is presented.	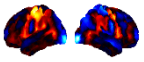
Reading Task	(Zhang, Bernhardt et al. 2022)	Participant will read sentences presented word by word in the center of the screen.	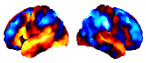
Memory Task	(Zhang, Bernhardt et al. 2022)	Participant is asked to recall a queued autobiographical memory.	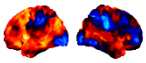
Zero-Back Task	(Turnbull, Wang et al. 2019)	When prompted, participant must select whether a shape on the left or right side of the screen matches one in the center.	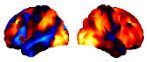
One-Back Task	(Turnbull, Wang et al. 2019)	When prompted, participant must select whether a shape on the left or right side of the screen from a previous slide matches one in the center.	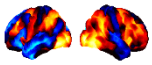
Easy Math Task	(Wang, Margulies et al. 2020)	Participant must select the correct answer to single digit addition problems.	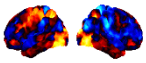
Hard Math Task	(Wang, Margulies et al. 2020)	Participant must select the correct answer to double digit addition problems.	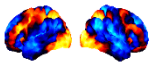
Movie Task Bridgeville	(Barch, Burgess et al. 2013)	Participants must watch a short documentary clip.	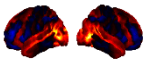
Movie Task Inception	(Barch, Burgess et al. 2013)	Participants must watch a thrilling movie clip.	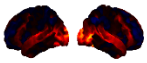
Two-Back Task Faces	(Barch, Burgess et al. 2013)	Participant must determine whether a stimulus (face) had been shown 2 trials ago.	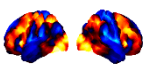
Two-Back Task Places	(Barch, Burgess et al. 2013)	Participant must determine whether a stimulus (scene) had been shown 2 trials ago.	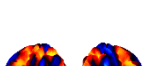

Chapter 4: Results

4.1 Creating Comparable Spaces

A total of 194 participants generated 7,334 mDES probes. All participants completed the entire battery. Average responses on each mDES item were calculated with standard deviations and are presented in Figure 6a. In order to map between experience sampling and brain activity, the mDES responses and fMRI maps needed to be put into a comparable space. In preparation for the conversion of mDES responses to a common space, PCA was performed on the data. The mDES data had a KMO value of 0.798 and passed Bartlett's test of sphericity. Bartlett's test of sphericity tests the hypothesis that the correlation matrix is an identity matrix, which would indicate that the variables are unrelated and therefore unsuitable for structure detection (Shrestha 2021). The KMO statistic is a measure of the proportion of variance among variables that might be common variance. The higher the proportion, the higher the KMO-value, the more suited the data is to factor analysis (Shrestha 2021). A KMO value above 0.6 is acceptable, and 0.8 is considered "meritous" (Nunnally 1978), so our data is acceptable. An analysis of the scree plot showed 4 components with eigenvalues greater than one (Fig. 6b). A varimax rotation was applied in order to simplify the resulting components and increase their interpretability in preparation for creating a space.

Based on eigen-values > 1, four components, accounting for 48.8% of the total variance, were retained for further analysis: 1) "Past Knowledge"— describing patterns of thought with the highest loadings on the items "Past", "Knowledge", and "Self"; 2) "Distracting Intrusive" — with the highest loadings on the items "Distracting", "Intrusive", and "Future"; 3) "Deliberate Problem Solving" — with the highest loadings on the items "Problem", "Deliberate", and "Detailed"; 4) "Sensory Engagement" — with the highest loadings on the items "Sounds",

"Images", and "Emotion". Item loadings on these components are presented as wordclouds in Figure 6c.

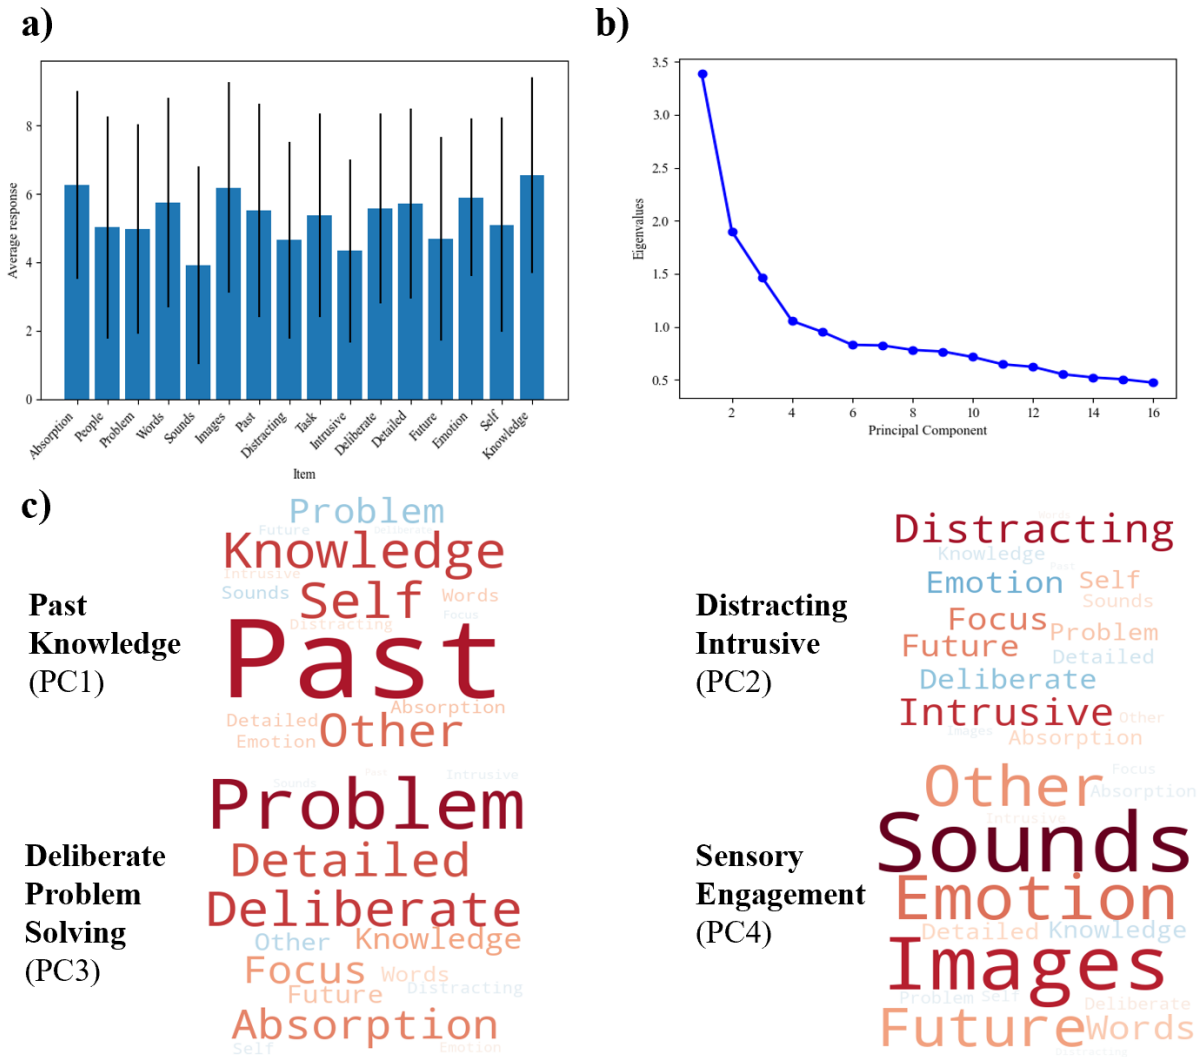

Figure 6: mDES Response Decomposition. a) Bar plot showing the means and standard deviations (error bars) of each item in the mDES. b) Scree plot of the PCA run on mDES data. Four components are above an eigenvalue of 1. c) Wordclouds representing the item loadings on the four patterns of thought identified in the thought data collected during the testing battery (n=7,334) using PCA. Each word represents an experience-sampling item (16 items). Font size represents the magnitude of the loading, and the color describes the direction. Warm colors reflect positive loadings while cool colors reflect negative loadings.

The averaged PCA loadings for each task illustrate the contributions of each task to our component structure (Fig. 7). Within component 1, our Past-Knowledge component, we see notably high scores on the Self, Other, and Memory tasks. In this component, we also see lower scores on the Go-NoGo task. Within component 2, our Distracting-Intrusive component, we see

generally more stable scores, with the other task having less distraction, while the Zero-back task has more. Within component 3, our Deliberate-Problem Solving component, we see higher loadings on the Easy and Hard Maths tasks, with lower loadings on both the suspenseful and documentary Movie tasks. The final component, Sensory Engagement, has high loadings on the Other, Memory, and Movie tasks, with lower loadings on the Math tasks.

In order to relate the brain maps to the mDES PCAs, they needed to be put into a comparable space. Using a similar strategy to that used by Turnbull (Turnbull, Wang et al. 2019), the brain maps were correlated to gradient maps using the StateSpace python package. The mean task-associated decomposed mDES responses are shown in Figure 7, and the raw results of the brain map dimensionality reduction are show in Figure 8. There are no standard deviations shown in Figure 8 because we had only one brain map for each task. These decompositions allowed for the task-associated brain maps to be plotted in a state space, where the distance between any given points in the space is representative of the difference between the maps. Tasks which have very similar brain maps will be very close within the space, while very different maps will have a larger distance between them (Fig. 9b). The decomposed mDES data are shown plotted in a similar space (Fig. 9a). The organization of both the task-associated mDES PCAs and the task-associated brain maps into two low-dimensional spaces facilitated the creation of similar representations.

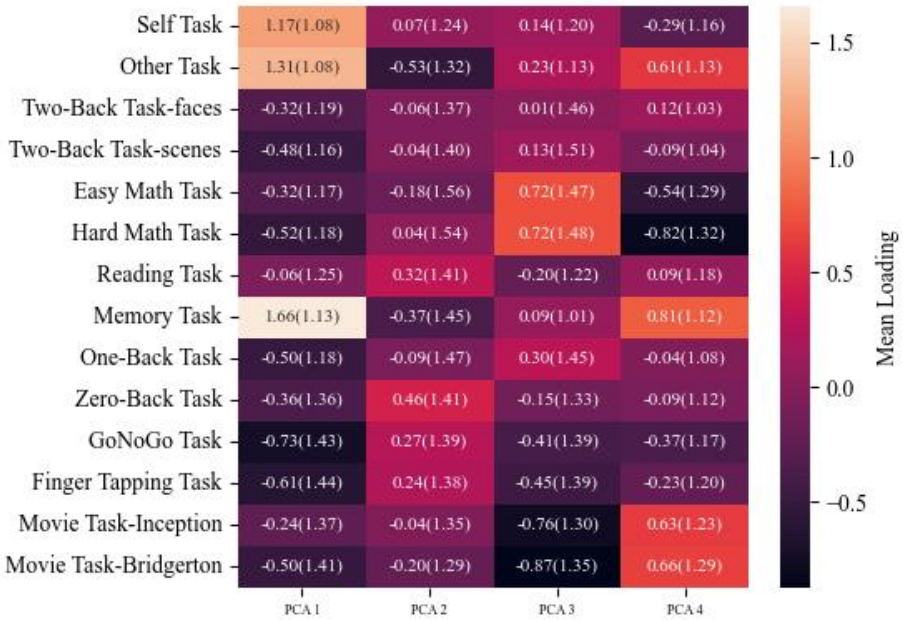

Figure 7: Mean task component loadings. Mean of all participant mDES responses after decomposition by PCA, grouped by task. Standard deviation is shown next to means in brackets. High scores denote that mDES reports during a specific task are highly associated with the specified PCA component.

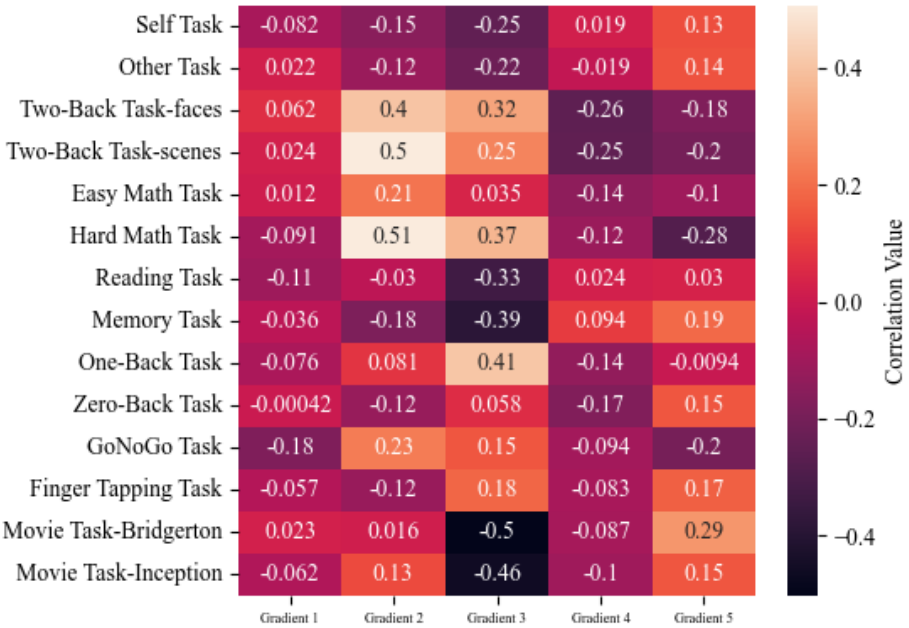

Figure 8: Task gradient correlations. Displayed scores are the result of correlating task-associated fMRI maps with gradients 1-5. High scores denote high correlations. There is notably little differentiation on gradient 1 compared with the other gradients. Standard deviations are not shown due to the use of single maps for each task.

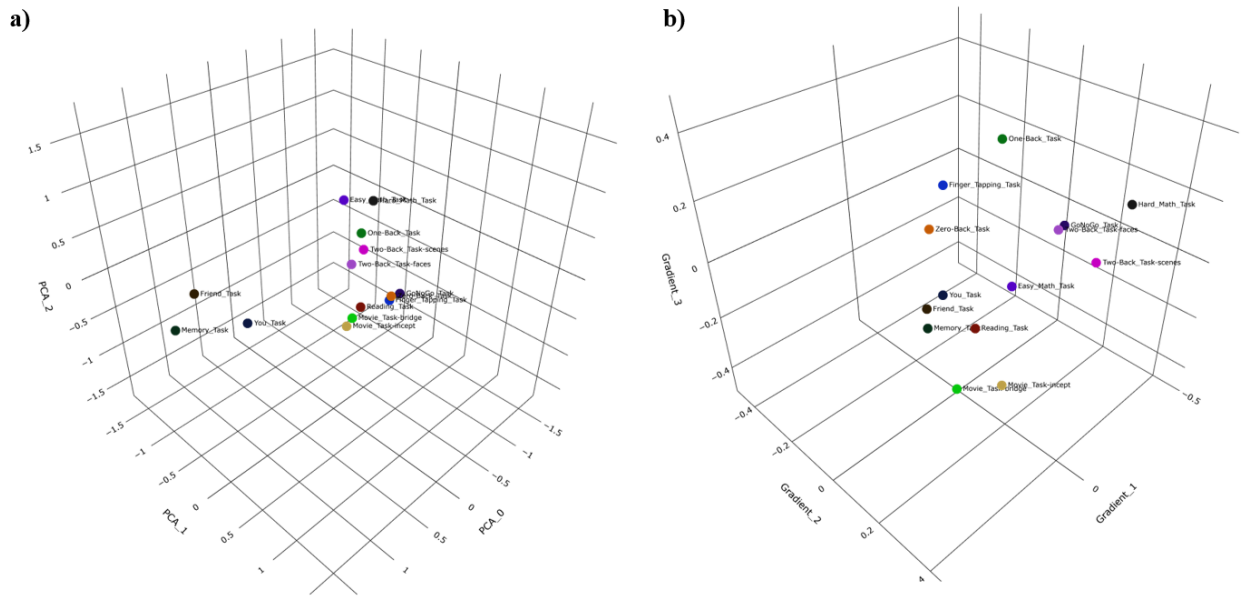

Figure 9: Tasks in common spaces. a) The PCA loadings across all blocks and participants of each task were averaged, and the average task component loadings were used to determine coordinates of tasks in this space. The axes represent the mean PCA scores across all participants in a task condition. b) Brain maps associated to task states were correlated to gradients after application of a common background mask calculated on all maps. Tasks are placed into the space based on the Spearman correlation to each of the gradients individually. Only gradients 1, 2 and 3 are shown for readability. The axes represent the correlation between a task-associated brain map and a gradient brain map.

4.2 Determining the Mapping Between Spaces

Having created a framework for comparing mDES and brain data, we evaluated whether the structure of one space can be used to predict another. This required an assessment of out-of-sample prediction error. To assess our out-of-sample prediction error, we used leave-one-out cross validation. This method would allow us to estimate how well this space would be able to predict a novel task not included in this experiment. On each iteration, we removed all the responses from one task and fit a linear regression on the remaining data. We then predicted the position of the left-out task for each probe and measured the average absolute error (Euclidian distance from true to predicted value) compared to the true location of the task. This was repeated for each task. The error measurements were then averaged, giving us the mean average error (MAE). This cross-validated approach allows us to emulate predicting on an unseen task. To control for possible data leakage (in which information about the left-out task is "leaked" to

the rest of the data), the PCA was recalculated on each fold, preventing variation from the unseen task to be represented in the PCA scores of the fitted data. An additional complication is that there are multiple mDES probes for each fMRI map, which can cause predictions of PCA components to have higher error scores.

To determine if this MAE is better than chance, we performed a permutation test. This test consisted of repeating the analysis 1000 times, but randomly shuffling the data each time so that probe data remained intact but was no longer associated with the task it was derived from. This broke the association between mDES probes and state space coordinates and generated a null distribution of MAE values. We then checked the value of our original MAE against the null distribution of MAE values. If our original was better than our null distribution, then we can predict better than chance. To illustrate this, the null distribution of MAE values resulting from the permutation test is shown, with the real accuracy value represented by the dotted red line (Fig. 10)

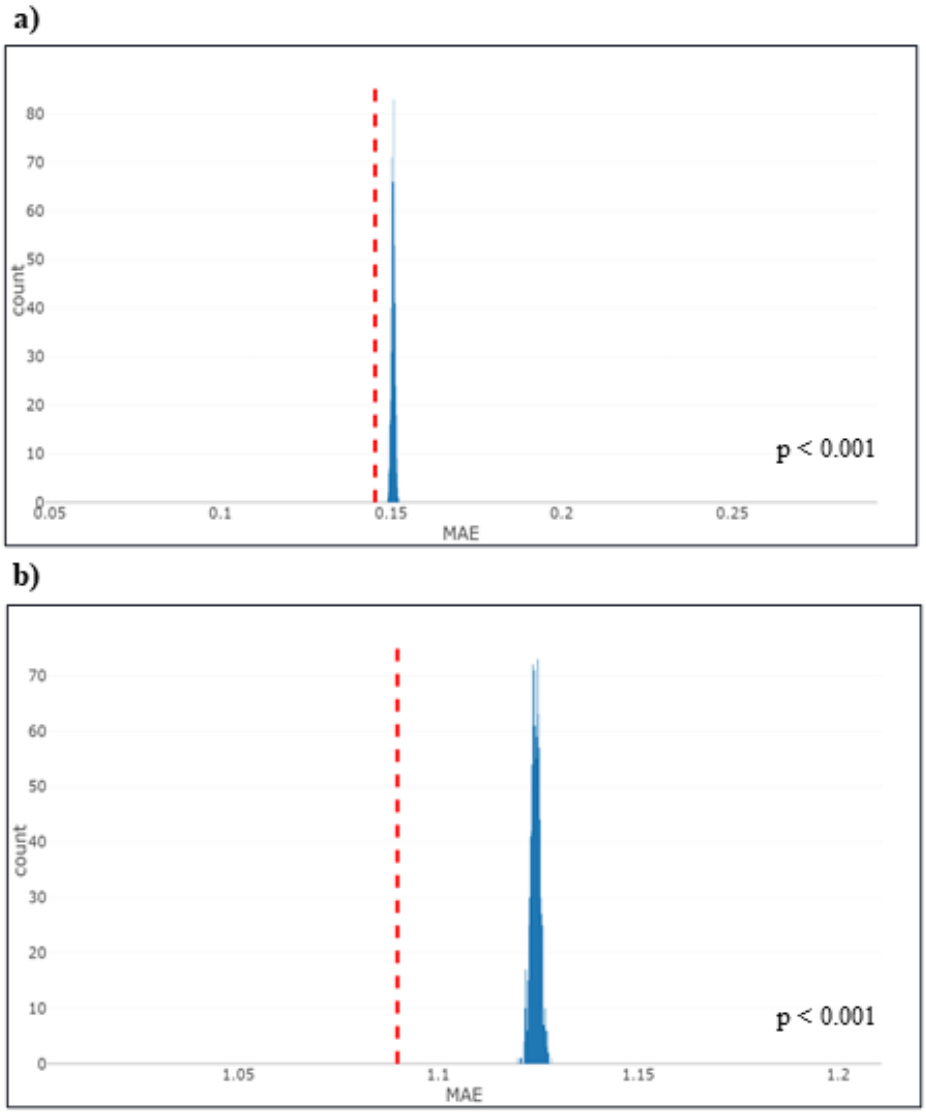

Figure 10. Results of permutation test using linear regression on gradient and thought component prediction. a) Coordinates on all gradients were predicted via a linear regression, with mean absolute error (MAE) metrics collected by cross validating across each task. The red line indicates the accuracy value of our non-permuted data. b) PCA scores on all components of mDES data were predicted via a linear regression, with MAE metrics collected by cross validating across each task. The red line indicates the accuracy value of our non-permuted data.

4.3 Investigating the Characteristics of the Mapping

The results of the permuted linear regression show that the true data has a much smaller error than the null data, implying that we can map across these two spaces better than chance (Fig. 10). This relationship can be evaluated further by breaking the prediction task down into individual gradient coordinates, or individual component scores. As seen in Figure 11, the prediction of gradient coordinates is successful for all gradients other than gradient 1, the

Unimodal-Transmodal gradient. For prediction of component scores, we see that our prediction is successful for all scores other than scores on component 2, the Distracting-Intrusive component.

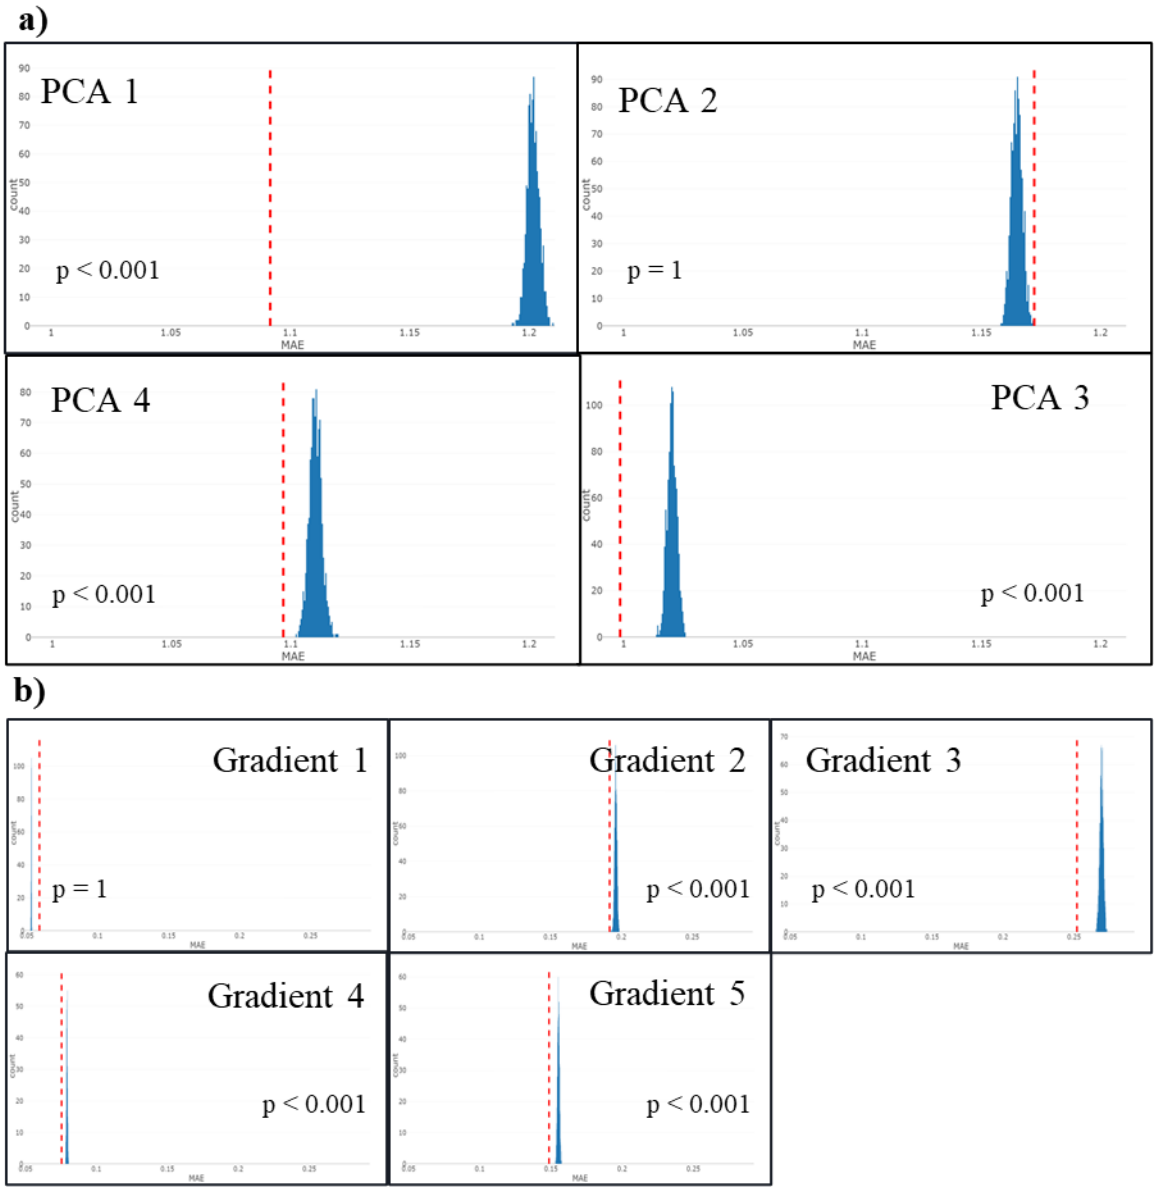

Figure 11: Dissection of charactersitcs of gradient and thought component prediction. a) Thought components 1-4 were predicted via a linear regression, with MAE metrics collected by cross validating across each task. P-value, depicted here as a red dotted line, is the significance of the observed accuracy versus null data (blue distribution). b) Gradients 1-5 were predicted via a linear regression, with MAE metrics collected by cross validating across each task. P-value, depicted here as a red dotted line, is the significance of the observed accuracy versus null data (blue distribution).

In order to understand which brain networks were the most important for these predictions, virtual lesions on brain networks were performed. This is accomplished by removing all voxels within a network, and consequently any information contained within them, from our analysis (Fig. 12). The networks used for this analysis were defined by Yeo (Thomas Yeo, Krienen et al. 2011). An additional masking step was also employed, which removed any regions missing from any task-associated brain map from all brain maps. The resulting mask removed a small dorsal section of the superior frontal gyrus.

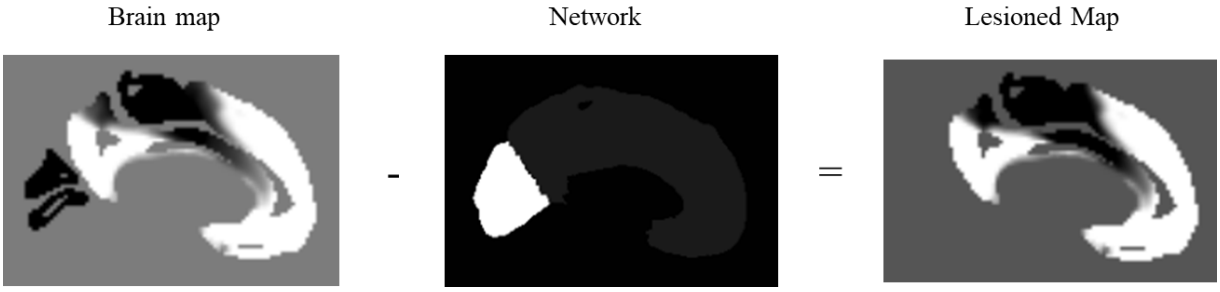

Figure 12: Network lesioning method. Illustrated here is the process for creating lesioned brain maps, showing the removal of a network from a fMRI map, and the resulting lesioned fMRI map.

We first explored the effect of lesioning on our prediction of all gradient coordinates (Fig. 13), and on all component scores (Fig. 14). Lesioning had no effect on any predictions with the notable exception of the lesioning of the Default Mode Network breaking our PCA score predictions. We also evaluated the robustness of this prediction under both lesioned and unlesioned conditions by bootstrapping the test (n=5000, subsample ratio=0.2), the results of which are seen in Figure 15.

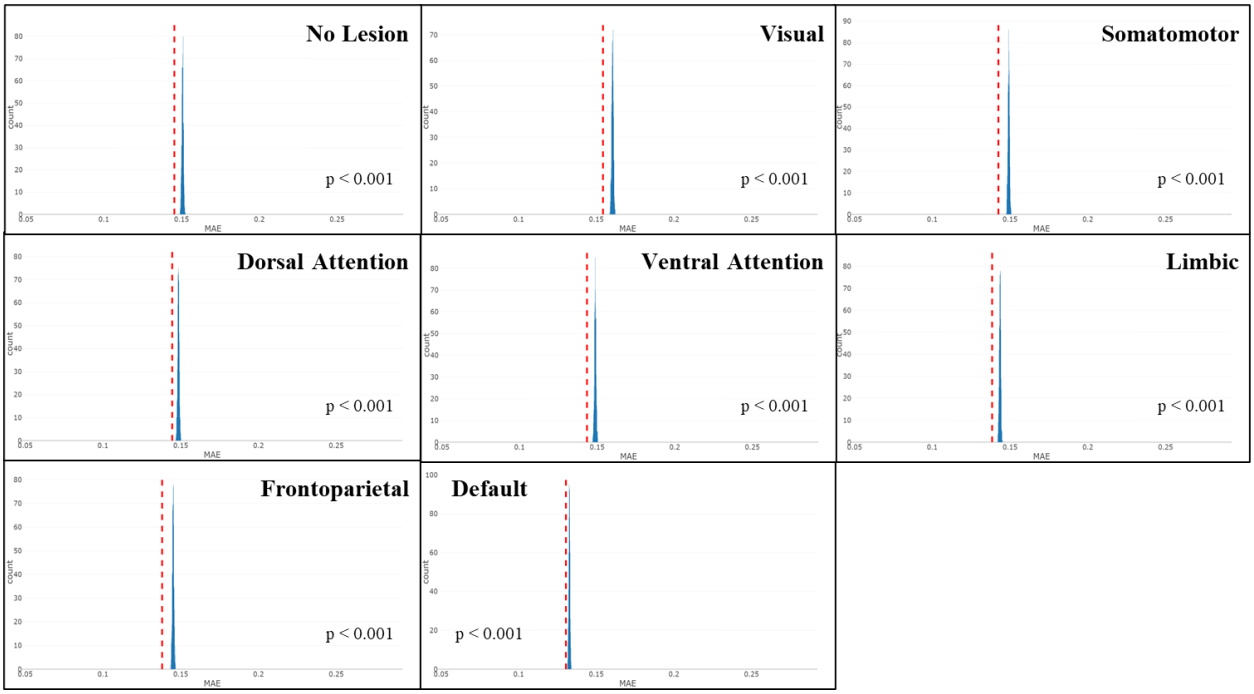

Figure 13: Results of permutation test using linear regression to predict gradient coordinates on all lesions. The whole pipeline was run on lesioned versions of the input maps, lesioned according to the Yeo networks. The linear regression predicted gradient coordinates based on mDES data decomposed with PCA. P-value, depicted here as a red dotted line, is the significance of the observed accuracy versus null data (blue distribution).

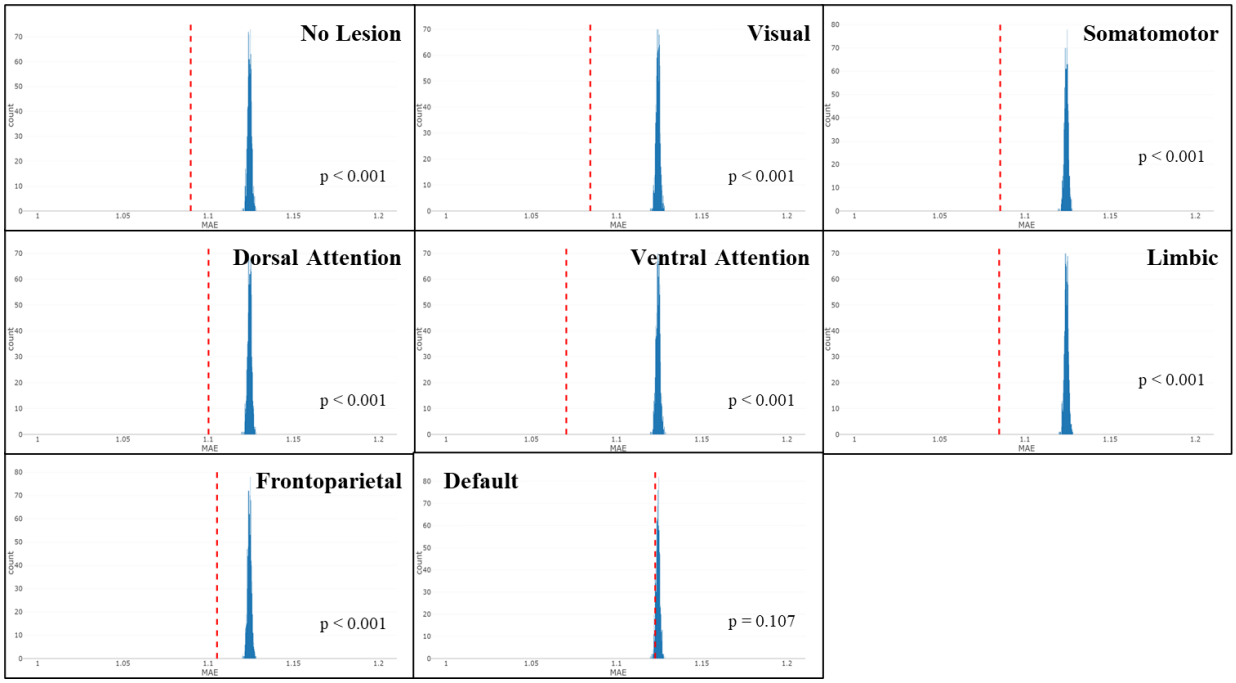

Figure 14: Results of permutation test using linear regression to predict PCA scores on all lesions. The whole pipeline was run on lesioned versions of the input maps, lesioned according to the Yeo networks. The linear regression was predicting mDES scores decomposed with PCA using the gradient coordinates of the task-associated brain maps. P-value, depicted here as a red dotted line, is the significance of the observed accuracy versus null data (blue distribution).

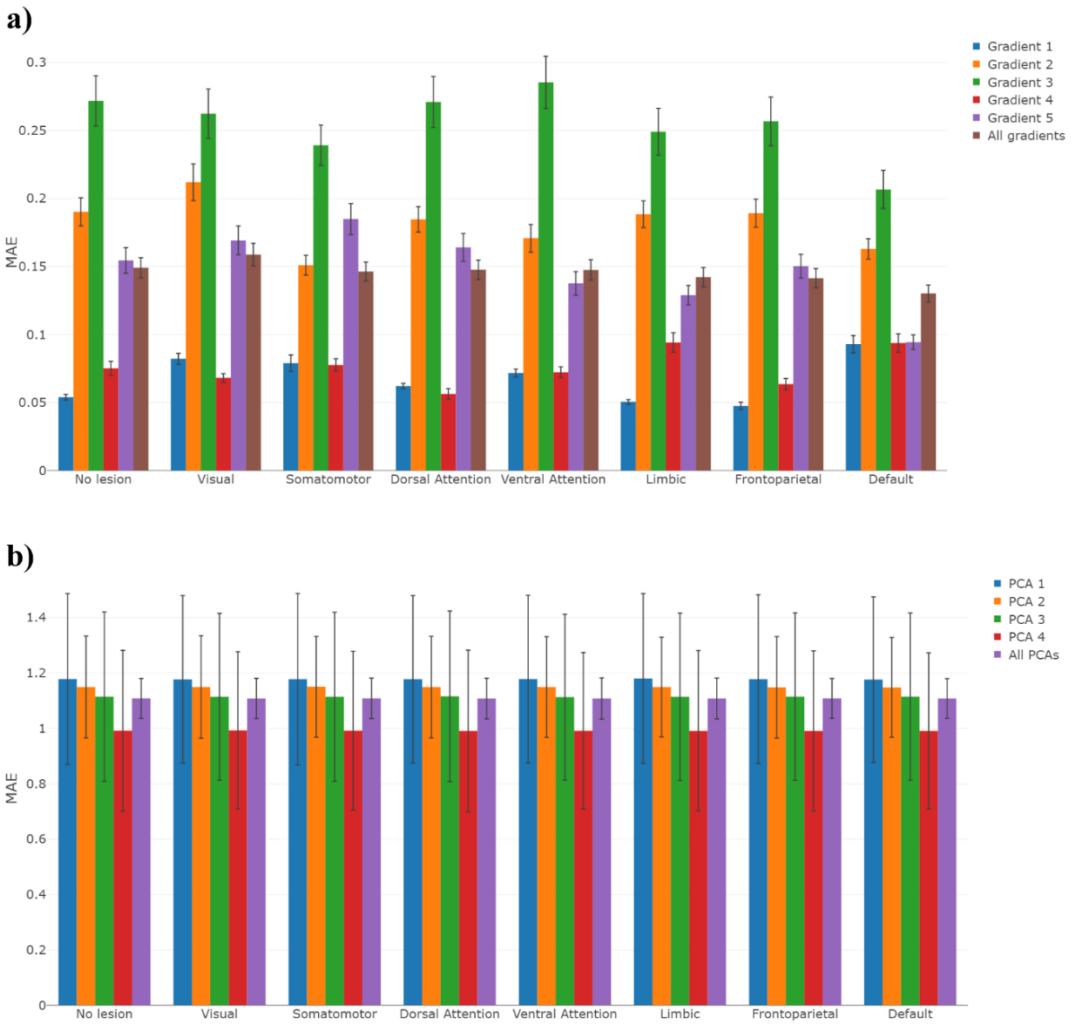

Figure 15: Bootstrapped MAE values from model fitting. a) MAE scores from a linear regression model are shown for our prediction of all gradient positions on multiple lesions. Colours are mapped to gradients as follows: blue represents gradient 1, orange represents gradient 2, green is gradient 3, red is gradient 4, purple is gradient 5, and brown is all combined gradients. 99.993% confidence intervals (Bonferroni corrected) are included. b) MAE scores from a linear regression model are shown for our prediction of all PCA scores on multiple lesions. Colours are mapped to PCA components as follows: blue represents PCA 1 (Past-Knowledge), orange represents PCA 2 (Distracting-Intrusive), green represents PCA 3 (Deliberate-Problem Solving), red represents PCA 4 (Sensory engagement), and purple represents all combined PCA components. 99.993% confidence intervals (Bonferroni corrected) are included.

The bootstrapped accuracy scores show that the prediction is generally resistant to outliers, with the individual components/gradients having similar confidence intervals (Fig. 15). With the relative robustness of the relationship established, we can now investigate the effects of lesioning on our predictions of individual gradient coordinates and PCA scores. The implication

of a lesioned network breaking our prediction is that the information contained within the network was essential for our prediction to succeed, or that the network is driving the prediction.

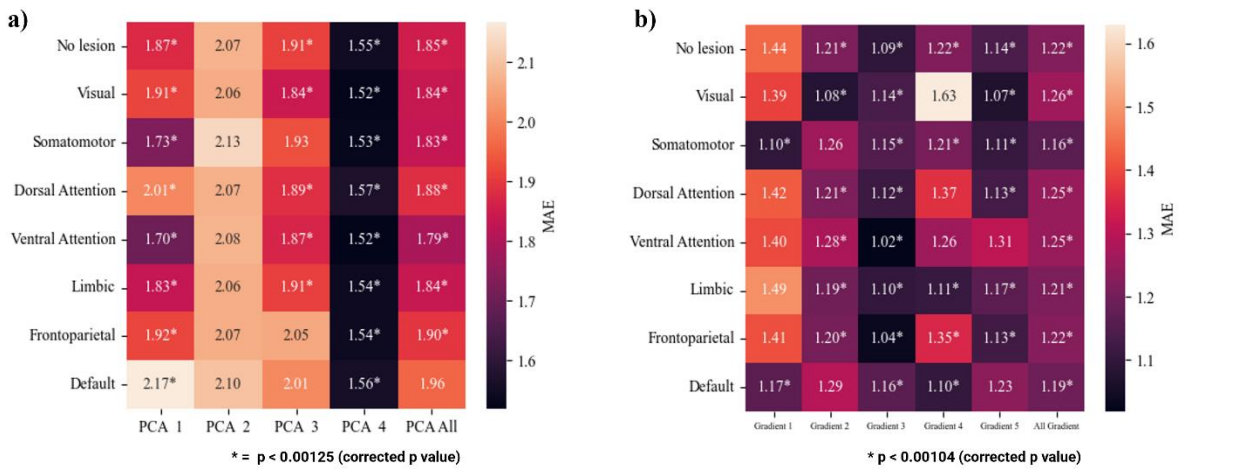

Figure 16: Results of network lesioning on permutation test. Resulting MAE scores from network lesioned permutation tests are shown. Significant results post Bonferroni correction are indicated with an asterisk, with the Bonferroni corrected p-values being 0.00125 and 0.00104 for a and b, respectively. a) Results are shown for permutation testing along 5 gradients, as well as the combined result of predicting all 5 gradient coordinates at once. b) Results are shown for permutation testing along 4 thought PCAs, as well as the combined result of predicting all PCA scores at once.

The results from the lesioned analysis revealed interesting patterns in the predictability of different components and gradient coordinates. In the PCA prediction analysis, component 1 (Past Knowledge) and component 4 (Sensory Engagement) consistently showed predictability regardless of lesions, whereas component 2 (Distracting-Intrusive) was never predictable. Interestingly, component 3 (Deliberate-Problem Solving) was predictable for all lesions except when the frontoparietal network, the somatomotor network, or the default mode network were lesioned, which resulted in a decrease in accuracy of 0.14 MAE, 0.02 MAE and 0.1 MAE, respectively, compared to the unlesioned analysis. Component 4 (Sensory-Engagement) also had notably better accuracy, likely due to easier differentiation of the component within movie/memory tasks. Furthermore, all combined components remained predictable across all

lesions except when the default mode network was lesioned, leading to a decrease in accuracy by 0.11 MAE compared to the unlesioned analysis (Fig. 16a).

In contrast, the results for gradient coordinate prediction were more intricate. Lesions had no effect on the prediction of all combined gradients. Notably, the prediction of Gradient 1 was impossible without any lesions; however, it became predictable when either the somatomotor or default mode network was lesioned. Remarkably, both lesions exhibited a substantial increase in accuracy compared to the unlesioned version, with a decrease of 0.32 MAE and 0.27 MAE respectively. Conversely, lesioning the somatomotor and default mode network disrupted the prediction of Gradient 2 coordinates, accompanied by a decrease in accuracy of 0.05 MAE and 0.08 MAE respectively, compared to the unlesioned analysis. The lesioning did not affect the prediction of Gradient 3 coordinates. However, lesioning the visual, dorsal attention, and ventral attention networks disrupted the prediction of Gradient 4 coordinates, resulting in accuracy decreases of 0.41 MAE, 0.25 MAE, and 0.04 MAE respectively, compared to the unlesioned analysis. Finally, the prediction of Gradient 5 coordinates was hindered by lesioning the ventral attention and default mode networks, leading to accuracy decreases of 0.16 MAE and 0.09 MAE, respectively (Fig. 16b).

An important caveat to the interpretation of these results is that the significance of each prediction is measured as the average accuracy of predictions made on left-out tasks versus the average accuracy of a null distribution generated by attempting many predictions on shuffled data. We cannot say that lesioning a network significantly decreases our accuracy compared to our unlesioned brain, we can only claim that we can no longer predict our mDES PCAs when a network is removed.

Chapter 5: Discussion

5.1 Summary of Results

The main objective of my Master's was to investigate whether contextual variation in the task being performed can help understand the relationship between ongoing thought and patterns of brain activity. Over the course of the study, we collected mDES responses from participants while in various task states (Table 1, Fig. 6a). We decomposed the responses to form 4 stable principal components of ongoing thought (Fig. 6, Fig. 8). This allowed us to form a space where we could plot the similarities and differences of task-related ongoing thought. We also created a similar space by correlating our task-associated brain maps with macroscale gradients of activity (Fig. 9). We then determined that there was an association between the task-associated ongoing thought space and the brain state space using a permutation test (Fig. 10). We then investigated this relationship further by examining the influence of each principal component and macroscale gradient on the predictability of the association (Fig. 11). We also virtually lesioned brain networks to determine which brain regions are most important for understanding the mapping between brain activity and patterns of ongoing thought (Fig. 13, Fig. 14, Fig. 16).

5.2 Experience Sampling

Our original objectives for the experience sampling analysis were to find components of ongoing thought consistent with previous experiments using mDES decomposition, and to determine if they could be associated with brain activity. The components we found were generally consistent with previous decompositions of mDES data (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013, Engert, Smallwood et al. 2014, Ho, Poerio et al. 2020, Konu, Mckeown et al. 2021, Mckeown, Poerio et al. 2021, Mulholland, Goodall-Halliwell et al. 2022). Broad examples include the replication of a Past-Knowledge component (Ruby, Smallwood et al.

2013, Ruby, Smallwood et al. 2013, Konu, Mckeown et al. 2021, Mulholland, Goodall-Halliwell et al. 2022), a Detailed-Deliberate component (Ho, Poerio et al. 2020, Konu, Mckeown et al. 2021, Mulholland, Goodall-Halliwell et al. 2022), a Distracting-Intrusive component (Engert, Smallwood et al. 2014, Konu, Mckeown et al. 2021, Mulholland, Goodall-Halliwell et al. 2022), and a Sensory/Modality component (Ho, Poerio et al. 2020, Mckeown, Poerio et al. 2021). Using a permuted linear regression, we found that we could predict all combined component scores based on spatial patterns of brain activity. Furthermore, examining the prediction of components individually, we could predict component scores on all components except for component 2, which is the Intrusive-Distracting component (Fig. 11). The bootstrapped accuracy scores (Fig. 15) also displayed relative stability within the error scores, indicating that the predictive relationship between our component scores and state space coordinates is robust to outliers in our sample. Finally, our network lesioning analysis found that the frontoparietal and default mode network contained necessary information for our prediction of component 3, our Deliberate-Problem Solving component (Fig. 16a). Additionally, lesioning of the default mode network disrupted the prediction of all combined component scores (Fig. 16a). Regardless of lesioning, our accuracy was best on predictions of component 4, the Sensory Engagement component (Fig. 16a).

5.3 Brain Activity

Using a permuted linear regression, we found that we could predict all combined gradient space coordinates, establishing a two-way predictive relationship between our mDES component scores and task-associated brain maps (Fig. 10). Examination of the individual gradients' predictability showed that we were able to predict position on all gradients apart from gradient 1, the Unimodal-Transmodal gradient (Fig. 13). Lesioning both the somatomotor and default mode

network made gradient 1 predictable, while also increasing the accuracy (Fig. 16b). The inverse effect is seen in gradient 2, the visual-sensorimotor gradient, as it is predictable while unlesioned, but the prediction breaks when either the somatomotor or default mode network is removed (Fig. 16b). Gradient 3, the task positive-task negative gradient, seemed to be predictable regardless of lesions but had some of the best accuracy values when the ventral attention and frontoparietal networks were lesioned (Fig. 16b). Prediction of gradient 4, the gradient that separates DMN regions, was broken when the visual, ventral attention, and dorsal attention networks were removed, with a notably large decrease in accuracy when the visual network was removed (Fig. 16b). Gradient 5, which separates the limbic and dorsal attention networks, could not be predicted after the removal of the ventral attention and default mode networks (Fig. 16b). Finally, the prediction of all combined gradient positions was successful regardless of lesion (Fig. 16b).

5.4 How These Results Fit with the Literature

Our findings not only contribute to our understanding of the association between thoughts and brain activity but also align with existing literature in this area. The four components identified in our study exhibit similarities to those found in other studies (Mulholland, Goodall-Halliwell et al. 2022) which found similar component structures for Past-Knowledge, Distracting-Intrusive, and Deliberate-Problem Solving, with notable differences within our last component, Sensory Engagement. Our findings are also similar to a study done by Konu (Konu, Mckeown et al. 2021), which used many of the same tasks as this thesis and had similar averaged task component loadings (Fig. 17). Although our component structure is slightly different, particularly with respect to Konu's component of "Self focus", the components of "Episodic social cognition", "Unpleasant intrusive", and "Concentration" mirror our components of "Past

Knowledge", "Intrusive Distracting", and "Deliberate-Problem Solving". A specific example which highlights the generalizability of mDES decomposition comes from the close mapping of Konu's "Concentration" component and our "Deliberate-Problem Solving" component. From our results (Fig. 8), we see that the "Deliberate-Problem Solving" component has low scores in the two Movie tasks, and higher scores in the Math tasks and One-back task. Our Math and Movie tasks were not included in Konu's study, but the One-back task is a working memory task, so our findings echo those of Konu. There was also a degree of similarity between the Konu's component of "Unpleasant intrusive" and our component of "Intrusive distracting", with both loading highly on negative emotion, future-focused thought and intrusive thinking. From other work, we see across many different iterations of mDES questionnaires, thinking about the past is strongly associated with thoughts about other people (Ruby, Smallwood et al. 2013, Ruby, Smallwood et al. 2013, Engert, Smallwood et al. 2014, Ho, Poerio et al. 2020, Mckeown, Poerio et al. 2021). Among later iterations of the mDES questionnaire which included questions about whether thoughts were intrusive, detailed, or deliberate, we see two well-conserved components of detailed task-focused thought and intrusive-distracting thought (Konu, Mckeown et al. 2021, Mulholland, Goodall-Halliwell et al. 2022).

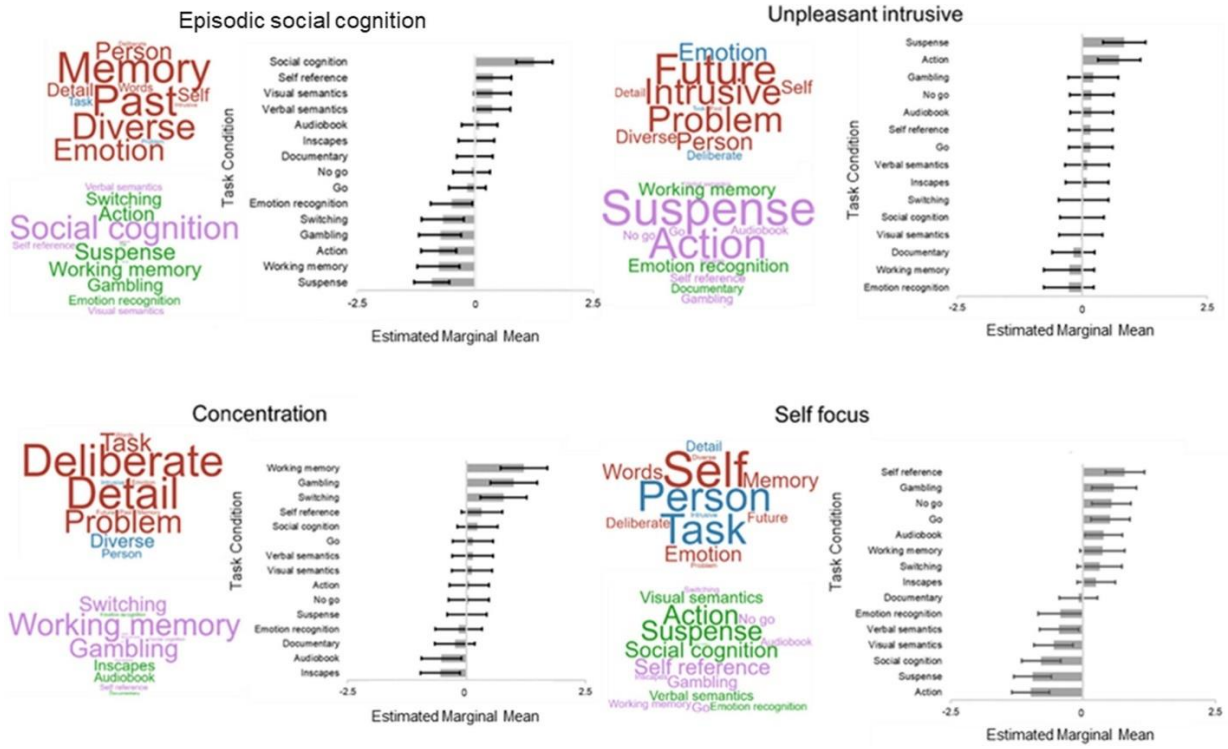

Figure 17: Results of a linear mixed model (LMM) examining the variance across task environments from previous work (Konu, Mckeown et al. 2021). In each panel, the top wordcloud reiterates the loadings on each thought pattern. The colour of the word describes the direction of the relationship; red = positive, blue = negative, and the size reflects the magnitude of the loading. The lower wordcloud highlights the loadings of this pattern in each task as described by the parameter estimates from the LMM. The colour of the word describes the direction of the relationship; purple = positive, green = negative, and the size reflects the magnitude of the loading. The bar plot shows the same data and reports the confidence intervals for these estimates ($p < .05$, corrected for family-wise error). Error bars, therefore, represent 99.7% CI and so control for family-wise error in these analyses.

Furthermore, our study's observation of the differentiation of brain maps across states, as demonstrated in Figure 9, aligns with other gradient-based analyses (Ito, Hearne et al. 2020, Turnbull, Karapanagiotidis et al. 2020, Wang, Margulies et al. 2020, Konu, Mckeown et al. 2021). Notably, Turnbull's differentiation of on and off-task states along gradient 3 (Turnbull, Wang et al. 2019), the task-positive to task negative gradient, aligned with the differentiation of our tasks on this gradient. The positive end of this gradient, which was on-task in Turnbull's study, was populated with the Hard Math, One-back, and Two-back tasks. The negative end of this gradient, which was off task in Turnbull's study, was populated by our Movie, Autobiographical Memory, and Reading tasks. These task-based differentiations also mirror the task's loadings on our Components 3 and 4, which are for deliberate problem-solving and

sensory engagement respectively. We see that positive deliberate problem solving is associated with the positive end of the gradient, and that sensory engagement is associated with the negative end. The study performed by Turnbull (Turnbull, Wang et al. 2019) was able to find some differentiating effects along gradient 1, which may be due to the experiments' partial focus on brain dynamics. When developing the gradients, Margulies primarily focused on the differentiation across Gradient 1 (Margulies, Ghosh et al. 2016), which proved challenging for us to predict in laboratory settings. This difficulty may arise from Gradient 1's differentiation of higher-order social cognition from simpler motor/visual tasks, the former being difficult to sample in a lab environment (Ito, Hearne et al. 2020, Viviani, Dommes et al. 2020, Bernhardt, Smallwood et al. 2022). Gradient 1 may also have more to do with the organization of brain structure, instead of differentiating function (Huntenburg, Bazin et al. 2018). Given our lack of differentiation along the first gradient, but successful differentiation along the other gradients, it may be the case that our approach was better suited to make predictions along functional gradients, instead of structural ones. While our study lacks fMRI time series data for participants, comparing states on the level of region-specific functional changes represents an intriguing avenue for future research, albeit requiring scanning participants during the task battery.

5.5 How These Results Fit with the Bigger Picture

Our work contributes to the broader understanding of the links between two measures of cognition, self-reports of ongoing thought patterns and brain activity, by adopting a predictive approach. This predictive nature sets our analysis apart from more common correlational approaches, as it estimates the out-of-sample generalization of the relationship (Poldrack, Huckins et al. 2020). The key differences come from the use of cross-validation and a measure of prediction error (Fushiki 2011). This is in contrast to descriptive work which links dimensions of

experience sampling to momentary brain states such as the correlation between reports of off-task thought and brain activity measured using fMRI (Turnbull, Wang et al. 2019). This enables our experiment to draw conclusions about the patterns of mDES data which represent off-task thought, but the alternative correlational approach can not test for how well this relationship generalizes to other tasks than the n-back task used in their experiment. By using a permuted linear regression, we can determine if we are able to predict on held-out tasks within our data better than chance (Anderson and Robinson 2001). The comparison of our prediction error on true data versus several shuffled sets of the determines our significance, where a low value indicates that there is a relationship we can predict. By demonstrating the stability and generalizability of the predictive relationship between subjective experiences and brain activity, our study provides a valuable method for testing the correspondence between subjective reports and objective measures of brain function. Moreover, the reliance on brain maps obtained and tested in different labs, with slight contextual variations, further strengthens the generalizability of our findings (Cao, McEwen et al. 2019). Despite the simplicity of our study design which employed a relatively straightforward task battery and analysis pipeline, we observed a significant predictive relationship (Fig. 10). This effect can be elaborated through additional investigations into the nature and robustness of the state space, as well as the generalizability of the predictive relationship between thoughts and brain activity.

A common theme among other studies that have decomposed mDES responses is the lack of presence of the modality question (“Were your thoughts in the form of images or words.") within components (Konu, Mckeown et al. 2021), or component structure that is strongly defined by the modality item (Konu, Turnbull et al. 2020). Although the questions used in our study were very similar to those in previous studies, the decomposition of the modality item had a large

impact on the component structure. Notably, the splitting of the modality item into three separate questions assessing whether thoughts were in the form of images, words, and sounds influenced the structure of our Component 4. This change was made following a study which showed that people high on an autism quotient were easily differentiated by the old modality question which asked if thoughts were in either images or words (Turnbull, Garfinkel et al. 2020). Breaking this question down into separate questions which ask the degree to which thoughts were in the form of 3 different senses (Images, words, sounds; Table 1) was done to enable people to more richly describe their ongoing thoughts, and to see if we can still see the patterns of component structure seen in previous studies.

In considering future directions for this study, several promising avenues emerge. We could first examine which of our tasks drove this effect the best. This would involve removing data from individual tasks to see which tasks, when removed, broke the prediction. This information would be valuable to anyone seeking to design a task battery with reduced length, which can still broadly sample cognition. The tasks which drive the effect are likely to be those with the most extreme values along both the gradients and the components, but an analysis to confirm this would be beneficial. Another interesting future direction would be to compare mDES scores to more traditional accuracy metrics from the tasks. In this battery we recorded task accuracy metrics where possible, but some tasks had no viable accuracy metrics. It should be possible to determine how individual accuracy metrics change the position of an individual's mDES scores in our space, and possibly create a fair exclusion criterion for low-effort subjects.

5.6 Limitations

While our exploratory analytic strategy has provided valuable insights into the comparison between brain activity and subjective reports, there are several limitations that

should be acknowledged. First, to enhance our strategy, we could consider increasing the number of tasks used in the study. By including tasks that exhibit a greater differentiation of states in the state space, we could offer a more comprehensive understanding of the relationship between brain activity and subjective experiences. This expanded task repertoire would enable us to capture a wider range of cognitive and emotional processes, providing a more nuanced exploration of the underlying neural mechanisms.

Second, conducting these tasks within a scanner and formally assessing the predicted answers would further refine our analysis. Utilizing neuroimaging techniques alongside the subjective reports would allow for a more direct assessment of the correspondence between brain states and subjective experiences. By obtaining objective measurements of brain activity in conjunction with self-reported data, we could validate the reliability and accuracy of the reported subjective states. The approach used in this experiment can be used to determine which tasks are best suited to discriminating experience.

Another limitation lies in the choice of our gradients, or for that matter, the choice of any standard brain map with which we created our space. In this paper, we only explored an approach using the first 5 gradients extracted by Margulies. However, further work could involve exploring the use of standard maps derived more strongly from structure, as proposed in a recent study (Valk, Xu et al. 2022). These maps could differentiate anatomical features such as density of certain receptors or maps representing how well regions are conserved across primates. Alternatively, we could investigate decomposing patterns of structure using PCA to create a more direct representation of brain states (Coifman and Lafon 2006, Tang, Chen et al. 2021, Pang, Aquino et al. 2023). By incorporating different mapping strategies, we can assess the

robustness and generalizability of our findings and ensure that they are not solely reliant on a single representation of brain activity.

In addition, our sample was not representative of the general population. We recruited participants from a university population and mostly consisted of people who identified as female which may limit the generalizability of our findings to broader demographics. To address this limitation, it would be valuable to expand the subject pool to include more diverse participants from various backgrounds and age groups and more representative gender identities. By incorporating a more representative sample, we can enhance the external validity of our results and gain insights into potential variations in brain activity and subjective experiences across different populations. Additionally, sampling participants in real-world settings using phone-based tasks could offer a more ecologically valid perspective on the relationship between brain states and subjective experiences. Conducting tasks outside of controlled laboratory environments would provide a better understanding of how brain activity and subjective reports manifest in naturalistic contexts. This approach would bridge the gap between laboratory findings and real-world experiences, ensuring the relevance and applicability of our research to everyday life situations.

Finally, the decision to select 4 components from our PCA was driven by their eigenvalues being over 1, which is a common cutoff, but not always ideal (Cliff 1988). The inclusion of more components in this analysis may increase accuracy while decreasing generalizability since the additional components will be representing smaller trends in our data. When added to this analytic pipeline, the components derived from smaller trends will be treated as equal to the other components, likely causing overfitting and decreasing our final accuracy. Once enough data has been collected with the current mDES questionnaire, a more

comprehensive analysis could be conducted to determine an optimal component structure. This would allow for component structure to be kept consistent across experimental designs and enable better comparisons between experimental results.

References

Alam, T. G., et al. (2018). "Meaningful inhibition: Exploring the role of meaning and modality in response inhibition." Neuroimage **181**: 108-119.

Anderson, M. J. and J. Robinson (2001). "Permutation tests for linear models." Australian & New Zealand Journal of Statistics **43**(1): 75-88.

Anderson, N. H. (1968). "Likableness ratings of 555 personality-trait words." Journal of personality and social psychology **9**(3): 272.

Andreasen, N. C., et al. (1995). "Remembering the past: two facets of episodic memory explored with positron emission tomography." The American journal of psychiatry **152**(11): 1576-1585.

Bansal, K., et al. (2018). "Data-driven brain network models differentiate variability across language tasks." PLoS computational biology **14**(10): e1006487.

Barch, D. M., et al. (2013). "Function in the human connectome: task-fMRI and individual differences in behavior." Neuroimage **80**: 169-189.

Barrett, L. F. and D. J. Barrett (2001). "An introduction to computerized experience sampling in psychology." Social science computer review **19**(2): 175-185.

Belkin, M. and P. Niyogi (2003). "Laplacian eigenmaps for dimensionality reduction and data representation." Neural computation **15**(6): 1373-1396.

Bernhardt, B. C., et al. (2022). Gradients in brain organization, Elsevier. **251:** 118987.

Binder, J. R., et al. (1999). "Conceptual processing during the conscious resting state: a functional MRI study." Journal of cognitive neuroscience **11**(1): 80-93.

Buckner, R. L. and M. E. Wheeler (2001). "The cognitive neuroscience og remembering." Nature Reviews Neuroscience **2**(9): 624-634.

Camilleri, J. A., et al. (2018). "Definition and characterization of an extended multiple-demand network." Neuroimage **165**: 138-147.

Cao, H., et al. (2019). "Toward leveraging human connectomic data in large consortia: generalizability of fMRI-based brain graphs across sites, sessions, and paradigms." Cerebral Cortex **29**(3): 1263-1279.

Christoff, K., et al. (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.

Cliff, N. (1988). "The eigenvalues-greater-than-one rule and the reliability of components." Psychological bulletin **103**(2): 276.

Cohen, J. D., et al. (2017). "Computational approaches to fMRI analysis." Nature neuroscience **20**(3): 304-313.

Coifman, R. R. and S. Lafon (2006). "Diffusion maps." Applied and computational harmonic analysis **21**(1): 5-30.

Cole, M. W., et al. (2013). "Multi-task connectivity reveals flexible hubs for adaptive task control." Nature neuroscience **16**(9): 1348-1355.

Craik, F. I., et al. (1999). "In search of the self: A positron emission tomography study." Psychological science **10**(1): 26-34.

de Caso, I., et al. (2017). "Knowing me, knowing you: Resting-state functional connectivity of ventromedial prefrontal cortex dissociates memory related to self from a familiar other." Brain and Cognition **113**: 65-75.

de Caso, I., et al. (2017). "That's me in the spotlight: neural basis of individual differences in self-consciousness." Social cognitive and affective neuroscience **12**(9): 1384-1393.

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.

Engert, V., et al. (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.

Fox, M. D., et al. (2005). "The human brain is intrinsically organized into dynamic, anticorrelated functional networks." Proceedings of the National Academy of Sciences **102**(27): 9673-9678.

Franklin, M. S., et al. (2016). "Unaware yet reliant on attention: Experience sampling reveals that mind-wandering impedes implicit learning." Psychonomic Bulletin & Review **23**: 223-229.

Fushiki, T. (2011). "Estimation of prediction error by using K-fold cross-validation." Statistics and Computing **21**: 137-146.

Ho, N. S. P., et al. (2020). "Facing up to the wandering mind: Patterns of off-task laboratory thought are associated with stronger neural recruitment of right fusiform cortex while processing facial stimuli." Neuroimage **214**: 116765.

Hong, S.-J., et al. (2020). "Toward a connectivity gradient-based framework for reproducible biomarker discovery." Neuroimage **223**: 117322.

Huntenburg, J. M., et al. (2018). "Large-scale gradients in human cortical organization." Trends in cognitive sciences **22**(1): 21-31.

Ito, T., et al. (2020). "A cortical hierarchy of localized and distributed processes revealed via dissociation of task activations, connectivity changes, and intrinsic timescales." Neuroimage **221**: 117141.

Jung, K., et al. (2021). "Response of multiple demand network to visual search demands." Neuroimage **229**: 117755.

Karapanagiotidis, T., et al. (2020). "The psychological correlates of distinct neural states occurring during wakeful rest." Scientific reports **10**(1): 1-11.

Kelley, W. M., et al. (2002). "Finding the self? An event-related fMRI study." Journal of cognitive neuroscience **14**(5): 785-794.

Konu, D., et al. (2021). "Exploring patterns of ongoing thought under naturalistic and conventional task-based conditions." Consciousness and cognition **93**: 103139.

Konu, D., et al. (2020). "A role for the ventromedial prefrontal cortex in self-generated episodic social cognition." Neuroimage **218**: 116977.

Kriegeskorte, N., et al. (2008). "Representational similarity analysis-connecting the branches of systems neuroscience." Frontiers in systems neuroscience: 4.

Lanzoni, L., et al. (2020). "The role of default mode network in semantic cue integration." Neuroimage **219**: 117019.

Livesey, A. C., et al. (2007). "Time perception: manipulation of task difficulty dissociates clock functions from other cognitive demands." Neuropsychologia **45**(2): 321-331.

Margulies, D. S., et al. (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.

Martinon, L., et al. (2022). "Catching thoughts: self-caught experience sampling preferentially captures characteristic features of off-task experiences across the life span." bioRxiv: 2022.2012. 2015.520676.

Martinon, L. M., et al. (2019). "The disentanglement of the neural and experiential complexity of self-generated thoughts: A users guide to combining experience sampling with neuroimaging data." Neuroimage **192**: 15-25.

Mckeown, B., et al. (2021). "The impact of social isolation and changes in work patterns on ongoing thought during the first COVID-19 lockdown in the United Kingdom." Proceedings of the National Academy of Sciences **118**(40): e2102565118.

Mourao-Miranda, J., et al. (2005). "Classifying brain states and determining the discriminating activation patterns: support vector machine on functional MRI data." Neuroimage **28**(4): 980-995.

Mulholland, B., et al. (2022). "Patterns of Ongoing Thought in the Real-World." bioRxiv: 2022.2010. 2005.510994.

Murphy, C., et al. (2019). "Hello, is that me you are looking for? A re-examination of the role of the DMN in social and self relevant aspects of off-task thought." PloS one **14**(11): e0216182.

Nunnally, J. C. (1978). Psychometric Theory: 2d Ed, McGraw-Hill.

Pang, J. C., et al. (2023). "Geometric constraints on human brain function." Nature: 1-9.

Peirce, J., et al. (2019). "PsychoPy2: Experiments in behavior made easy." Behavior research methods **51**: 195-203.

Poldrack, R. A., et al. (2020). "Establishment of best practices for evidence for prediction: a review." JAMA psychiatry **77**(5): 534-540.

Raichle, M. E. (2015). "The brain's default mode network." Annual review of neuroscience **38**: 433-447.

Raichle, M. E., et al. (2001). "A default mode of brain function." Proceedings of the National Academy of Sciences **98**(2): 676-682.

Ruby, F. J., et al. (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.

Ruby, F. J., et al. (2013). "Is self-generated thought a means of social problem solving?" Frontiers in psychology **4**: 962.

Shine, J. M., et al. (2019). "Human cognition involves the dynamic integration of neural activity and neuromodulatory systems." Nature neuroscience **22**(2): 289-296.

Shrestha, N. (2021). "Factor analysis as a tool for survey analysis." American Journal of Applied Mathematics and Statistics **9**(1): 4-11.

Smallwood, J., et al. (2008). "Going AWOL in the brain: Mind wandering reduces cortical analysis of external events." Journal of cognitive neuroscience **20**(3): 458-469.

Smallwood, J. and J. W. Schooler (2015). "The science of mind wandering: Empirically navigating the stream of consciousness." Annual review of psychology **66**: 487-518.

Smallwood, J., et al. (2021). "The neural correlates of ongoing conscious thought." IScience **24**(3).

Smallwood, J., et al. (2021). "The neural correlates of ongoing conscious thought." IScience **24**(3): 102132.

Tang, Y., et al. (2021). "Dimensionality reduction methods for brain imaging data analysis." ACM Computing Surveys (CSUR) **54**(4): 1-36.

Thomas Yeo, B., et al. (2011). "The organization of the human cerebral cortex estimated by intrinsic functional connectivity." Journal of neurophysiology **106**(3): 1125-1165.

Trautwein, F.-M., et al. (2016). "Stimulus-driven reorienting impairs executive control of attention: evidence for a common bottleneck in anterior insula." Cerebral Cortex **26**(11): 4136-4147.

Turnbull, A., et al. (2020). "Word up–Experiential and neurocognitive evidence for associations between autistic symptomology and a preference for thinking in the form of words." Cortex **128**: 88-106.

Turnbull, A., et al. (2020). "Reductions in task positive neural systems occur with the passage of time and are associated with changes in ongoing thought." Scientific reports **10**(1): 9912.

Turnbull, A., et al. (2020). "Reductions in task positive neural systems occur with the passage of time and are associated with changes in ongoing thought." Scientific reports **10**(1): 1-10.

Turnbull, A., et al. (2021). "Age-related changes in ongoing thought relate to external context and individual cognition." Consciousness and cognition **96**: 103226.

Turnbull, A., et al. (2019). "Left dorsolateral prefrontal cortex supports context-dependent prioritisation of off-task thought." Nature Communications **10**(1): 3816.

Valk, S. L., et al. (2022). "Genetic and phylogenetic uncoupling of structure and function in human transmodal cortex." Nature Communications **13**(1): 2341.

Van Essen, D. C., et al. (2013). "The WU-Minn human connectome project: an overview." Neuroimage **80**: 62-79.

Vanderwal, T., et al. (2008). "Self, mother and abstract other: an fMRI study of reflective social processing." Neuroimage **41**(4): 1437-1446.

Viviani, R., et al. (2020). "Segregation, connectivity, and gradients of deactivation in neural correlates of evidence in social decision making." Neuroimage **223**: 117339.

Wang, X., et al. (2020). "A gradient from long-term memory to novel cognition: Transitions through default mode and executive cortex." Neuroimage **220**: 117074.

Wen, T. and T. Egner (2023). "Context-independent scaling of neural responses to task difficulty in the multiple-demand network." Cerebral Cortex **33**(10): 6013-6027.

Woolgar, A. and R. Zopf (2017). "Multisensory coding in the multiple-demand regions: vibrotactile task information is coded in frontoparietal cortex." Journal of neurophysiology **118**(2): 703-716.

Zhang, M., et al. (2022). "Perceptual coupling and decoupling of the default mode network during mind-wandering and reading." Elife **11**: e74011.