Essays on attention and creative thought

by

Jocelyn Yuxing Wang

B.A. Philosophy
Yale University (2018)

Submitted to the Department of Linguistics and Philosophy in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2024

Signature of arthur: ……………………………………………………………………………………..
Department of Linguistics and Philosophy
September 1, 2024

Certified by: ………………………………………………………………………………………….
E.J. Green
Miller Associate Professor, Johns Hopkins University
Thesis Supervisor

Certified by: ………………………………………………………………………………………….
Alex Byrne
Laurance S. Rockefeller Professor
Thesis Supervisor

Accepted by: …………………………………………………………………………………………
Brad Skow
Laurance S. Rockefeller Professor
Chair of the Committee on Graduate students

Essays on attention and creative thought

by Jocelyn Yuxing Wang

Submitted to the Department of Linguistics and Philosophy on August 23, 2023, in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

In the mental life of an ordinary person, creative thoughts, as well as other non-rigid forms of thought, such as mind wandering, are both pervasive and important for our cognitive endeavors. The goal of my dissertation is to provide a theory of these non-rigid forms of thought by understanding some of the cognitive mechanisms that underlie them, as well as to understand how these underlying mechanisms contribute to our epistemic lives more generally in all kinds of reasoning.

Chapter 1 (based on co-authored work with Azenet Lopez) begins with a puzzle that arises from research on mind wandering: since during mind wandering we plausibly prioritize the information relevant to the concurrent tasks less, why does mind wandering sometimes improve rather than impair concurrent task performance? I resolve the puzzle by rejecting the standard conception of attention, according to which the more focused one's attention is, the better it is at improving task performance. I instead argue that certain tasks are better performed with a more diffuse rather than focused mode of attention. I offer a conception of "diffuse attention" that generalizes from external to internal forms of attention and conceptualize mind wandering as an instance of it.

Chapter 2 turns to provide an account of creative thinking, which is closely related to mind wandering. I argue that previous accounts in philosophy about the generation of creative thought are incomplete due to overlooking the role of what I call "memory gists". Memory gists are memory contents that represent more abstract or qualitative features that are extracted from the specific, surface level features in the memory representations that were initially encoded in memory. I argue that generating and using memory gists in memory search enables highly creative people to form connections between memory contents that are not usually associated with each other by revealing their commonalities shared in their gists. Moreover, I argue that different mechanisms underlie online and offline generation of memory gists: the former involves the mode of diffuse attention that I conceptualized in Chapter 1, while the latter involves memory consolidation during sleep or wakeful rests.

The active role that memory plays in creative thinking raises some questions about how to conceptualize the function of memory in our epistemic lives more generally. I explore this topic further in Chapter 3, where I reject the traditional view in epistemology that memory merely functions to preserve previously acquired information, such as information acquired through perception. I argue instead that one of the functions of memory is to improve our understanding of what was represented in the contents that we previously acquired. This is possible thanks to the fact that during memory consolidation, our memory system further processes previously acquired information, and generates representations about relationships between different components of the subject under consideration. My work thus contributes to the ongoing project of understanding memory as an active process instead of merely performing the role of storing information, and highlights understanding as one of the epistemic values that memory generates.

Thesis supervisor: E.J. Green
Title: Miller Associate Professor, Johns Hopkins University

Thesis supervisor: Alex Byrne
Title: Laurance S. Rockefeller Professor

Professor Alex Byrne has certified this thesis on behalf of Professor EJ Green, who has left MIT for his position at Johns Hopkins University. I would like to thank Professor Byrne for acting in this capacity.

Acknowledgments

I am grateful to my primary advisor, EJ Green. EJ has been an extremely generous and responsible mentor since the day we started working together, from the early days of offering to read papers with me to familiarize himself with the topics that I work on, to later providing detailed feedback on countless drafts of my work, and continuing to do so after he moved to Johns Hopkins. I would also like to thank the rest of my dissertation committee: Alex Byrne, Zachary Irving, and Roger White. This dissertation would not have been possible without their guidance and support. I am also grateful to Azenet Lopez, my beloved co-author for one of my dissertation chapters, for being one of my first companions when I started exploring philosophy of mind and philosophy of cognitive science, and for being a great collaborator and a wonderful friend.

Thank you, also, to the members of the MIT Philosophy community for your support throughout my PhD. To Selina Güter: thank you for being a great interlocutor, and for being a great friend in grad school. To Tamar Schapiro: thank you for helping me become a better teacher. I would also like to thank friends that I have made in the broader philosophy community throughout graduate school who have helped me grow intellectually, including Yuan Tian, Malte Hendrickx, Aaron Glasser, Seth Goldwasser, Lu Teng, and many more. And to Sara Aronowitz: thank you for having faith in my projects and for your insight and guidance.

I am also thankful to my colleagues at Minorities and Philosophy (MAP) from my time working there, especially Jingyi Wu and Angela Sun. Working and organizing for MAP was one of the highlights of my graduate school experience. Being a minority in a predominantly white profession is not easy, and I am grateful to have formed many connections through the MAP community. I am also indebted to friends I have made in the Asian community in philosophy, whose support has been crucial to my career in philosophy so far. I would like to thank Hannah Kim and Jen Nguyen, both of whom generously advised me on navigating grad school, job market, and academia more generally. I am also grateful to Yuan Yuan, whose support and encouragement has been crucial to my journey since the day we met at Yale.

Finally, I am very thankful to those in my life throughout my PhD beyond the world of philosophy academia. They say "it takes a village", and I have certainly found my village both within and beyond my profession. Thank you to Enid Cruz Colón for the support and companionship. ¡Gracias por todo! And to my wonderful friends in Boston: Adela Yiyu Zhang, Eugenie Lai, Leah Kihn, Josué Pérez Campos, Nikko Jeffreys, Duc Hoang, Aryan Galani, Keely New, Steffany Villasenor. Thanks to you all, Boston became a second home to me.

And of course, special thanks goes to my parents, and to my feline companion, Kiki: thank you for being the cutest cat one could ask for.

Contents

Acknowledgments 4
Contents 6
Introduction 8

Chapter 1: Mind Wandering as Diffuse Attention 15
1. Introduction 15
2. Diffuse Attention, Internal and External 18
3. Mind Wandering as Internal Diffuse Attention 23
4. Interaction between Mind Wandering and Other Tasks 26
4.1. A Broad Focus of Spatial Attention Improves Temporal Resolution 28
4.2. Interactions between Mind Wandering and Attention in the Primary Task 32
4.3. What Kind of Tasks are Better Performed while Mind Wandering? 35
5. Diffuse Attention and the Prioritization View 37
6. Conclusion 41

Chapter 2: Memory Gist as Mechanism for Creative Insight 42
1. Introduction 42
2. How do creative ideas emerge? A discussion of the philosophical literature 47
2.1. Online generation of creative ideas 48
2.2. Offline generation of creative ideas 52
2.3. Summary 55
3. Online gist generation through internal diffuse attention 55
3.1. Diffuse attention: internal and external 58
4. Offline gist generation through memory consolidation 61
4.1. Gist extraction during memory consolidation and creative insight 62
4.2. Use of memory gists generated offline in memory retrieval 65
5. Discussion 66

Chapter 3: Memory as a generative source of understanding 70
1. Gaining understanding without conscious thinking: a mystery? 70
2. Understanding and the function of memory 74
2.1. The functions of memory 77
3. What is memory consolidation 80
4. How does memory consolidation improve understanding 83
4.1. Extracting general patterns 84
4.2. Integrating piecemeal information 87
4.3. Goal-sensitivity of memory consolidation 91
5. Varieties of epistemic generation in memory 93

5.1. Epistemic generation in memory and the inaccessibility condition 93
5.2. Justification, understanding and generation of different epistemic values 98
6. Conclusion 101

Bibliography **103**
List of figures **116**

Introduction

What does the epistemic life of the ordinary person look like? If you read traditional philosophical discussions of cognition, you might conclude that our thinking is highly rigid: we spend our time applying inference rules, reasoning step-by-step from accepted premises, and reflectively deliberating about and planning our actions. However, much of our epistemic lives are nothing like this. We come up with ideas while daydreaming on the train to work. When we "hit a wall" on difficult problems, we try to get "unstuck" by going on a walk and letting our minds wander. We plan, not step-by-step, but rather by brainstorming through disparate options and actions. Empirical research suggests that these more spontaneous forms of thought occupy up to almost half our waking lives (Kam et al., 2021) and boost creativity (Baird et al. 2012; Irving et al., 2022). How should we understand these non-rigid forms of thinking given their prevalence and significance? How should we understand the epistemic role of the cognitive mechanisms that contribute to non-rigid thoughts? These are the questions that my dissertation addresses.

Attention and mind wandering

A paradigm form of non-rigid thought is mind wandering: the kind of thoughts that are not constrained by a goal, and meander "hither and thither without fixed course or certain aim" (Christoff et al., 2016). When you are in a boring meeting at work for example, your mind might wander from what to get for lunch, to the dishes that you were too busy to wash after dinner last night, to the upcoming visit from your cousin, etc. What mode of attention is involved in mind wandering, and how does it affect the mode of attention employed in concurrent tasks? You might expect that mind wandering harms your performance in concurrent tasks, because you are no longer prioritizing the information relevant for the concurrent tasks as much. For example, as you mind wander in the meeting, you would expect your performance in taking in and processing the content of the meeting to decrease, as you are no longer prioritizing the information delivered by the speakers as much.

Surprisingly, while mind wandering does harm task performance in some cases, it is not *always* detrimental for task-performance (Brosowsky et al. 2021, Baird et al. 2012), in some cases, mind wandering even correlates with performance *improvements*. In Chapter 1, I present one of these cases where memory improves task performance. The task involved is a rapid serial visual presentation task (RSVP), where participants must detect targets within a string of rapidly succeeding distractors. Mind wandering participants do better on these tasks, compared to fully attending participants (Olivers & Nieuwenhuis, 2005; see also Thomson et al., 2015).

I argue that we can explain these puzzling benefits of mind wandering by appealing to a distinctive form of attention employed in mind wandering. According to the standard conception of attention, to pay attention to a certain piece of information is to prioritize this piece of information, and the more prioritized this piece of information is, the better we are in processing it, and the better we are in performing related tasks that require the processing of this information. I will call this view "the prioritization view". In accordance with this view, if during mind wandering one is prioritizing the perceptual input needed for a task at hand, it is only to be expected that performance will suffer, compared to full attention conditions.

Since during mind wandering we plausibly prioritize the information relevant to concurrent tasks less, why does mind wandering sometimes improve rather than impair concurrent task performance? My solution to the puzzle is to reject the prioritization view, with a proposal that enriches the existing taxonomy of different modes of attention with what I call "internal diffuse attention". I argue that certain information is better processed with *diffuse* rather than focused attention, even though diffuse attention does not prioritize information as much as focused attention does (Prettyman 2023, 2014; Gopnik, 2009).

Diffuse attention has traditionally only been discussed in the context of perception. In the last two decades, psychologists have been proposing that at least in perceptual attention, there is a *distributed* mode of attention besides focused attention (Treisman, 2006; Gopnik, 2009, Srinivasan et al., 2009; Demeyere & Humphreys, 2007; Chong & Evans, 2011). This attentional mode has a broader scope, kicks in more quickly, and targets different information than focused attention, even though it sometimes can be used alongside focused attention. Recently in philosophy, Adrienne Prettyman provided an account of this mode of diffuse attention employed in perception (Prettyman, 2023).

In contrast to these previous accounts of diffuse attention in perception, I offer a broader conception of diffuse attention that generalizes from external to internal forms of attention and captures deep commonalities between the two. This conception will also provide us with the benefit of better explaining the interaction between tasks that require internal and external forms of diffuse attention. Following the prior literature in psychology and philosophy, I believe that attention in general has both external and internal forms (Fortney 2018, 2020; Chun et al., 2011). In addition to attending to objects represented in our perception of the external world, we can also direct our attention to stimulus-independent contents, including thoughts, memories, and mental imagery.

I propose that mind wandering critically involves a diffuse form of internal attention. This account of mind wandering enables me to explain the interaction between mind wandering and other tasks, and especially its benefits for concurrent tasks. Mind wandering also induces both internal and external diffuse attention in concurrent tasks, and can thereby improve the agent's performance in concurrent tasks when these tasks happen to be the type that are better performed with diffuse attention. Some of the tasks that fall into this category include, for example, RSVP perceptual task, as well as the generation of creative thoughts, which often involves internal diffuse attention.

How do creative thoughts emerge?

As I have suggested in Chapter 1, a wide range of tasks can plausibly be better performed with diffuse attention, including not only perceptual tasks like the RSVP task, but also cognitive tasks like creative thinking. In Chapter 2, I focus on the case of creative thinking. More specifically, I am interested in the following questions: how does the generation of creative thoughts involve internal diffuse attention? Does it involve the same type of internal diffuse attention as the type in mind wandering? And moreover, what other processes work alongside internal diffuse attention in the generation of creative thoughts? These are some of the questions that I tackle in Chapter 2. In Chapter 2, I provide an account of the underlying mechanisms of the generation of creative thoughts. My account tries to answer the following question: why are some people better at coming up with creative ideas while others struggle? Why do creative writers, such as Oscar Wilde, generate brilliant metaphors like “We are all in the gutter, but some of us are looking at the stars”, while run-of-the-mill influencers on social media can only come up with tropes?

I argue that previous accounts in philosophy about the generation of creative thought are incomplete due to overlooking the role of what I call “memory gists”. Memory gists are memory contents that represent more abstract or qualitative features that are extracted from the specific, surface level features in the memory representations that were initially encoded in memory. I argue that generating and using memory gists in memory search enables highly creative people to form connections between memory contents that are not usually associated with each other by revealing their commonalities shared in their gists.

Moreover, I argue that different mechanisms underlie online and offline generation of memory gists. Online generation of memory gists involves internal diffuse attention directed towards contents in our memory. However, the difference between the type of internal diffuse attention deployed in creative cognition versus that deployed in mind wandering is that, in creative cognition, internal diffuse attention is guided and constrained by memory gist. The targets of internal diffuse attention here are not just any unconstrained content, but rather contents that share the features represented in the gist. Disparate contents that are very different in their surface-level features can still share the same abstract features represented in the gist, and this enables creative people to attend to and retrieve contents distant from each other across the conceptual space.

Meanwhile, offline generation of memory gists involves memory consolidation during sleep or wakeful rests, while diffuse attention still plays a role in the retrieval of these gists. Memory consolidation is a process that humans go through during both sleep and wakeful rests (Lewis & Durrant, 2011; Robin & Moscovitch, 2017; Rångtell et al., 2017), where our memory system extracts and highlights underlying commonalities between different memory contents that might have some superficial differences. For example, someone who has been to many different kids' birthday parties might, over time, lose their memory of the specific details of each event, such as what kind of cake they had at a specific party. However, through replaying the overlapping parts of the different episodes in memory consolidation, they gradually come to form the beliefs about what kids' birthday parties are like in general: they tend to involve many kids, balloons, and a cake (Lewis & Durrant, 2011).

Drawing on evidence from empirical studies on memory and sleep (Wagner et al. 2004), I argue that memory consolidation plays an important role in the offline generation of memory gists: the consolidation process enables the memory system to extract underlying commonalities between contents that appeared to be very different on the surface. When the agent returns to consciously thinking about the problem, these non-occurrent representations can then be boosted into conscious cognition with diffuse internal attention, and lead to creative insights.

What is the epistemic role of memory?

Chapter 3 builds on the account I provided in Chapter 2 of how memory contributes to creative thinking, and explores the epistemic role of memory more generally in all kinds of reasoning. Given that memory plays an active role in processing previously acquired information through processes such as memory consolidation, how do we understand the role it plays in our epistemic lives? Does it generate new epistemic values in addition to merely preserving epistemic values, and what kinds of epistemic values does it generate? These are some of the questions that I tackle in this chapter.

According to a traditional view in philosophy of memory sometimes called the "storehouse view", memory merely functions to preserve previously acquired information, such as information acquired through perception (Burge, 1997; Goldman, 2009). I draw on empirical evidence about the active processes that happen during memory consolidation to argue against the storehouse view. I argue instead that one of the functions of memory is to improve our understanding of subjects and systems represented in the contents that we previously acquired. Memory therefore generates new epistemic values, especially understanding, instead of merely preserving epistemic values generated through other means, such as through perception and conscious cognition.

What is understanding? Following the literature on understanding, I assume that to understand something, the agent needs to see how a wide range of things can fit together (Riggs, 2003; Grimm, 2011), through "grasping [...] explanatory and other coherence-making relationships in a large and comprehensive body of information" (Kvanvig 2003, 192). Moreover, the object of understanding is sensitive to the goal of the agent, and can change even if the subject matter remains the same. For example, when studying basil in a culinary context, one needs to understand its aroma and flavor, and when studying it in a botanical context, one needs to understand what kind of environment is conducive to its growth. One can therefore improve their understanding through accepting new representations that are both *structured*, in the sense that the representations represent relationships between different components of the object, and *goal-sensitive*, in the sense that they represent and highlight features that are crucial given the goal of the agent.

By focusing on understanding, my account differs from the other accounts in the current literature that also argue that memory generates new epistemic values, but focus on the claim that memory generates epistemic justification, instead of improving understanding. While I am open to the claim that memory also generates other epistemic values, improving understanding better captures one of the crucial causal role functions that memory plays. I argue that in fact, during memory consolidation, our memory system systematically generates representations that satisfy the two features of being structured and goal-sensitive by further processing information that the agent previously acquired.

To make this argument, I draw on evidence for two kinds of processes that occur during memory consolidation. First, as I have already argued in Chapter 2 regarding the role of memory consolidation in generating memory gists, memory consolidation extracts general patterns by uncovering abstract features shared by disparate items that the agent was not aware of before. Second, memory also recognizes associative and inferential relations between disparate contents that were acquired in different contexts, and agents often were not aware of these relations before memory consolidation (Ellenbogen et al., 2007).

To illustrate, suppose you just moved to Somerville and are researching different restaurants in the area. You learned from different contexts that restaurant C is better than D, A is better than B, and B is better than C. Later, when your friend asked you whether A or D is better, you really don't know the answer because you are having difficulty putting it all together. Chances are you will much more likely produce the right answer if you come back to this question the second day after a night's sleep. This is because empirical evidence shows that memory consolidation generates new representations about relationships between different contents. Thanks to these two processes during consolidation, memory generates the specific type of representations that improve understanding. By drawing empirical research on memory consolidation, a process that has received very little attention from the philosophy literature, my account therefore argues that memory plays an active role in our epistemic lives, and especially in improving our understanding.

Chapter 1: Mind Wandering as Diffuse Attention

Abstract

This paper reconciles an inconsistency between the benefits of mind wandering and a prominent conception of attention in philosophy and cognitive psychology, namely, the prioritization view. Since we prioritize the information in a task less if we are doing it while mind wandering compared to solely concentrating on it, why does our performance in the task sometimes improve when we are mind wandering? To explain this, we offer a conception of diffuse attention that generalizes from external to internal forms of attention and captures deep commonalities between the two (Chun et al. 2011). We conceptualize mind wandering as an instance of internal diffuse attention that also induces diffuse attention in other concurrent tasks. Moreover, *pace* the prioritization view, certain tasks are in fact better performed with diffuse attention, which prioritizes information less than focused attention. Our account of mind wandering as an instance of the more general category of diffuse attention improves upon current leading philosophical views of mind wandering (such as Irving 2016) by better explaining the interaction between mind wandering and other tasks.

1. Introduction

In this paper, we solve a puzzle about mind wandering and attention. Often, mind wandering is thought of as detrimental to our performance in another task. Recall the last time your mind started wandering as you were reading a difficult academic article, or trying to follow a talk at a conference. As your mind starts wandering, your attention is diverted from the task, leading to impaired processing of the information in the article or the talk. You start to lose track of the arguments in the article, or the research that the speaker is presenting (Smallwood and Schooler 2006). Indeed, many studies show that mind wandering hinders performance in a variety of tasks that require attending to external targets, such as reading, learning a sequence, or performing a go-no go perceptual task (Brosowsky et al. 2021, Christoff et al. 2009, Smallwood and Schooler 2006).

Intuitively, you are prioritizing the task you are working on less if your mind is wandering, compared to when you are focusing solely on the task. This idea is articulated in a standard characterization of attention in cognitive psychology, according to which paying attention to some target consists in *selecting* this target and *prioritizing* its processing with respect to other, competing processing (Carrasco 2011, Fiebelkorn & Kastner 2020). Moreover, selecting and prioritizing information are often linked to how attention performs what is often taken to be its crucial function: improving processing of the information that is attended to (Wu, 2014). For example, attention makes perceptual representation less noisy and can be more quickly processed (Anderson & Drucker, 2013), and it is often assumed that attention has these processing benefits thanks to the fact that it selects and prioritizes certain information over others.

In philosophy, a version of this view that attention improves processing by prioritizing information is expressed by Sebastian Watzl's (2017, 2023). Watzl takes attention to play the role of introducing a priority structure to mental representations by organizing them into an ordered set: attention structures our stream of consciousness by prioritizing some mental representations over others. This enables attention to facilitate processing of prioritized information. Watzl argues that priority of a certain information is closely tied to its usability by other mental processes downstream (Watzl, 2023).

We interpret this view, which we will refer to as "the prioritization view of attention", as making the following two commitments: (1) Attending to information *P* requires selecting and prioritizing *P*, and (2) more prioritization of information *P* improves processing of *P* and also improves the related task-performance. In accordance with this view, if during mind wandering one is prioritizing the perceptual input needed for a task at hand less, we would only expect that performance will suffer, compared to conditions where subjects are solely focused on the task. However, mind wandering is not always detrimental for task-performance (Brosowsky et al. 2021, Baird et al. 2012). Surprisingly, in some cases, mind wandering even correlates with performance *improvements*. A notable example involves rapid serial visual presentation tasks (RSVP), where participants must detect targets within a string of rapidly succeeding distractors. Mind wandering participants do better on these tasks, compared to fully attending participants (Olivers & Nieuwenhuis, 2005; see also Thomson et al., 2015).[1] As we discuss in detail below, performance improvement was measured by the reduction of the effect known as attentional blink (Raymond et al. 1992), in which participants tend to miss the second of two rapidly succeeding targets. In this way, we have three seemingly inconsistent claims:

1. When your mind is wandering while performing an unrelated task *T* that requires information *P*, you are not prioritizing *P* as much as you would if you were solely concentrating on the task T.
2. More prioritization of information *P* improves processing of *P* and also improves performance in task *T*.
3. For some tasks, performance is better during mind wandering than when you solely concentrate on the task.

Since claim 3 is clearly in tension with claims 1 and 2, a revision of at least one of these is needed. We propose to solve this puzzle by rejecting 2, which is part of the commitments of the prioritization view of attention. Against the prioritization view, we argue that certain information is better

[1] Thomson et al. (2015) found a correlation between a *disposition* towards mind wandering and performance in RSVP tasks: participants who tended to engage more in mind wandering in everyday life did better than those with lesser tendencies towards mind wandering. This supports the robustness of the connection between mind wandering as a distinctive way of using attention and the resulting distinctive way of processing perceptual information.

processed with *diffuse* rather than focused attention, even though diffuse attention does not prioritize information as much as focused attention does (Prettyman 2022, 2014; Gopnik 2009). More importantly, we propose an account of diffuse attention that expands beyond previous discussions of diffuse attention that focus on perception. Our more expansive conception of diffuse attention includes mind wandering as an instance of internally directed diffuse attention (Chun et al. 2011; Fortney 2018, 2020).

We propose that by unifying mind wandering with other mental processes all under the category of diffuse attention, we can better understand their interactions. Thus, an advantage of our view is that it provides new insights into the effect of mind wandering on other tasks, i.e., the kind of tasks that mind wandering impairs, improves, or neither impairs nor improves. Our account is thus an improvement on leading philosophical views of mind wandering, such as Irving's (2016) unguided attention view. A further advantage is that our view also provides new insight on the nature of attention: *pace* the prioritization view, certain contents, whether internal or external, are better processed when they are less prioritized.

The paper is structured as follows. We start by introducing the literature on diffuse attention in perception (§2) and use it to motivate the view that there is an internal analogue to external diffuse attention (§3). Then, we suggest that mind wandering has the signature characteristics of internal diffuse attention and show how this view can explain the effect of mind wandering on other tasks (§4). Finally (§5), we explain how our view improves our understanding of mind wandering, specifically, beyond the unguided attention view (Irving 2016, 2021) and of attention more generally by revising the standard prioritization view. We conclude (§6) by laying down some questions that our account poses for future empirical research.

2. Diffuse Attention, Internal and External

In the last two decades, psychologists have been proposing that at least in perceptual attention, there is a *distributed* mode of attention besides focused attention (Treisman 2006, Gopnik 2009, Srinivasan et al. 2009; Demeyere & Humphreys 2007; Chong & Evans 2011). This attentional mode has a broader scope, kicks in more quickly, and targets different information than focused attention, even though it sometimes can be used alongside focused attention. In the words of Gopnik (2009), attention can be either a focused "spotlight", or a less focused "lantern". With "lantern" attention, the scope of attention is larger and enables the processing of a larger scope of content. The properties that are better processed with this distributed, "lantern" attention include the gist of a scene (Li et al. 2002; Oliva & Torralba 2006) and statistical properties like mean sizes or colors of groups of objects (Chong & Treisman 2005; Treisman 2006; Srinivasan et al. 2009; Demeyere & Humphreys 2007; Chong & Evans 2011). Li et al. (2002) also show that subjects can perform two perceptual tasks (a demanding letter discrimination task in their focal attention and a natural scenery categorization task in their peripheral vision) simultaneously just as well as they can perform the two separately. Other studies also suggest that a spatially broader focus of attention increases sensitivity to lower spatial frequencies, while narrowing the focus of attention increases sensitivity to higher frequencies (Oliva & Torralba, 2006). For example, in gist perception, subjects can categorize whether a low-resolution image shows a kitchen or a street based on configurational properties of the image without processing the local properties in the specific items in the image. Finally, a broader focus of attention is associated with increases in temporal resolution: when subjects are attending to a broader portion of space, they are better able to process the temporal details of a perceived event, such as the order of appearance of two rapidly succeeding targets (Yeshurun & Levy 2003, Mudumba & Srinivasan 2021).

Some theories conceptualize this more distributed type of perceptual attention as *diffuse attention*. The first (and, to our knowledge, only) systematic account in philosophy is due to Adrienne Prettyman (2022). Her Global Selection view characterizes diffuse attention by the kind of targets it selects, namely, global objects and properties. For instance, a whole scene is a global object, and the set of spatial relations amongst the elements of the scene is a global property. Prettyman uses the example of contemplating a large painting (think, for instance, of Rubens' *The Great Last Judgment* or Siqueiros' *La Marcha de la Humanidad*). Appreciating these artworks requires spreading your attention widely, to pick out the whole painting as a global object, without fixating on any of the local elements it depicts. You also need to direct your attention to global properties, like the composition of shapes and colors.[2]

While we agree with Prettyman that diffuse attention often selects global objects and properties, we offer an alternative conception of diffuse attention that does not consider this to be the distinctive feature of diffuse attention. Instead, we offer a broader conception of diffuse attention that generalizes from external to internal forms of attention and captures deep commonalities between the two. This conception will also provide us with the benefit of better explaining the interaction between tasks that require internal and external forms of diffuse attention.

Following the prior literature in psychology and philosophy, we believe that attention in general has both external and internal forms (Fortney 2018, 2020, Chun et al. 2011). In addition to attending to objects represented in our perception of the external world, we can also direct our attention to stimulus-independent contents, including thoughts, memories, and mental imagery.[3] The Global Selection view does seem to capture distributed visuospatial attention in perception, since statistical and configurational properties plausibly fall under the category of global properties. However, if there are non-perceptual forms of diffuse attention, the feature identified by the Global Selection view might not apply to those. In some stimulus-independent processes such as visual

[2] Nanay (2016) holds a similar view. He focuses on aesthetic attention, and proposes that this way of directing attention is focused on one object (the painting), but distributed over lots of the object's properties (shapes, colors, etc.).

[3] We follow Fortney (2018, 2020) in using "stimulus-independent" in the sense that the stimulus-independent contents do not have direct causal relation with external stimuli. Thus, both visual memory and imagination count as stimulus-independent, despite them having causal dependencies on external stimuli in a very attenuated sense.

memory and imagination, one might attend to contents diffusely in a way that can be captured by the Global Selection view (e.g. attending to a street scene as a whole in one's visual working memory, or imagining a street scene as a whole). However, one can also use internal attention to attend to contents in cognitive processes such as deliberation and making judgments. For example, if you are trying to understand an argument that you remember from a paper you read, you can attend to each premise in turn and hold it in your mind for a few moments.[4] If there is a diffuse mode of this kind of internal attention, the feature identified by the Global Selection view might not apply to it, as it is not obvious that there are global objects or properties that could be the targets of such form of attention.

While characterizing diffuse attention as selecting global objects and properties does not generalize well to internal attention, there are certain features shared by paradigm cases of external diffuse attention that do generalize well to internal attention. One is having a broad scope of focus, which we understand in a general sense to encompass attending to non-pictorial contents. In addition to attending to a broad scope of content in space and time, one could also attend to a broad scope of content in the conceptual space, as is often the case in internal diffuse attention. Conceptual space refers to a representation of contents on the conceptual level where the similarity relation between contents is modeled as the distance between points in a geometric space (Bellmund et al., 2018; Tenenbaum & Griffiths, 2001; Gärdenfors 2000). Contents that are less similar are therefore represented as more distant from each other in the geometric space. This feature is different from how the Global Selection View characterizes diffuse attention. A broad focus does not require that there are global objects or properties that are attended to. It only requires that attention is distributed over a wide range of information, but the information need not be unified into a single global object or global properties.

One example of attending to a broad scope of internal contents is creative thinking. During creative thinking, attention is often directed inwards towards disparate contents in memory that could

[4] See Levy (2022) for further argument and examples.

serve to solve an open-ended problem (Benedek et al. 2017). Think, for instance, about what you do in brainstorming to try and come up with a good sofa design. You would attend to a wide variety of sofas of different styles for inspiration (Carruthers 2020). The focus of your attention is much broader than, say, when you are trying to recall a specific sofa design to see if it will fit with the style of your other living room furniture.[5]

Another feature shared by paradigm cases of external diffuse attention is reduced inhibition of information. For example, in the perceptual dual-task experiment from Li et al. (2002), instead of inhibiting all the information except the targets in one's focal attention, subjects are still attending to information in their peripheral vision (albeit with less prioritization) to complete the secondary task. This feature of exhibiting reduced inhibition of information is also shared by creative thinking. During the brainstorming process, you would often refrain from inhibiting any information that is of some faint relevance to the problem, especially if you are trying to come up with a novel idea that is not typically associated with the problem.

Another case that involves internal attention that shares both of these two features, i.e. having a broad focus and exhibiting reduced inhibition of information, is open-monitoring meditation (Lutz et al. 2008; Fujino et al. 2018). Open-monitoring meditation involves the non-reactive monitoring of the content of experience from moment to moment (Lutz et al. 2008). Meditators must maintain awareness of whatever contents pop out in their minds at any given moment, from bodily feelings and sensations to thoughts and emotions, without further engaging with any of these contents.[6] Meditators therefore attend to a broad scope of contents spread out in the conceptual space. Importantly, these contents are not regarded as distractors, but instead as contents for observation pertaining to the present moment (Fujino et al. 2018). Therefore, meditators refrain from inhibiting any of the contents that they observe. This is different from, say, when you try to focus your attention

---

[5] See Author (Manuscript) for a detailed discussion of diffuse attention in creative thinking.
[6] Though in open-monitoring meditation attention is not *exclusively* directed at internal content, it very plausibly involves internal attention in addition to external attention as one of its components.

on a conference talk, in which case you need to actively suppress irrelevant or interfering information (e.g., the color of the speaker's shirt or the notifications on your phone).

We now propose a general conception of diffuse attention characterized by these two features: (1) having a broad focus, and (ii) exhibit relaxed inhibition, or no inhibition of information. This conception of diffuse attention encompasses both internal and external attention. External diffuse attention, which psychologists and philosophers have been focused mostly on so far, is only one species of this more general notion of diffuse attention. We thus offer the following taxonomy of attention:

1 Focused and external Watching a specific soccer player's performance in a match Tasting the fruity flavor in the wine	**2 Focused and internal** Calculating tips in your mind Recalling an argument from a specific paper
3 Diffuse and external Attending to busy street scene as a whole Watching a meteor shower in the wild	**4 Diffuse and internal** Brainstorming about a sofa design Contemplating your feelings as they come and go

Table 1. A taxonomy of attention

We propose that mind wandering fits in the fourth box of this taxonomy, as we will now argue.

3. Mind Wandering as Internal Diffuse Attention

The two core features of diffuse attention, i.e. attending to a wide range of content, and reduced or absence of inhibition of information, aptly characterize mind wandering. When your mind starts wandering, for example, during a conference talk, your thoughts drift towards a wide range of content, from what to have for dinner, to when to do laundry, or to your next vacation. Compare such a condition with calculating tips in your mind, or recalling a detail of your last meeting with your collaborator . There is a sense in which your internal attention has a narrower and more constrained focus in the latter cases, compared to how this focus is in mind wandering.

Moreover, in support of the idea that mind wandering is accompanied by relaxed or even absent inhibition, some psychologists interpret mind wandering as a result of omission in exercising top-down control over what internal content to attend to (McVay and Kane, 2009). Psychology studies show that subjects with lower working memory capacity tend to engage more in mind wandering. This is likely because they are less capable of exercising top-down controlled attention to sustain contents in the working memory, which lends support to understanding mind wandering as having relatively less inhibition of information.

One might worry that our view cannot capture the dynamic nature of mind wandering suggested in Irving (2016) and Christoff et al. (2016). Both Irving and Christoff et al. emphasize that the content of mental states in mind wandering changes dynamically over time, such that one's thoughts move "hither and thither without fixed course or certain aim". A similar idea was proposed in Sripada (2018), who proposes that mind wandering can be captured as an exploratory train of thought that jumps from one topic to another. All of these accounts allow that in mind wandering, one might be paying focused attention to some thought at any given time. Thus, the distinctive feature of mind wandering is not that it does not involve a narrow focus of attention, but rather that the focus of attention meanders through different contents across time. One might worry that our theory of mind wandering as diffuse attention does not capture this feature.

To address this worry, we suggest that attentional states can be diffuse along two dimensions:

synchronically and diachronically. The former refers to breadth of focus and range of contents at a time, while the latter refers to the same things but over time. Internal synchronically diffuse attention can occur when you try to remember a scene of a busy crossroad that you glanced through very quickly through the window, as you try to bring up a wide range of information in the scene. It can also occur when you try to attend to many thoughts at the same time in your mind. Internal diachronically diffuse attention can occur in creative thinking. For example, when you are trying to come up with a new sofa design, your attention is distributed over a wide range of content through a stretch of time as you explore various ideas (Carruthers 2020). We propose that this diachronic dimension of diffuse attention captures the dynamic nature of mind wandering, as it is manifested in the drifting of attention among a wide range of content during a time stretch.[7] In section 5, we will return to a discussion of Irving's account and Christoff's account, where we argue that not only does our account capture the dynamic nature of mind wandering that these accounts also capture, but our account also improves upon these accounts by better explaining the interaction between mind wandering and other mental processes.

Another potential worry is that there are certain cases of mind wandering where one's mind in fact wanders about *closely related* topics instead of meandering among different ones. *Prima facie*, in such cases the subject seems to pay focused attention both synchronically and diachronically, because the range of attended content is not too wide neither at a time nor over time. However, our model would still characterize such cases as diffuse attention. In these cases, even though the mind is wandering about closely related topics, the pattern of thoughts are importantly different from rumination or goal-directed thought, insofar as there is no inhibition. It is a mere coincidence that in some cases the mind wanders about closely related topics. But if attention were then to shift to other

[7] Of course, an important difference between creative thinking and mind wandering is that the former is guided, however loosely, towards a goal, while the latter is not. This might suggest that creative thinking sometimes requires a smaller temporal focus than mind wandering, as one would need to suppress ideas completely unrelated to the creative problem. To account for this, we think that the distinction between diffuse and focused attention might come in degrees, and more constrained forms of creative thinking are closer to focused attention than mind wandering is.

less related content, there would be no inhibition of such content.[8] In this sense, attention is still diffuse.

With this account in hand, and having addressed some initial worries, we can now explain the kind of tasks that mind wandering improves, impairs, or neither improves nor impairs.

4. Interaction between Mind Wandering and Other Tasks

Several studies have shown that mind wandering improves performance in rapid serial visual presentation (RSVP) tasks. In RSVP tasks, a string of numbers or letters are presented for a tiny fraction of a second at the spot where participants are fixating their eyes (see Figure 1.1). Often, participants must give reports about the letters or numbers ("did you see a '5'?"). Studies find a limitation in this task: if participants are asked to detect and identify two targets (T1 and T2) presented within one such string (e.g., two letters amongst a string of digits), they are unable to do this when T1 and T2 appear too closely together. More specifically, if participants successfully detect T1, and if T2 appears within 200-400 ms after T1, participants will miss T2. This effect is known as the attentional blink (Raymond et al. 1992).[9]

[8] According to Irving (2016), the mark of mind wandering is that attention is then unguided: there is no felt need to keep attention "in track" and/or "pulling it back" if we note that our mind has drifted off. Though this counts as absent inhibition in our view, we do not think that inhibition must be characterized by a *felt* need to redirect attention. Inhibition can also be unconscious, as the literature on unconscious attention demonstrates (Kentridge 2011, Norman et al. 2013).

[9] The specific mechanisms underlying the attentional blink are still unclear. See Chun et al. (2011) for discussion.

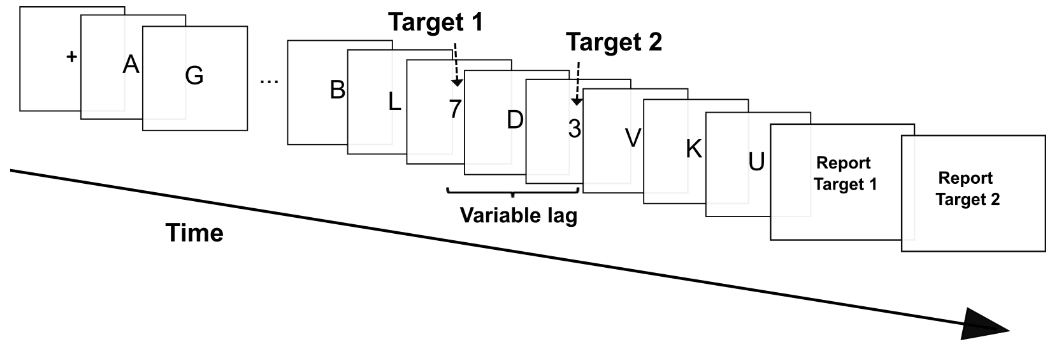

Figure 1.1: RSVP task from Olivers and Nieuwenhuis (2005).

Interestingly, a study from Olivers and Nieuwenhuis (2005) shows that mind wandering participants, who are seemingly prioritizing the stream of images less, show reduced attentional blink, i.e. they show an improved accuracy in detecting T2 while maintaining the same level of accuracy in detecting T1 as participants who only focus on the RSVP task. Olivers and Nieuwenhuis (2005) asked participants to detect two target digits presented amongst a string of letters, under one of three conditions: focusing solely on this detection task, listening to music, or mind wandering (free association starting off from thinking about their holidays or dinner plans). The experimenters found that in the second and third conditions where attention was not solely focused on the detection task, subjects are just as accurate in detecting the first target as when they focus solely on the detection task, but their detection accuracy of the second target significantly improves.

While Olivers and Nieuwenhuis (2005) instructed participants to let their mind wander, another study on the effect of mind wandering on RSVP tasks investigates the effect of the general disposition for mind wandering outside the lab on attentional blink in the RSVP task. They also found a link between mind wandering and improved RSVP performance. Participants who are more disposed to mind wandering in their everyday lives also exhibit a lesser attentional blink compared to participants who are less disposed to mind wandering (Thomson et al. 2015). Though here we focus on the results from Olivers and Nieuwenhuis (2005), the offered considerations can be applied to the results from Thomson et al. (2015) as well.

Olivers and Nieuwenhuis attribute the improved performance in the second and third conditions to a diffusion of attention. They suggest that this diffusion might consist in a "widening of attention", so as to "include" both the perceptual task and the concurrent task. They then add that attention could have also "widened temporally", so that the second target in the perceptual task was also "included."

We think that this proposal is on the right track, but it requires some further articulation. It remains unclear how mind wandering induces diffuse attention in the perceptual task. For say that mind wandering involves a "diffusion of attention". If our considerations from sections 2 and 3 are on the right track, this diffusion of attention is primarily internal. It remains to be explained how *internal* diffuse attention influences *external* information processing. Specifically, more needs to be said about how mind wandering interacts with attention in the perceptual task.

Here, we provide more evidence that the improvement in performance in the RSVP task is in fact due to diffuse attention. Moreover, we argue that mind wandering subjects switch to diffuse attention in the perceptual task because mind wandering is itself a form of diffuse attention. As we will now argue, internal and external attention interact in such a way that when one of them switches to diffuse mode, the other one plausibly does so as well.

Our argument has two parts. First, we show how performance improvements in RSVP tasks can be attributed to diffuse *external* attention. Then, we show how mind wandering, as diffuse *internal* attention, can bring about diffuse external attention in RSVP tasks.

4.1. A Broad Focus of Spatial Attention Improves Temporal Resolution

One plausible explanation for why subjects with reduced attentional blink in the RSVP task can detect the second target shortly after the first one is that their visual processing has higher temporal resolution. Thus, provided that our conceptualisation of broadly distributed spatial attention as diffuse external attention is on the right track, studies linking increased temporal resolution to broader distributions of spatial attention provide evidence for our claim that diffuse external attention is responsible for the reduced attentional blink. We will now discuss two such studies.

Yeshurun and Levy (2003) studied the effects of spatial attention on temporal resolution. They presented participants with either one disk in their visual periphery, staying there for several milliseconds, or two successive disks separated by a varying interval (see Figure 1.2). Participants had to report whether they saw one or two disks. This was harder when the interval between the disks was shorter, for it required participants to detect smaller temporal gaps. The crucial finding was that directing participants' attention to the location of the disks impaired their discrimination performance, compared to a baseline condition where attention was not directed to any specific location in the display. This suggests that *focused* attention impairs temporal resolution.

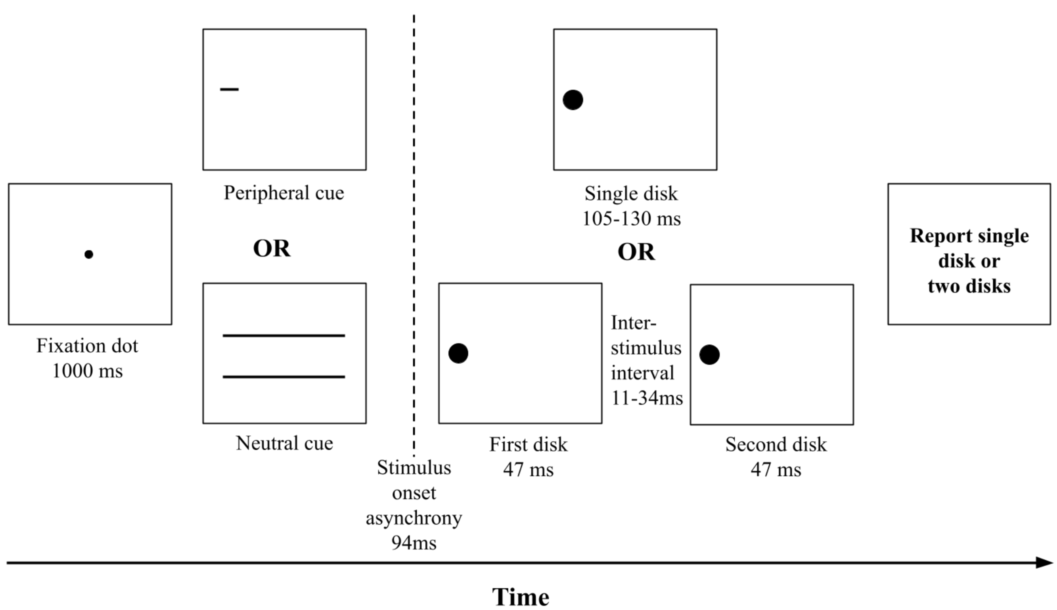

Figure 1.2: Temporal resolution task from Y. Yeshurun and L. Levy (2003).

As Yeshurun and Levy note, these results fit well with other findings revealing that focusing attention on a spatial location increases the apparent duration of stimuli presented at that location. They also note that temporal resolution plausibly trades off with spatial resolution, so that increases in the former correlate with decreases in the latter, and vice versa. For instance, while spatial resolution is best at the fovea and worse at the visual periphery, the opposite occurs with temporal resolution, which is better at the visual periphery. Furthermore, better temporal resolution is associated with larger neuronal receptive fields, as opposed to spatial resolution which requires smaller receptive fields.

For our present purposes, these findings are significant because they suggest a trade-off between focused and diffuse attention, insofar as these two modes of attention typically improve perceptual processing in distinctive *and opposing* ways. It is well known that focused spatial attention increases spatial resolution throughout the entire visual field (Anton-Erxleben & Carrasco 2013), and that it also shrinks down receptive fields (Moran & Desimone 1985). If temporal resolution requires less spatial resolution and/or large receptive fields, it is not surprising that it gets worse with focused attention. But if there was a mode of attention that enlarged rather than shrunk receptive fields, then, even if it also decreased spatial resolution, that mode of attention could well improve the ability to detect the gaps between two stimuli that appear very closely together in time.

We propose that diffuse external attention, in the form of broadly distributed spatial attention, effectively does this job. Participants in Yeshurun and Levy's task must maintain their gaze fixated on the center of the display, but they also know that the stimuli for their task will appear at a random location in their visual periphery. Because they know this, it is plausible that they use a strategy of distributing their attention evenly throughout the display. Thus, their performance in the baseline condition could be attributed to diffuse external attention.

More direct evidence for a connection between improved temporal resolution and diffuse external attention as broadly distributed spatial attention comes from a recent study by Mudumba and Srinivasan (2021). In this study, participants saw arrangements of “local” letters, which together composed a “global” letter; for example, many small H’s arranged to form a big ‘E’ (see Figure 1.3). At the start of each trial, participants were given the instruction of reporting either the local or the global letter. Alongside with the letter display, two discs appeared briefly, one after another, with varying intervals between the two of them. One disc appeared in the upper portion of the display and the other in the lower portion. Participants then pressed the “up” and “down” arrows on the keyboard, to indicate the order in which they saw the discs. Finally, participants reported the target letter by pressing the corresponding key.

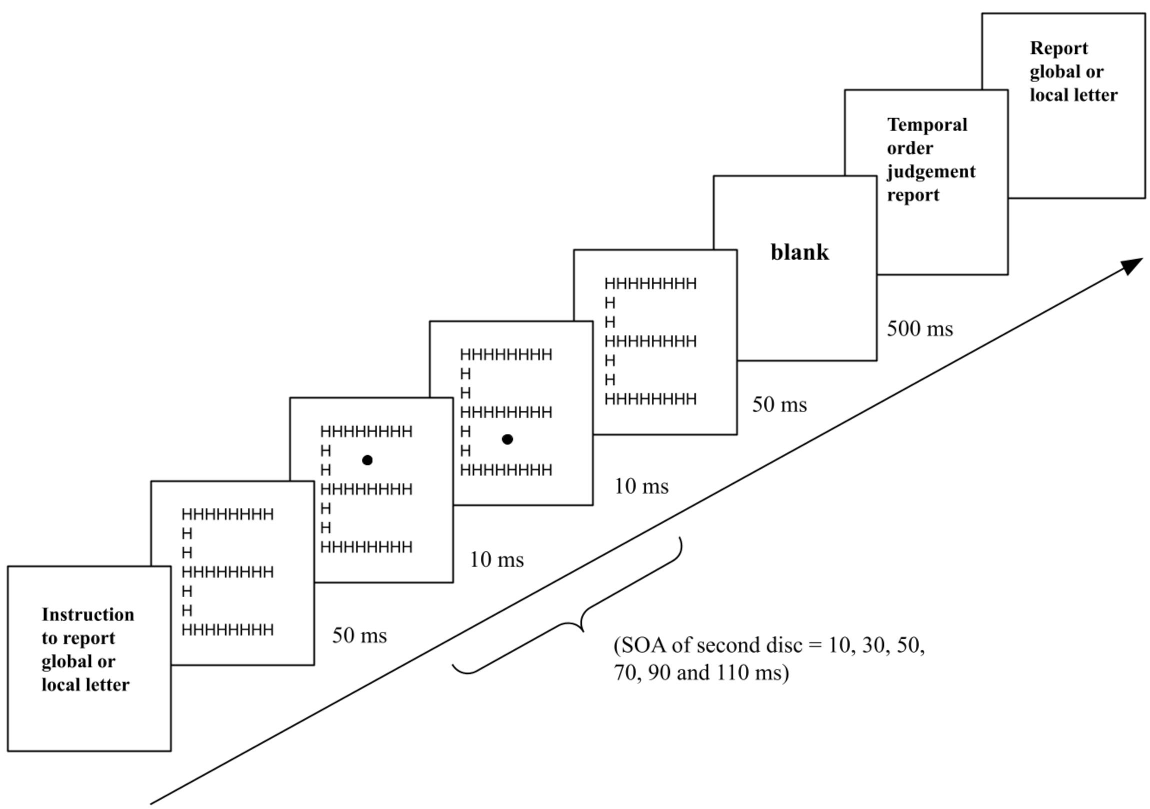

Figure 1.3: Task from Mudumba and Srinivasan (2021).

Mudumba and Srinivasan treated the global letter condition as requiring a broader scope of attention, since reporting the big letter composed of smaller letters requires attention to be spread out over a larger spatial region. They hypothesized that if a broad focus of attention improves temporal resolution, participants would report the order of appearance of the two discs more accurately in the global than in the local letter condition, where attention was more narrowly focused on smaller portions of the display. This is just what they found. So they concluded that broadening the spatial scope of attention effectively improves temporal resolution.

By showing how a widely distributed focus of spatial attention can improve temporal resolution, these studies provide evidence that diffuse external attention can improve our ability to discriminate objects appearing in rapid succession. Hence, these studies support our proposed explanation of the improved performance of mind wandering participants in a RSVP task, as attributable to a diffuse mode of external attention.[10] It remains to be seen why one should take these participants to deploy diffuse external attention.

4.2. Interactions between Mind Wandering and Attention in the Primary Task

So how does mind wandering induce diffuse attention in the primary task? In other words, why should internal diffuse attention induce external diffuse attention?

Several studies suggest that there is interaction between internal and external diffuse attention as we have characterized them in at least one direction, i.e. from external diffuse attention to internal diffuse attention. In a study by Friedman et al. (2003), participants are asked to first focus on certain

[10] It is very likely, however, that improvement in temporal resolution does not give the complete explanation for the effect that mind wandering has on the RSVP task. When T2 appears immediately after T1 within 100ms, non-mind wandering subjects in RSVP tasks can usually capture T2 with a much higher accuracy than when T2 appears between 200-500ms (Chun and Potter 1995). This effect is called “Lag-1 sparing”. This effect would not occur if the attentional blink occurs solely because of a limitation in temporal resolution. However, we note that the accuracy of capturing T2 within 100ms after T1 is still lower than when T2 occurs after more than 600ms when attentional blink no longer occurs. Therefore, it is plausible that limitation in temporal resolution still accounts for the attentional blink to some extent.

perceptual targets, and then perform a creativity task of either generating original uses of a brick within limited time, or generating unusual category exemplars. In one experiment, participants were asked to repeatedly search for the number "3" among several digits either scattered throughout a big area on the screen or displayed close to each other in a small area. According to our proposed taxonomy, these two conditions correspond to conditions of diffuse and focused external attention, respectively. Then, participants were given a minute to write down the most creative use for a brick that is "neither typical nor virtually impossible". Their answers are then rated for creativity.[11] Participants who repeatedly searched through the bigger area provided the most original answers. In another experiment, participants were presented with a selected map of an American state, and asked to either attend to the entire state, or attend to a red star in the center of the map. Again, we respectively take these as conditions of diffuse and focused external attention. Then, participants were asked to type as quickly as possible the name of the most unusual exemplar they can think of of a given category (birds, fruits, furniture, etc). Participants that attended to the broader perceptual target, i.e. the entire map, provided exemplars that were ranked as more original than those that attended to the star. Friedman and collaborators interpret the result as broad perceptual attention inducing internal attention that is broad in the conceptual space. Given our characterization of diffuse attention, we interpret this result to provide evidence that diffuse external attention induces diffuse internal attention.

In another study by Wegbreit et al. (2012), participants are asked to first perform a visual perceptual task, and then solve Compound Remote Associate (CRA) problems. In each CRA problem, participants are presented with three words (e.g. carbon, cat, right), and must think of one word that forms a compound word or common two-word phrase with each of the three words (e.g. "copy" forms "carbon copy", "copycat", and "copyright"). For the visual perceptual task, one group

[11] An answer such as "throw it out the window" is ranked as not creative, and "to grind [it] up and use [it] as makeup" is ranked as relatively creative.

is asked to perform a center-focused flanker task, where they need to focus on the target letter in the center, and ignore the distracting letters or numbers that surround it. Another group is asked to perform a rapid object identification task of identifying a series of visually degraded pictures of animals flashed by quickly. These are conditions of more focused versus more diffuse external attention, according to our taxonomy.

Results show that those who performed the center-focused flanker task showed higher tendency to solve the CRA problems through an analytic method, i.e. by generating words that can form compounds with one of the three given words, and then test out, one by one, whether these words can form compounds with the other two. Meanwhile, those who performed the rapid object identification task tended to use insight instead: they suddenly arrived at the right answer by attending to weak associations among all three words all at the same time. This suggests that diffuse external attention can make internal attention more diffuse, in that the subjects can attend to weak semantic associations among multiple lexicons at the same time.

Given that diffuse external attention seems to induce diffuse internal attention, it is reasonable to think that the interaction goes the other way as well, i.e. that diffuse internal attention also induces diffuse external attention. This proposal is further supported by studies suggesting that internal and external attention share the same underlying mechanism. Prominently, visual working memory studies suggest that sustaining attention to contents in the "mind's eye" and to objects in the external world might recruit the same top-down control mechanism, differing only in whether attended contents are internal representations or external objects (Chun et al. 2011). While this interaction concerns *focused* internal and external attention, we suggest that something similar might happen for *diffuse* attention. Particularly, we think that the lack or relaxation of inhibition in internal attention plausibly elicits a relaxation of inhibition in external attention as well. This idea sits well with some psychological views attributing mind wandering to an omission of top-down control over what content to attend to. Studies show that subjects with lower working memory capacity tend to engage more in mind wandering (McVay and Kane 2009). This is likely because they are less capable of exercising top-down controlled attention to sustain contents in working memory.[12] [13]

So far we have offered reason for thinking that diffuse internal attention can induce diffuse external attention, and that diffuse external attention is a plausible explanation of why mind wandering participants are better able to detect two targets within an RSVP string, compared to fully attending participants. To strengthen our case for these views, we will now review evidence that mind wandering's improvement of RSVP performance is not an isolated case. We will briefly discuss some other tasks that are either facilitated by mind wandering, or at least compatible with mind wandering, and we will explain how our account of mind wandering as diffuse attention can readily apply to these other cases.[14]

4.3. What Kind of Tasks are Better Performed while Mind Wandering?

Besides the RSVP task, what other tasks are better performed while mind wandering? Clearly, not all tasks are better performed while mind wandering. Abundant studies show that mind wandering in

[12] Dixon et al. (2014) suggests another way to think about the interaction between internal and external attention, according to which internal and external attention interfere with each other when both forms of attention require high inhibitory control. Since diffuse attention only requires relaxed inhibition, our theory is compatible with Dixon's theory and might also complement it by explaining how low inhibition internal and external attention interact with each other.

[13] One might worry that our proposal stands in tension with findings on visual information sampling during mind wandering. In a recent experiment, Krasich et al. (2018) found that mind wandering subjects make fewer and longer fixations on a visual scene they are asked to memorize, compared to attentive subjects. This suggests that their spatial attention might not be as broadly spread out as we propose. In response, we emphasize that attention and gaze behavior can be dissociated, as evinced by a large body of empirical research (see Carrasco 2011 for a review). Thus, a fixed gaze is compatible with different attentional states, and in this specific case, subjects might be simply "zooming out" while blankly staring at a target. Later studies show that different mechanisms underlie the longer fixation durations observed in mind wandering subjects and in subjects paying attention under high processing demands (O'Neil et al. 2022), which lends further support for the "blank stare" interpretation.

[14] While we won't discuss in length here, our model also fits well with extant psychological theories of attentional blink. In general, psychological models explain the attentional blink either as a result of overexerting cognitive control (Taatgen et al., 2009), or as a result of the first target taking up all of the attentional resources (Chun and Potter, 1995). In both models, detecting the first target triggers a protection rule to ensure its processing and prohibit processing of any new target in the meantime. Our explanation in terms of diffuse attention would imply disabling the protection rule to allow processing of the second target while the first one is still being processed. This explanation is also suggested by Taatgen et al (2009), though their model only considers the effect of a secondary perceptual task, and not an internal attention task.

fact impairs task performance in many cases, such as when the primary task is a go-no go perceptual task (Christoff et al. 2009), or a reading comprehension task (Smallwood and Schooler, 2006). This makes it seem as a rather curious accident that mind wandering participants do better in one very specific kind of task. However, we think that there is support for a more general claim. There is a more general class of tasks, including the RSVP task, that are better performed (or at least not impaired) while mind wandering. We suggest that this class of tasks are tasks that are better performed with diffuse instead of focused attention.

One notable example are tasks that require creativity. A study by Baird et al. (2012) shows that mind wandering during the incubation stage of creative thinking improves performance in these tasks, even when the contents that subjects entertain during mind wandering are not explicitly related to the task. Participants were first given a creative thinking problem ("generate as many unusual uses as possible for a common object"), and then some were asked to work on a non-demanding task which causes increased mind wandering, while others worked on a demanding task that did not cause mind wandering. When both groups returned to the creative thinking task, the mind wandering group performed much better than the other group. Assuming, as we have suggested above, that the creativity task requires some level of internal diffuse attention (e.g., to spread the focus of attention widely, to activate distant or unusual nodes in semantic association networks), this study provides evidence that mind wandering facilitates better performance in some concurrent tasks that require diffuse attention.

We think that one commonality between the RSVP task and creative thinking task is that they are tasks that are better performed with diffuse attention. While we will not commit to arguing that *all* tasks of this category can be performed better while mind wandering, we think that these studies at least suggest that mind wandering tends to have a positive effect on performance in these tasks.[15]

[15] We also do not claim that the mode of attention employed in the other task is the only factor that affects how it interacts with mind wandering. A study by Brosowsky et al. (2021) shows that mind wandering impairs explicit but not implicit sequential learning. In their experiment, subjects are asked to follow a sequence displayed on the screen

5. Diffuse Attention and the Prioritization View

To close our discussion, we will now explain how our proposal improves on existing views of mind wandering. We will also return to our initial puzzle and explain how our views motivate some important qualifications on the prioritization view of attention.

We have mentioned before that our theory accounts for the dynamic nature of mind wandering identified in recent theories of mind wandering as unguided attention (Irving, 2016), spontaneous thought with mild constraints (Christoff et al., 2016) or exploratory thinking (Sripada, 2018). In addition to capturing the same feature of mind wandering as these previous theories, our account of mind wandering as (internal) diffuse attention also has a distinctive advantage over these alternatives. The advantage is that our account can easily explain why mind wandering is beneficial for some tasks that depend on external information, including but not necessarily limited to RSVP. This is thanks to the fact that our account understands mind wandering as a form of internal attention, and draws on the interaction between internal and external attention. Previous views of mind wandering, on the other hand, cannot readily explain the interaction between mind wandering and tasks that rely on external contents represented in perception.

For example, neither Sripada's account nor Christoff et al.'s account gives much consideration to the role of attention in mind wandering. As a result, it is not obvious how an exploratory mode of *thinking* (Sripada 2018) or the lack of constraints in *thought* (Christoff et al. 2016) shall facilitate the processing of incoming visual information, as in the RSVP task.

---

by hitting the corresponding buttons on a keyboard, but only the explicit learning group is informed that they will be asked to reproduce the sequence later. We think that in this case, the difference maker between the implicit and explicit learning group in terms of their interaction with mind wandering is more likely how information is encoded in working memory, instead of the modes of attention involved in learning. The encoding of information for explicit learning might be more demanding on working memory capacity, and so more affected by mind wandering. This result suggests that factors other than mode of attention employed also affect how tasks interact with mind wandering.

Conceptualizing mind wandering as an attentional phenomenon helps bridging this gap, as attention spans internal and external processing. In this respect, Irving's (2016) unguided attention view is already a step in the right direction. However, Irving's account does not consider the interaction between internal and external attention, and therefore cannot explain the beneficial effects that mind wandering has for certain perceptual tasks, such as RSVP. In comparison, by situating mind wandering in the taxonomy of internal and external attention, our proposal provides a more integrative view where the processes underlying mind wandering are understood as interacting with other processes in the mind.

In addition, our characterization of mind wandering as internal diffuse attention, along with the proposed two-dimensional taxonomy of attention (as external or internal and focused or diffuse) has more general implications beyond explaining mind wandering. In particular, our taxonomy has the additional benefit of furthering our understanding of the nature of attention. Though diffuse attention has been previously discussed in philosophy and psychology, previous theories focus mostly on diffuse attention in visual perception. We provide a more general account, in which diffuse attention can be both internal and external. Since this more general account casts a wider range of psychological phenomena as attentional, it pushes us to hone in into what the defining characteristics of attention are.

For example, one important point of clarification concerns the function of attention. Specifically, our theory puts pressure on the prioritization view. Recall that this view has two commitments: (1) attending to information *P* requires selecting and prioritizing *P*, and (2) more prioritization of information *P* improves processing of *P* and also improves the related task-performance. Our view rejects the second commitment. In our view, more prioritization does not always improve processing and performance. When we attend to things with diffuse attention, no single piece of information is as prioritized as it would be if we were to attend to it with focused attention, since we are giving as much priority to it as we do to many other contents at the same time.

To phrase it in Watzl's terminology of "priority structure", the corresponding priority structure of diffuse attention is flat, rather than s*piky*, which would be the case with focused attention.[16]

However, as we have demonstrated with empirical evidence, for some specific internal or external contents, paying diffuse attention to them, which involves less prioritization, leads to improved (or at least comparably good) processing compared to when we pay focused attention to them.[17] This suggests that the function of attention might in fact be something more general than how it is characterized by the prioritization view. More specifically, while it might be true that attention functions to improve processing of the attended information (Marchi 2020), this does not necessarily have to be achieved through selection and prioritization. At least, *pace* the prioritization view, it is not necessarily the case that for all kinds of content, the more we prioritize it, the better we can process it.

Finally, we would like to consider a possible objection: someone who interprets selection and prioritization in a very strict way might think that our account rejects (1) as well, and object to our account on that ground. The thought goes as follows: not only are selection and prioritization part of the definition of attention, but that a very specific type of selection and prioritization is required for attention, i.e. the type required for, in Watzl's terms, "spiky" attention. According to this view, attention requires focusing on a narrow scope of contents, and strictly prioritizing it over other contents. This objector might go on to argue that what we call "diffuse attention" on our view is

[16] Watzl (2017, p. 73).

[17] Notably, recent accounts of skilled action and choking phenomena reach similar results. Choking occurs when a skilled performer reflects on their ongoing behavior, thus hindering their performance. Since the standard view of skilled behaviors is that they are automatized and no longer require attention, one could think of these as examples where attention does not improve and in fact impairs performance. However, Bermudez (2017) has recently argued that skilled behavior still depends on an amount of attentional control. In his view, choking and related phenomena occur when attention is misdirected from its proper target, the basic action, to focus on the automated action components or factors external to the action (e.g., audience pressure). A skillful soccer player must maintain her "head in the game", but to do so she does not need to keep track of the angle of her feet. While a novice player must highly prioritize this information, the expert player must deprioritize it to perform at her best. Importantly, deprioritizing information of these action components is compatible with her playing attentively. In our view, her attentive playing is explained in terms of the appropriate breadth of her attentional focus, which "zooms out" from already automatized details.

better characterized as inattention. They might think that the mind wandering subject is not attending to any thoughts at all, because no single thought is strictly prioritized over others.

In response, we would like to note that it would be somewhat controversial to understand the notion of selection and prioritization in attention so narrowly. Watzl, for example, understands prioritization in a general sense, and takes “flat” attention to be a type of attention as well. According to his view, attention merely requires taking one thing to be *at least as* prioritized as another, instead of being prioritized *over* another. Different things can “have a tie” in the priority structure, which is the case when one’s attention is distributed among a wide scope of contents (Watzl, 2017, Ch. 5). We agree with Watzl in understanding prioritization in a more general sense, such that the mind wandering agent is still prioritizing the thoughts that they entertain as their mind wanders, even though they are prioritizing a wide range of thoughts, and none of them is prioritized *over* each other, and the same reason applies to why the agent performing the RSVP task while mind wandering is attending to the targets in the RSVP task, instead of processing them without paying attention to them at all.

While we agree with Watzl in terms of how to understand the scope of attention, there is a crucial difference between our account and current accounts like Watzl’s that already adopt a broad conception attention that includes “flat” attention: since we reject (2) of the prioritization view, we do not think of diffuse attention as a “lesser” version of focused attention in terms of its processing benefits. Instead, we show that diffuse attention has its own distinctive processing benefits, and that *pace* the prioritization view, different kinds of attention achieve their own distinctive processing benefits through different kinds of prioritization, be it spiky or flat.

6. Conclusion

We have proposed that by understanding mind wandering as a state of internal diffuse attention, we can explain the interaction between mind wandering and other tasks, and especially the benefits of mind wandering for other tasks that are best performed with diffuse attention. With this, we have provided a solution to a puzzle about mind wandering and attention, which has the further advantages of deepening our understanding of both attention and mind wandering.

Admittedly, our proposal has limitations. To more thoroughly investigate whether mind wandering has the tendency to improve tasks that are better performed with diffuse attention, we will need more evidence about how mind wandering interacts with other tasks that fall into this category beyond the ones we discussed. Some of these other tasks include magnitude estimation (Chong & Evans 2011) and temporal resolution tasks (Yeshurun & Levi 2003; Mudumba & Srinivasan 2021). We will also need to investigate what other specific tasks (besides RSVP) can benefit from mind wandering, and whether these tasks can be conceptualized as recruiting diffuse attention, as we suggested for RSVP and creative thinking. Answering these questions will require more empirical investigations as well as philosophical interpretations of the results.

Finally, while we have made a case against the prioritization view of attention, more theoretical work is needed for establishing either an alternative to the prioritization view, or a revised version of it. In either case, the new theory will need to avoid the bias of taking focused attention to be the paradigm case of attention, which likely underlies the prioritization view. Instead, the new account of attention will need to take seriously the fact that at least certain contents are better processed with internal or external diffuse attention, and understand the function of attention in light of this fact.

Chapter 2: Memory Gist as Mechanism for Creative Insight

Abstract

Why are some people better at generating creative ideas than others? This paper argues that previous accounts in philosophy give merely incomplete answers to this question due to overlooking the role of what I call "memory gists". Memory gists are memory contents that represent more abstract or qualitative features that are extracted from the specific, surface level features in the memory representations that were initially encoded in memory. I argue that generating and using memory gists in memory search enables highly creative people to form connections between memory contents that are not usually associated with each other by revealing their commonalities shared in their gists. Moreover, I argue that different mechanisms underlie online and offline generation of memory gists: the former involves a mode of attention that is the internal analog of diffuse attention in perception, while the latter involves memory consolidation during sleep or wakeful rests. My account therefore highlights the important roles that internal diffuse attention and generative processes in our memory system play in creative thinking.

1. Introduction

What enables some people to easily generate creative metaphors while others struggle? One intuitive answer is that more creative people are able to think past how we usually associate concepts, and instead draw out highly non-obvious but fitting connections between concepts that are not often associated with each other. Take, for instance, Oscar Wilde and one of his famous quotes, "We are all in the gutter, but some of us are looking at the stars".[18] Arguably what renders Wilde's metaphor more creative than the run-of-the-mill positivity quote on social media is the remarkable connection

[18] This quote is originally from the play *Lady Windermere's fan* (Wilde 1995).

that Wilde draws between the effect of a positive mindset in the face of adversities and the experience of seeing the gorgeous starry sky from the gutter. “Stars” and “gutter” are not often associated with “positivity” and “adversity”. The former are highly concrete words about natural or urban scenes, while the latter are very abstract terms. However, a creative writer is able to connect them in a metaphor, and capture the deep similarity between two experiences despite their drastic differences on the surface. What, then, are the underlying psychological mechanisms that enable Oscar Wilde to draw the connections that less creative people have great difficulties coming up with?

Beyond metaphors, the ability to form connections between contents in memory such that the connections are both unusual, and also appropriate for the task or context, seems to be crucial for the generation of creative ideas in general. On the one hand, it is generally accepted in the creativity literature that thinking creatively requires someone to “think outside the box”, and hence get past the ideas most often associated with the problem and reach into conceptually more distant domains. Connecting concepts that are distant from each other, or concepts that are distant from the problem, facilitates creative thinking, since such connections tend to be unusual and therefore novel (Finke et al. 1996). On the other hand, not just any combination of conceptually-distant ideas will do, but the combination also needs to be appropriate. A sentence like “We are all in the piano, but some of us are fishing” would not be creative, but rather just nonsensical. This intuition about the two features of creative idea generation is captured in the standard conception of creativity, where how creative an idea is is conceptualized as a combination of its novelty and its usefulness for solving the task (Boden, 2004; Barron, 1955; Stein, 1953).

Not everyone has Wilde’s flair for composing creative metaphors, and people’s ability to generate creative ideas more generally varies a lot. What then are the underlying mechanisms that explain individual differences in the ability to generate creative ideas? In other words, what underlying mechanism allows more creative people to retrieve and connect conceptually distant contents in a way that is also appropriate for the context? This paper provides one answer to this question by focusing on the role of memory gists. I will argue that overlooking the role of memory gists prevents previous accounts in philosophy from providing an adequate answer to this question.

"Gist" is a terminology that psychologists use to refer to "the essence of a set of memory contents" (Lewis & Durrant, 2011). It is "fuzzy and qualitative", such as remembering that one quantity is bigger than another without remembering the exact numbers (Abadie et al. 2013), and it often contains "global patterns or relationships" extracted from the initial experience (Reyna & Brainerd, 1995).[19] Discussion of memory gists can be traced to the fuzzy trace theory, according to which our long term memory stores both a verbatim memory trace, and a gist trace (Reyna & Brainerd, 1995). The former contains the details of the original experience, while the latter contains a schematic summary extracted from the original experience without the specific details. The verbatim memory trace is often lost later on, while the gist trace is preserved.

For example, in the famous Deese–Roediger–McDermott paradigm, upon reading a list of words related to healthcare ("nurse, patient, hospital, ambulance, etc"), subjects often form a gist trace, based on which they believe that they have read a list of words that share the commonality of being common words related to healthcare. Later, when tested on what words they saw, many have the illusion that they saw other healthcare-related words that weren't on the list (e.g. "doctor"). This is because they use the gist trace to help them retrieve specific words (Roediger & McDermott, 1995).

In addition to memory gists formed at the time of encoding, others have also argued that memory gists can be formed in the later process of memory consolidation and retrieval, such as through hippocampal–neocortical interactions (Robin & Moscovitch, 2017). These memory gists can

[19] Existing literature disagrees on whether gist can only be extracted from a single memory episode (Gilboa & Marlatte, 2017), or can be extracted from both a single memory episode, and multiple memory episodes (Lewis & Durrant, 2011). My notion of "gist" follows the latter. Gist extracted from multiple memory episodes also differs from how some use the notion of "schema", which are not only extracted from multiple memory episodes but also "detail-rich and highly elaborate" (Gilboa & Marlatte, 2017). The contents of gist are abstract, generic and qualitative, and do not contain as much detail as the initially encoded contents in memory.

be formed from a single memory episode, as well as from multiple memory episodes. For example, someone who has been to many different kids' birthday parties might, over time, lose their memory of the specific details of each event, such as what kind of cake they had at a specific party, but through replay the overlapping parts of the different episodes in memory consolidation, they gradually come to form the gist of what kids' birthday parties are like in general, which might involve many kids, balloons, and a cake (Lewis & Durrant, 2011).

In accordance with discussions about gists in psychology, I understand memory gists to be memory contents that represent more abstract or qualitative features that are extracted from the surface level features in the memory representations that were initially encoded in memory. How do memory gists facilitate the generation of creative ideas? I separate my discussion into online and offline generation of memory gists. The former happens during retrieval as the subject is consciously thinking about the problem, while the latter happens during memory consolidation when the subject is not consciously thinking about the problem, such as during sleep or wakeful rests. In both cases, I argue that the generation of gist memory facilitates generating creative ideas by connecting disparate contents in memory that share the same abstract features. Memory gists can therefore guide our memory retrieval, such that disparate contents that share the same gist in the context of the problem can be retrieved together.

I also identify two underlying mechanisms that are crucial for generating memory gists online and offline respectively. Online generation of memory gists involves a mode of attention that I call "internal diffuse attention". Many have argued that humans employ diffuse instead of focused attention to process global objects and properties represented in visual perception (Prettyman, 2023; Chong & Treisman, 2005). Moreover, in addition to attending to external contents in perception, we can also attend to internal, stimulus-independent contents in memory, imagination and cognition, such as when we are daydreaming, when we try to imagine a scenario, or when we try to solve a math problem (Chun et al., 2011). There is also evidence that internal and external attention share some of the same control mechanisms, despite being directed towards different contents (Chun, 2011). I will argue that generating creative ideas involves an internal analog of external diffuse attention in perception.

Meanwhile, offline generation of memory gists involves memory consolidation during sleep, or during wakeful rests when individuals are not consciously thinking about the problem that requires a creative solution. In these cases, memory consolidation sometimes generates more abstract gist contents extracted from the more specific contents. The creative individual can then later retrieve the gist contents when they return to thinking about the problem, and use the memory gists to generate new creative insights. Memory consolidation has received very little attention from the philosophical literature[20], and few philosophers have explored the role that memory consolidation plays in our epistemic lifes.[21] My argument for the role of memory consolidation in offline gist generation in creative thinking therefore also highlights the crucial role of memory consolidation in our cognitive economy, which has been overlooked by most of the philosophical literature.

Before I dive into the specifics of my account I would also like to clarify its scope. First, according to the seminal "GENEPLORE" model of creativity proposed by Ronald Finke and co-authors (Finke et al. 1996), creative problem solving comes in two different stages: in the early generation stage, subjects generate a wide range of preliminary ideas, and in the later exploration stage, subjects then evaluate the ideas generated in the first stage and explore ways to refine them. Recent work in philosophy has also suggested that both the initial generation of ideas and the further refinement and reorganization of ideas are crucial for improving both the novelty and usefulness of the ideas in creative thinking (Langland-Hassan, 2020; Stokes, 2014, 2017). My account only focuses on the generation stage of creative thinking, and I leave it open that in the exploration stage, memory gists might not play roles that are as significant as in the generation stage. Second, my account does

[20] With the notable exception of Aronowitz (2019).
[21] For example, the debate between generativism and preservationism in the discussion of the epistemic role of memory has barely explored the implication of memory consolidation (Frise, 2023).

not purport to explain *every* case of creative thinking. Instead, I only argue that memory gists and the attention and memory processes that underlie its generation and retrieval are crucial for many cases of creative thinking, and that they have been overlooked by the current philosophical literature on creativity. Overlooking the role of memory gists causes the accounts in the philosophical literature to give incomplete explanations for individual differences in the ability to generate creative ideas.

Here is a roadmap for the rest of the paper. In section 2, I discuss theories about the online and offline generation of creative ideas in the philosophical literature, and argue that overlooking the role of memory gists leads to a gap in their explanation of individual differences in the ability to generate creative ideas. In sections 3 and 4, I present my positive account. In section 3, I provide my account of how memory gists are generated online, which I argue involves internally directed diffuse attention. I also explain what this mode of attention is. In section 4, I provide my account of how memory gists are generated offline, which I argue involves memory consolidation. In section 5, I conclude by comparing my account with several existing accounts of creativity in psychology, as well as highlighting the contribution and further implications of my account for the philosophical understanding of attention and memory.

2. How do creative ideas emerge? A discussion of the philosophical literature

In this section, I will discuss some theories in the philosophical literature that explain how creative ideas are generated in terms of some underlying mechanisms that they identify. I argue that while it is highly plausible that these mechanisms are important for generating creative ideas, their accounts only give incomplete explanations for individual differences in the ability to generate creative ideas because they overlook another important mechanism: the generation and use of memory gists. Following Carruthers (2020), I will divide my discussion into mechanisms that underlie online and offline generation of creative ideas. Online generation of creative ideas refers to when the agent generates a creative idea while intentionally solving the problem (e.g. when someone is intentionally trying to come up with a creative metaphor), while offline generation of creative ideas refers to when the agent is not intentionally solving the creative thinking problem but a creative idea either pops up into conscious cognition, or a non-occurrent mental representation is formed and later facilitates a creative insight when the agent returns to thinking about the problem.

2.1. Online generation of creative ideas

I start with a discussion of several accounts of online generation of creative ideas that focus on cognitive control and inhibition (Carruthers, 2020; Irving, 2016, 2021). In his discussion of online generation of creative ideas, Carruthers (2020) identifies inhibiting prepotent ideas as an important mechanism. This is because the ideas that pop up first are often commonly associated with the prompt, and therefore not unusual. Inhibition of prepotent ideas gives more time for more ideas to be activated and emerge into consciousness, and the ideas that emerge later also have a higher chance to be creative, since they are not as usually associated with the problem given, and so conceptually more distant from the prompt. I agree with Carruthers (2020) that suppression of prepotent ideas is a crucial mechanism for online generation of creative ideas, but Carruthers' (2020) account remains incomplete: suppose a creative writer, like Oscar Wilde, and a not so creative writer that writes positivity posts on social media are both trying to come up with a metaphor about positivity. Even if the less creative writer inhibits the prepotent idea as much as the more creative one does, they might still have trouble coming up with a creative metaphor: they might simply have no ideas popping into their mind, because they cannot think of anything other than words that are commonly associated with "positivity"; or they might think of some unusual sentences but do not convey the appropriate message, and therefore not as useful.

One way to explain the difference between these two people requires analyzing the differences in the ways in which they conduct their memory search, beyond inhibiting prepotent ideas. My account based on memory gists provides a plausible explanation for this difference. As I will argue in my account, the more creative person is able to search through a wide range of contents in their memory to form connections that are both unusual and useful for the task, because they are able to generate and use memory gists to guide their search. More creative people generate memory gists that enable them to see the underlying commonalities that the content contained in the prompt or under consideration shares with other contents that are not often associated with them, which thereby allow them to search through a wide range of contents to form connections that are both unusual and useful.

Another account from Irving (2016, 2021) instead focuses on *reduction* of inhibition as the primary underlying mechanism of creativity. Zac Irving offers an unguided attention account not of creative idea generation, but of mind wandering. However, he also proposes a framework for relatively unstructured forms of thought broadly conceived, including creative idea generation, where certain features of his theory of mind wandering also apply. Irving argues that mind wandering can be thought of as a form of unguided attention, where the focus of the subject's attention drifts from one topic to another without guidance or purpose. Moreover, Irving argues that at least certain types of creative idea generation, such as brainstorming, are highly similar to mind wandering (Irving 2016; Christoff et al. 2016). According to Irving's theory, brainstorming differs from mind wandering insofar as it has an overarching goal that the agent is working towards (e.g. to come up with a creative sofa design), while mind wandering does not have a goal. However, brainstorming still bears a lot of similarity to mind wandering. On this, Irving writes,

> [O]n my view, conceptions of a goal can have hierarchical structure. For example, algebra problem solvers usually know what sub-goals to perform, and in what order. In contrast, a

brainstormer may have only the broadest conception of her overarching goal. Certain cases of goal-directed attention (e.g. algebraic reasoning) may thus be guided in a more structured way than others (e.g. brainstorming), and therefore seem more dissimilar to mind-wandering. (Irving 2016)

Here, Irving is saying that in more structured forms of thought such as in solving an algebraic problem, there are subgoals nested within the overarching goal of solving the problem (e.g. add the last digits, then do the division, etc), such that there are strong constraints over what to attend to. However, in brainstorming, there are no subgoals nested within the overarching goal of solving the creative problem, and therefore the agent's attention remains relatively unconstrained within the scope of working towards the overarching goal. For example, when brainstorming about a creative sofa design, almost any thought somewhat relevant to sofas is permissible, since there is not a set structure or procedure to follow to solve the design problem. My mind can wander through different styles, materials, colors, so long as what I am thinking about is still relevant to the creative task, and I exercise relatively little constraint over what I attend to within this scope.[22] Therefore, according to Irving, at least certain forms of creative idea generation (especially brainstorming) share the underlying mechanism with mind wandering of having relatively low level of inhibition towards what content to attend to.

My response to Irving's account is similar to my response to Carruthers' account: I believe that Irving's account reveals some mechanisms that underlie creative idea generation, but it remains incomplete in explaining individual differences in the ability to generate creative ideas. Consider the following result from Yoed Kenett and his collaborators on the different memory retrieval patterns between more and less creative individuals (Kenett et al., 2014, 2016; Beaty & Kenett, 2023). In one

[22] Irving also expresses a similar view in his collaborated work (Christoff et al. 2016), where he argues that creative thinking lies on a spectrum in between mind wandering and more structured forms of thought in terms of the strength of deliberate constraints.

of their experiments, subjects are divided into a high-creativity group and a low-creativity group based on their self-reported results on a questionnaire about their creative achievements (Kenett et al., 2016).[23] Both groups are then asked to complete a verbal fluency test, where they are prompted to generate as many animal categories as they can think of within 60 seconds. Based on the behavioral data from both groups, Kenett and collaborators then plotted the semantic network for each group as a whole. In the networks, each node represents an animal category produced by a sufficiently large percentage of participants in the group, and the existence of an edge between two nodes indicates that the group as a whole tends to retrieve the two categories together (i.e. a significant number of participants in that group who generated one of the category also generated the other one).[24]

Comparing the networks generated by the less creative and more creative group, Kenett and collaborators found that those who self-reported more creative achievements in the questionnaire generated less modular semantic networks. In other words, there are fewer sub-groups isolated from each other in the network generated by the more creative group. Instead, the sub-groups are all meshed together through a web of connections that link up categories in different sub-groups. To use one word as an example, in the less creative group, “goat” is connected with sheep, pig, chicken and cow, all of which are farm animals. Meanwhile, for the more creative group, “goat” is also connected with some non-farm animals, including bat and turtle.

While the verbal fluency test is not itself a creativity test, the difference between the memory retrieval behavior between the high creativity group and the low creativity group likely reflects how they behave in creative thinking as well. The less creative group likely gets stuck with mundane ideas that are most often associated with the prompt, and are therefore semantically close to the prompt.

[23] In this specific study, Kenett et al. further divide subjects into four groups based on another measure of creativity besides self-reported creative achievement (i.e. fluid intelligence) in order to study the relation between semantic network and the two measurements separately. However, for the purpose of this paper and for the sake of simplicity, I will restrict my discussion of this study to its findings on the relation between semantic networks and self-reported creative achievement.

[24] For the purpose of studying structural properties of the network, the weight of the edges are binarized. In other words, after filtering out weak associations, all the edges are assigned a uniform weight of 1, instead of having different weights that reflect the strength of association between the connected categories.

This is akin to how they get stuck with other farm animals once they think of “goat”. Meanwhile, the more creative group can leap between ideas that are semantically very distant. This result is difficult for Irving’s account to explain. In the verbal fluency test, it is likely that the high creativity group and the low creativity group have equally lowered their inhibition of thoughts within the category of all animals, but they still produce very different memory retrieval patterns, which likely mirrors how they behave in creative thinking as well. One might think that as Carruthers (2020) has pointed out, the low creativity group has trouble inhibiting the ideas that are mundane and pop up first, which explains why they get stuck with closely associated contents. However, it is also possible that even if the low creativity individuals in fact inhibit prepotent ideas, they simply cannot think of more distant ideas: nothing more is popping into their mind, while the mind of the more creative person leaps easily between semantically distant contents. To explain this possibility, one would need to identify other factors that affect individuals’ memory retrieval patterns. As I will elaborate on in Section 3, my own account does just that by highlighting the role of memory gist generated during memory retrieval in online generation of creative ideas. The more creative person can use the memory gist that they generated to leap beyond the nearby contents and search for a wide range of potentially useful contents that fit the gist. Using memory gists therefore enables them to think past the mundane ideas, and connect ideas not often associated with each other in a way appropriate to the task.

2.2. Offline generation of creative ideas

I will now turn to discuss two accounts of offline generation of creative ideas in the philosophy literature, and argue that they also provide only incomplete explanations for individual differences in creative idea generation, due to overlooking the role of memory gists. Carruthers (2020) also provides an account of the mechanisms that underlie offline generation of creative ideas, which he discusses separately from online generation of creative ideas. In particular, Carruthers takes the following to be a crucial mechanism for generating creative ideas offline: even when we are not intentionally thinking about the problem that requires a creative solution, the relevance network continues to monitor mental representations in the service of our current goals, which sometimes leads to a creative idea to pop up. More specifically, when we have previously been thinking about a problem or attempting to reach a certain goal, but at the moment focusing on an unrelated task or simply mind wandering, the relevance network continues to monitor mental representations that can be activated by what we perceive in our surroundings, and evaluate whether these representations can be used for the purpose of our background goals. This can sometimes lead to a creative idea in relation to our background goals, and we experience the idea as popping up in our mind even when we are not intentionally thinking about the problem.

While I agree with Carruthers that this might be one of the important mechanisms of offline generation of creative ideas, his account remains incomplete: this account cannot explain the difference between why some people are creative, and others are merely scatter minded. As Carruthers acknowledges, this process underlies many mental transitions to background goals that do not necessarily involve creative solutions.

> Notice that the stochastic processes appealed to here are quite general, and underlie many forms of *non*-creative mentation as well. For instance, seeing a face on the metro that vaguely resembles that of a colleague might result in one remembering that one had promised to send that colleague an email about a prospective student. [...] But neither is the result in any way new or surprising. What is distinctive of *creative* idea-generation is that associatively-evoked representations emerge into consciousness as providing a novel solution to some existing problem, goal, or need (Carruthers 2020).

We can imagine a scatter minded person who is constantly interrupted by intrusive thoughts about their background goals even when they are doing unrelated tasks, but these thoughts might be extremely mundane, e.g. “I need to send an email”, or “I need to pick up my package”. The scatter minded person can have thoughts that are useful for their background goals, but are by no means unusual. Therefore, to generate creative thoughts offline, other processes need to be involved. Indeed, my account identifies the offline generation of memory gists as an important process involved in offline creative thinking, which will to some extent explain what makes the creative person different from the scatter minded one. As I will elaborate on in Section 4, offline generation of memory gists involves the process of memory consolidation. Memory consolidation is a process that humans go through during both sleep and wakeful rests (Lewis & Durrant, 2011; Robin & Moscovitch, 2017; Rångtell et al., 2017), where our memory system highlights underlying commonalities between contents that sometimes appear to be very different on the surface. This process thereby serves as a mechanism for offline generation of *creative* ideas.

Another account of offline creativity is proposed by Dustin Stokes (2007), which focuses on the role of incubation in broadening the scope of ideas that are attended to during creative idea generation. Stokes appeals to studies on “incubation effects” that show that subjects who take a break from a creative task and perform an unrelated task show greater success on the creative task than those who did not perform an unrelated task (Smith & Blankenship, 1989, 1991). Stokes argues that during the incubation period of creative idea generation, performing unrelated tasks allows subjects to weaken their focus on particular ideas that might not lead to a promising answer. Meanwhile, they also activate a wider range of connections, such that the new connections can emerge into consciousness quickly when the subject returns to the problem. The resulting new ideas might be more novel or useful than the one that the subject was fixating on before. My response to Stokes' account is similar to my response to Carruthers' in that I think it is incomplete. Broadening the scope of attention is sometimes helpful for generating creative ideas, but in most cases, certainly often not enough. To generate creative ideas, we need ideas that are not just unusual but also useful. More creative people do not blindly generate unusual ideas, but are instead specifically skilled at activating those mental representations that satisfy both conditions for creativity.[25] As I will explain later, their ability to generate memory gist offline provides just this explanation that we need.

2.3. Summary

To sum up my discussion, previous accounts of both online and offline generation of creative ideas in philosophy have trouble providing complete explanations for individual differences in generating creative ideas, and this is at least partially due to the fact that they have overlooked the role of memory gists in generating creative ideas. My own positive account highlights the role of memory gists in both online and offline generation of creative ideas, but it does not purport to show that it is the *only* mechanism that facilitates the generation of creative ideas. The other mechanisms identified by these aforementioned accounts likely play a role as well. My claim is simply that these previous accounts also need to pay attention to the role of memory gists in order to provide a complete explanation for the question of individual differences in creative idea generation. In the next two sections (section 3 and 4), I introduce my account of online (section 3) and offline (section 4) gist generation in turn.

3. Online gist generation through internal diffuse attention

I will now provide my account of how the generation and use of memory gists in memory retrieval facilitate the online generation of creative ideas. I will start with an illustration of how the use of memory gists enables creative people to generate ideas that are both unusual and also appropriate for

[25] It is worth noting that the Smith and Blankenship studies that Stokes (2007) draws on specifically focuses on the role of broadening the scope of activation for overcoming “mental block” or fixation of ideas that lead to dead ends. In these cases, certainly thinking about anything else will be more useful than what the agent was fixating on before, but once the subject does overcome the unhelpful fixation, they then again need a strategy to activate not just any contents, but those that fit the conditions of being unusual and also useful.

the task, using metaphor generation as an example. I then argue that online generation of memory gists involves the mode of attention that I call "internal diffuse attention".

While there is relatively little literature on how we generate metaphors, metaphor generation likely shares at least some underlying processes as metaphor comprehension, which has been relatively well-studied. As indicated by theories of memory comprehension, metaphors often involve highlighting underlying commonalities between contents that are not often associated with each other, such as by showing that their superficial features have some similarities or mappings on a more abstract level (Gentner, 1983). For example, understanding the metaphor "Juliet is the sun" requires understanding the similarity between Juliet's specific, surface-level features, e.g. being attractive or a desirable companion to Romeo, and the sun's specific, surface-level features, e.g. providing heat and light to beings that depend on it.[26] One way to understand the mapping is to find features that both Juliet and Romeo have on a more abstract level, e.g. being nurturing and life-sustaining. Similarly, coming up with creative metaphors might often involve extracting more abstract features from the content that one is describing, and then finding another content that seems on the surface unrelated to what is being described, but shares this more abstract feature in the context being considered. For example, to generate the metaphor "Juliet is the sun", one would need to extract the more abstract feature of being life sustaining or extremely valuable from Juliet's specific features, and finding another concept that shares this feature with Juliet, i.e. the sun.

The notion of "memory gist" provides us a way to understand this process: we can understand the more abstract feature that we extract from the initial content's specific features as the gist of this content in the context we are currently considering. This gist memory enables us to then retrieve concepts that are, on the surface, not alike the one we are describing, but share the features

[26] More specifically, according to Gentner's structure-mapping theory, to understand a metaphor, one would need to find similarities between the target and the vehicle based on their shared underlying structure with concepts that they are related to, i.e. to understand "Juliet is the sun", the reader would need to understand the similarity between Juliet's relationship to Romeo, and the sun's relationship to living beings that depend on it. This is in line with my suggestion here.

represented in the gist. In the case of “Juliet is the sun”, the memory gist is the more abstract feature of being life sustaining or extremely valuable, which facilitates the retrieval of another content (i.e. the sun) with specific features that are on the surface very unlike Juliet’s (e.g. the former being a hot and gigantic planet) but share the same abstract features in their gist in this context.

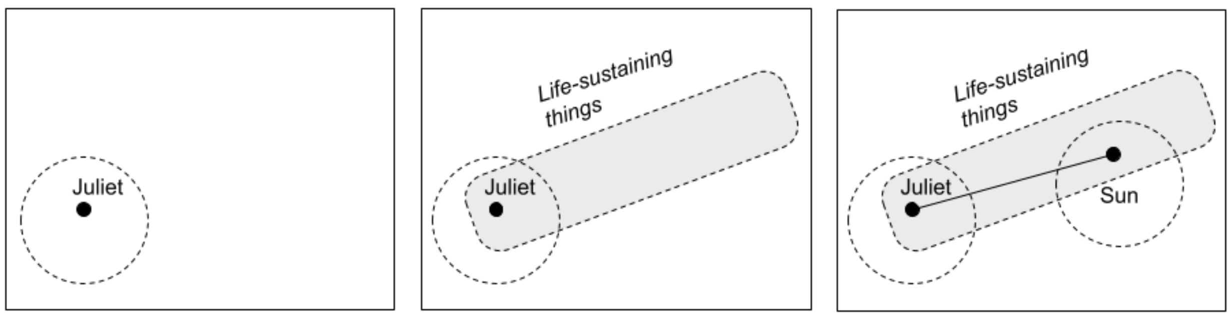

Figure 2.1: Using abstract features represented in gist to guide memory search.

Through the memory gists that are formed on the fly during retrieval, creative individuals can then activate concepts that are semantically distant, but share the same gist in this context. On the one hand, the gist provides these individuals with a shortcut through which they can leap out of the group of contents that are often associated with the content that they initially started with (i.e. Juliet). The abstract nature of the features represented in the gist therefore enables creative individuals to activate distant contents that lead to unusual ideas. This is because contents that share abstract features can be very different in their surface-level features. My account can therefore explain the result from Kenett et al. (2016) where more creative individuals tend not to get stuck within the same subcategory of animals in the verbal fluency task. They are better at generating unusual contents thanks to the fact that they can switch from thinking about the specific features to the more abstract features, which guides their memory search into more distant contents. On the other hand, since the activated contents are not just *any* random contents, but are those that fit the gist, these contents also satisfy the usefulness condition, because the gist is generated in accordance to the context of the task. This explains why a creative person would generate the metaphor “Juliet is the sun”, but not “Juliet is a raccoon” or “Juliet is Manchester”, because “raccoon” and “Manchester” are simply not activated, as they do not fit the gist.

How then do we understand the process of online memory gist generation? To answer this question, I will now introduce the notion of “internal diffuse attention”, a mode of attention akin to the type of diffuse attention that some philosophers and psychologists have identified in the case of perception, but directed internally instead of externally.

3.1. Diffuse attention: internal and external

For several decades, psychologists have proposed that there is a distributed mode of attention that is distinct from focused attention (Srinivasan et al. 2009; Oliva & Torralba, 2006; Chong & Treisman, 2005). This mode of attention has been conceptualized recently in philosophy as diffuse attention (Prettyman, 2023). It has a broader scope than externally directed focused attention, kicks in faster, and is better suited for processing more holistic features. We employ diffuse attention in perception to process global and statistical properties of a group of objects, such as their mean sizes or color (Albrecht & Scholl 2010; Chong & Treisman, 2005), as well as what psychologists call the “gist” of a scene (Oliva & Torralba, 2006), i.e. the holistic, identifying features of a scene that do not contain the particular details and local objects.[27]

While previous research in philosophy and psychology has primarily conceptualized diffuse attention as directed externally in perception, I propose that we understand diffuse attention such that

[27] I do not argue that the “gist” in scene perception and memory gists are fundamentally the same kind of content, but the former is potentially a subcategory of the latter. “Gist” in scene perception usually refers to global properties of *one* particular scene, but memory gists can also be extracted from different memory episodes or contents.

it can be directed not only externally in perception, but also internally towards contents in our memory.[28]

While the targets of external diffuse attention in perception are often represented in an imagistic format, the target of internal diffuse attention can sometimes be thoughts or memory contents that have a discursive format (such as when you are coming up with a creative metaphor).[29] In cases where the targets of your internal diffuse attention are thoughts or memory representations that have a discursive format, internal diffuse attention is directed towards a wide range of contents in virtue of the fact that they share the same underlying abstract features.[30] This mode of internal attention is diffuse in the sense that it is directed towards a range of contents that are distant from each other in the conceptual space. Conceptual space represents contents as points in a continuous geometric space, where the similarity relationship between contents is modeled as the distance between points in the geometric space (Bellmund et al., 2018; Tenenbaum & Griffiths, 2001; Gärdenfors, 2000). Contents that are less similar are therefore represented as more distant from each other in the geometric space. In the above case of metaphor generation, for example, generating Juliet’s gist in this specific context, i.e. being nurturing and life-sustaining, therefore enables the creative person to attend diffusely to a wide range of contents across the conceptual space that share this abstract feature, including not only Juliet but also the sun, shelter, food, etc.

[28] The existing work on internal attention in philosophy mostly focuses on the role of internal attention towards contents that are already conscious in cognition or working memory (Fortney, 2018; Phillips, 2012). Here, I adopt a broader conception of internal attention such that the contents that are attended to can also include contents in long term memory that are not conscious prior to being attended to. This broader conception of internal attention, where the pre-attended content can be unconscious long-term memory, is widely accepted in psychology (Dixon et al., 2014; Chun et al., 2011). This conception is also compatible with the view that attention is sufficient for consciousness (Cohen et al. 2012), if we accept that the nonconscious contents become conscious through being attended to: attention boosts these contents into consciousness (De Brigard, 2012; Posner, 1994).
[29] My view is also compatible with the possibility that there is in fact a spectrum from purely iconic to purely symbolic representations (Greenberg, 2023).
[30] I will use “feature” and “property” interchangeably throughout this paper. Following Tenenbaum and Griffith (2001), I will use “feature” in most of my discussions with respect to the conceptual space. However, Prettyman (2023) uses “property” in her discussion of external diffuse attention, so I will use “property” in discussion of her work.

My account also does not claim that diffuse attention is the only mode of attention involved in online gist generation. Instead, many studies in psychology have shown that creative thinking requires not just the ability to attend to a wide range of contents in a distributed manner, but also the ability to switch between this more spread out attention and more focused attention. Studies have also shown that creative people are better at switching between attending to local and global objects in visual perception (Zabelina et al., 2016), indicating that the ability to switch flexibly between focused and diffuse attention plays a role in creative thinking (Zabelina, 2018; Gabora, 2018).

Therefore, my account acknowledges that creative thinking requires not just the ability to attend diffusely to internal contents, but also the ability to switch flexibly between focused and diffuse attention. In the case of generating a creative metaphor like "Juliet is the sun", the creative person would need to switch from using focused attention to attend to Juliet, to using diffuse attention to attend to a wide range of contents that share the same abstract feature represented in the gist, i.e. being nurturing and life-sustaining. These contents might include the sun, shelter, food, etc. The creative person will then again switch back to focused attention to individually select some of these contents and compare them with Juliet to further refine the metaphor.

By arguing that there is an internal analog to externally directed diffuse attention in perception, I enrich the taxonomy of different modes of attention, and call attention to this diffuse mode of internal attention. Much like its external analog, diffuse internal attention is particularly good at selecting a broad range of contents in virtue of the fact that they share certain underlying features. This distinctive processing benefit of diffuse internal attention therefore makes it particularly beneficial for certain modes of reasoning that rely on selecting disparate contents with underlying commonalities, including creative thinking. My account also allows me to provide a framework that helps us explain the interaction between internal and external attention. Studies have shown that perceptual tasks that require a defocused mode of attention improves people's ability to generate creative ideas (Wegbreit et al., 2012; Friedman et al., 2003). In one study (Friedman et al., 2003), subjects who were primed to attend to broad visual targets subsequently performed better in some creative thinking tasks, including a divergent thinking task that required them to come up with unusual uses of a brick. Given my framework, one explanation would be that the perceptual task likely requires external diffuse attention, which then induces internal diffuse attention as well, which facilitates better performance in the creativity task.

4. Offline gist generation through memory consolidation

I now turn to a discussion of the role of memory gists in offline creative thinking. The mechanism for offline gist generation differs from that of online gist generation. In particular, I will argue that memory consolidation during sleep and wakeful rest plays an important role in generating memory gists. In contrast to my account of online gist generation, I will not argue that diffuse attention plays a role in offline gist generation. This is because many in the attention literature in philosophy adopt the view that attention is sufficient for consciousness.[31] Given that some of the underlying mechanisms of offline creativity only generate non-occurrent mental representations, it will be highly controversial to argue that diffuse attention is involved in offline cases. However, I will still argue that diffuse attention plays a role in the *retrieval* of memory gist when the agent returns to thinking about the problem.

My emphasis on the role of memory consolidation is supported by the existing evidence that memory consolidation contributes to the generation of more abstract properties that are represented in memory gists. Memory consolidation is the process where new memory that is prone to decay gets stabilized and integrated into long term memory through active replay during sleep (Fischer et al., 2006; Drosopoulos et al., 2007; Diekelmann et al., 2009) or waking rests (Squire et al., 2015; Wamsley, 2022). During this process, newly encoded memory contents can be processed and

[31] See Cohen et al. (2012) for a discussion of this view about attention and consciousness.

reorganized, and overlapping features of different contents are enhanced while some of their individual differences are forgotten. Through memory consolidation, our memory system thereby builds new connections between contents based on their shared commonalities (Lewis & Durrant, 2011; Robin & Moscovitch, 2017; Rångtell et al., 2017).

Below, I will present a study that shows that the extraction of memory gist during sleep facilitates the emergence of a creative insight. In this task, gist extraction generates beliefs about common features shared across different contents, and subjects were not aware of such commonalities prior to the process of gist extraction. Moreover, the common, abstract feature represented in the gist directly supports the emergence of a creative insight in solving the task. I will then discuss the role of attention in cases where gists are extracted offline.

4.1. Gist extraction during memory consolidation and creative insight

A study by Wagner et al. (2004) shows that cognitive processes that happen during sleep play an important role in generating a creative insight. In the experiment, subjects are asked to perform a number reduction task (NRT). In the task, subjects need to process several strings that are written in a ternary numeral system containing 1,4, and 9. The letters are processed through the following rules: (1) two identical digits produce the same digit (i.e. 1 and 1 result in 1); (2) two non-identical digits produce the third remaining digit (i.e. 1 and 4 result in 9). Subjects are asked to calculate the final result for each string as fast as possible.

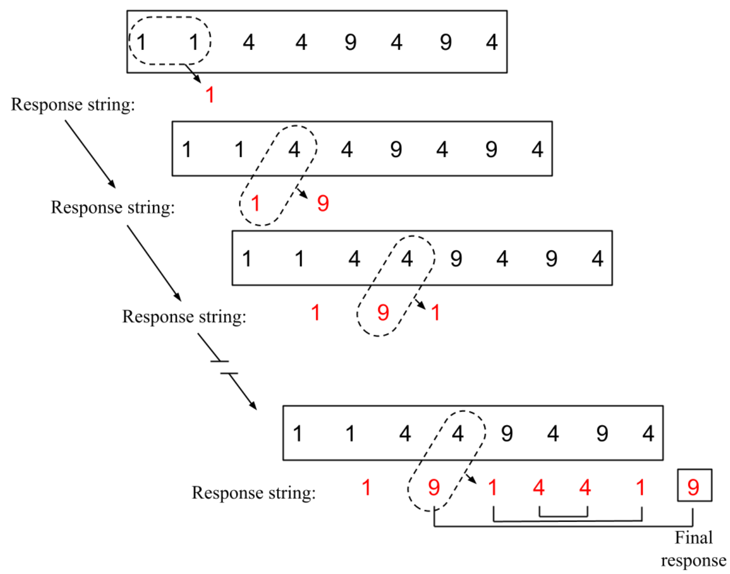

Fig. 2.2: Illustration of the number reduction task from Wagner et al. (2004)

While not announced to the subjects, there is a global pattern underlying the different strings: all the strings are produced such that the last three numbers ("419" in the example) in the response string mirror the previous three responses ("914" in the example). Once subjects notice this pattern by themselves, they can just report the second letter in the response string as the final answer, which fastens the calculation dramatically.

After performing the task, subjects are divided into a "sleep" group that gets retested after a night's sleep, and two "no-sleep" groups, one of which gets tested after staying awake throughout the night, and another gets tested the same day but after a long interval.[32] Results show that the "sleep" group is much more likely to notice this pattern after returning to the task compared to both "no-sleep" groups, and because of this, they as a group solve the task much faster when they return to it compared to the "no-sleep" groups.

[32] Note that participants in the "no-sleep" group that stays awake during the day and gets retested later in the day self-reports similar levels of fatigue as the "sleep" group during the retest, which rules out the effect of fatigue on their retest performance.

As I have mentioned before, memory consolidation during sleep can generate gist contents, and while memory consolidation happens both during sleep and wakeful rests, there is some evidence that sleep is especially effective in consolidating certain types of memory contents that involve associative relations (Studte et al., 2015). Therefore, a plausible explanation of the result from Wagner et al. (2004) is that in this case, memory consolidation during sleep generates memory gist contents that facilitates the emergence of a creative insight into how to solve the problem when the subjects are retested. In this case, the gist contains representation of a more general feature that is shared by the different contents in the representations that were initially encoded. In particular, the different number lists that the subjects generated for each problem share the more general feature of having mirroring structure between the last three numbers and the three numbers before them. It was difficult for subjects to come to recognize this fact prior to going to sleep because the number lists look very different on the surface, but the process of gist extraction enables them to recognize the underlying pattern, and therefore leads to a creative insight in solving the problem.

In addition to the Wagner et al. (2004) study, another more recent study shows that gist generation during memory encoding might support creative capacities. Thakral et al. (2021) show that individuals who tend to generate false memories in the Deese-Roediger-McDermott (DRM) paradigm perform significantly better at several different creativity tasks compared to those who are more accurate in their word recognition in the DRM paradigm. Given that false memories in the DRM paradigm is likely caused by gist generation, this study provides some evidence that gist generation in the memory system supports creative thinking. While in this case, in the DRM paradigm, the memory gists are likely generated at the time of encoding instead of during consolidation, the result stills lends some plausibility to the role that gist generation within the memory system without intentional intervention from conscious cognition (whether at the time of encoding or during consolidation) facilitates the generation of creative ideas.

Finally, my account does not purport to explain *every* case of offline creativity. It is certainly plausible that certain cases of offline creativity merely require activation and recombination of highly specific contents, without generating more abstract gist contents. For example, a study by Cai et al. (2009) shows that rapid eye movement (REM) sleep improves individuals' ability to creatively use contents they read in an unrelated task as solutions in a word puzzle. In this case, REM sleep facilitates generation of a creative solution later on simply by reorganizing and integrating information acquired in different contexts, without generating more abstract gist contents. I acknowledge that there are in fact many cases of offline creativity like this that do not involve gist generation through memory consolidation or diffuse attention in retrieval. However, my point is simply that, still, recognition of task-appropriate connections between distant contents often rely on abstract gist contents, and therefore my account identifies at least *one* of the crucial mechanisms for generating creative ideas.

4.2. Use of memory gists generated offline in memory retrieval

While diffuse attention does not play a direct role in the generation stage of memory gists in cases where gist generation happens offline, it is likely still involved at least sometimes in the retrieval of memory contents guided by memory gists in solving the problem. Imagine if, in the Wagner et al. (2004) study, someone does in fact form the appropriate gist representation of the feature shared by the different number strings during sleep. However, when they are later retested, they fixate on the specific numbers in each string, or fixate on trying to recall how they solved each individual problem last time. Due to narrowly focusing their attention (internally and externally), they fail to recall that the various different strings all share the same feature. In this case, despite having generated the appropriate gist representation offline, the individual fails to use it to guide their memory retrieval when they return to the problem due to not using diffuse attention.

I also do not argue that diffuse attention is *always* involved in memory retrieval when the memory gists are generated offline. In fact, sometimes individuals experience a highly specific idea as "popping out" either when they return to intentionally thinking about the problem, or when they are engaging in unrelated tasks. In this type of case, a wide range of contents might have been partially activated during the offline process thanks to the fact that they all fit the same gist. However, only one of them is boosted into conscious cognition through focused attention (De Brigard, 2012).[33]

5. Discussion

To sum up, I have argued that memory gists play crucial roles in many cases of online as well as offline generation of creative ideas. Memory gists enable creative individuals to connect conceptually distant contents in their long term memory that share the same underlying abstract features represented in the gist. The fact that they all fit the gist also ensures that the generated idea is not only unusual but also useful. I have identified different mechanisms that are crucial for online and offline gist generation respectively. Internal diffuse attention is crucial for online gist generation. In this case, gists are generated on the fly during memory retrieval. Meanwhile, memory consolidation is crucial for offline gist generation. However, when gist is generated offline through memory consolidation, diffuse attention is sometimes still involved in memory retrieval guided by the gist.

I will now turn to discuss several accounts in psychology and neuroscience that share some of the commonalities with mine. I then explain how my account differs from them in terms of the goal of my account. This is followed by a discussion of the philosophical implications of my account for how we understand attention and memory.

---

[33] I thank an anonymous referee for prompting me to think about this issue.

Several accounts in the psychology literature highlight the role of flexible use of executive control in cognition and memory retrieval in creative thinking (Beaty & Silvia, 2012; Beaty et al., 2014; Beaty & Schacter, 2018; Benedek & Neubauer, 2013; Benedek et al., 2017; Frith et al., 2021). In addition to arguing that cognitive control plays a role in inhibiting early responses (Beaty & Silvia, 2012) as Carruthers (2020) has already identified, studies also show that cognitive control is involved to guide memory search by ensuring that the agent is generating concepts in the categories that are appropriate for the task, and inhibits when the agent is attending to information irrelevant or not appropriate for the task (Beaty et al., 2014; Frith et al., 2021).

Let me start with an explanation of how the executive control theories are compatible with my account, before I explain how my account differs from both the executive control theories. In online creative thinking, I do not take diffuse attention to only involve bottom-up processes. Instead, acknowledge that top-down control also plays a crucial role in generating as well as maintaining the appropriate gists throughout memory search, as well as in switching between diffuse and focused attention.[34] Moreover, given the existing evidence that executive control can also be involved in non-conscious processing, the executive control accounts can provide some explanation for offline creative thinking as well (Fogel et al., 2022; Samanta et al., 2021). My account does not rule out that executive control might be involved in gist generation in memory consolidation, or that there are offline processes that contribute to creative thinking that do not involve gist, since I do not claim to give an exhaustive account of all cases of creative thinking.

What differentiates my philosophical account from the group of theories in psychology that emphasize the role of executive control is a difference in our explanatory goals. Beyond pointing out a gap in the philosophical literature on creativity, my account also makes distinctive contributions to

[34] In fact, diffuse attention in creative thinking is likely often a result of co-operation between bottom-up and top-bottom attentional mechanisms. As evidenced by recent studies, the default mode network that underpinned internal bottom-up attention and the frontoparietal control network (FPCN) that underpins top-bottom control can be activated at the same time and work cooperatively (Andrews-Hanna et al., 2014; Christoff et al. 2016).

the understanding of the nature of attention and memory, as well as their roles in our cognitive economy. These goals go beyond the contribution of the psychology theories that focus on explaining the mechanisms of creativity. First, by introducing the notion of "internal diffuse attention" and highlighting its role especially in online creative thinking, my theory contributes to the philosophical understanding of the nature of attention. By arguing that there is an internal analog to externally directed diffuse attention in perception, I enrich the taxonomy of different modes of attention, and call attention to this diffuse mode of internal attention. With a few exceptions (De Brigard, 2012; Phillips, 2012; Fortney, 2018), internal attention in general remains not very well theorized in philosophy, and current discussion of internal attention has been clustered around focused attention. My account improves our understanding of internal attention by distinguishing between diffuse and focused attention, and by understanding internal attention in parallel with external attention in perception. More future theorizing is needed to understand how to conceptualize internal diffuse attention, as well as how to understand its relationship with external diffuse attention.

My account also suggests that diffuse attention, internal or external, comes with its distinctive processing benefit. They are particularly good at processing more global targets, which make them suitable for certain tasks, such as creative thinking, and scene recognition in perception. This potentially puts pressure on the traditional discussion of the function of attention that only focuses on the role of focused attention.[35] My account thereby goes beyond the accounts of creativity in terms of executive control in psychology by providing additional theorizing about the nature of attention.

Second, while executive control might also explain some cases of offline creative thinking to the extent that executive control can be involved in offline processes, accounts based on executive control differ from the focus of my account of offline creative thinking, which focuses on the role of memory. In fact, by highlighting the role of memory consolidation for offline creative thinking, my account also draws out some important implications for the role of memory in our epistemic lives.

[35] See for example, Watzle (2023).

Traditionally, philosophers have taken memory to take a relatively passive role of accurately preserving previously acquired information (Goldman, 2009). However, more recently philosophers have attributed a more active function to memory. In particular, several have argued that memory makes changes to previously encoded information, or generates new information on its own, not for the purpose of preserving previously acquired information, but rather for the purpose of providing useful information for future reasoning.

For example, according to a view by De Brigard (2014) that focuses on episodic memory, in addition to (or sometimes instead of) functioning to accurately preserve previous experiences, episodic memory also functions to simulate possible past events (i.e. what could have happened) in the service of simulating possible future events. These simulations of possible future events can then be used in decision making. Another view by Aronowitz (2019) argues that memory contributes to future reasoning by building models based on encoded information, which involves reorganizing and simplifying previously acquired information. These models can then facilitate future reasoning and decision making.

By highlighting the role of memory consolidation in offline creative thinking, my account thereby provides another case for the active role of memory in our cognitive economy, but specifically focuses on the active role of memory in creative thinking, which the philosophy of memory literature has not considered. Gist generation through memory consolidation provides creative individuals with crucial information that they can then use for generating creative insight. My account therefore suggests that at least one of the functions of the memory system includes generating abstract or qualitative contents that are especially beneficial for the types of reasoning that heavily rely on these contents, including but not limited to many cases of creative thinking.

Chapter 3: Memory as a generative source of understanding

Abstract

What is the function of memory? In this paper, I reject the traditional view that memory merely functions to preserve previously acquired information, such as information acquired through perception. I argue instead that one of the functions of memory is to improve our understanding of subjects and systems represented in the contents that we previously acquired. This is possible thanks to the fact that the memory system further processes previously acquired information, especially during the memory consolidation process. Through this process, the memory system improves the agent's performance in forming new beliefs and other representations based on the information that they already have. More specifically, the consolidation process systematically generates the specific type of representations that contribute to understanding: these representations make accessible underlying relationships between different components of the represented subject or system and do so in a way that is sensitive to the agent's goals. This paper therefore contributes to the ongoing project of understanding how memory generates new epistemic values, but departs from those that focus on epistemic justification and expands the scope of epistemic values under discussion by including and foregrounding understanding.

1. Gaining understanding without conscious thinking: a mystery?

In everyday life, we often hear anecdotes of scientists, researchers, or ordinary people emerging from sleep with a sudden insight that provides them with a newly developed or improved understanding of a subject matter, or a problem that they have been pondering about during the day. They suddenly grasp some hidden connections between things that they thought were unrelated, or suddenly come to see an underlying pattern that was hiding underneath the superficial features they observed initially. For example, consider the following scenario:

> **A detective with fresh eyes** Dave the detective is working on a murder case and trying to identify who the murderer might be based on all the evidence he has been given. He examines all the evidence that the police collected again and again, and tries to think through what it might indicate, but it is just not “clicking”. After he went to sleep and woke up the second day, Dave looked at the evidence again with fresh eyes, and suddenly everything clicked! Dave came to the realization that a piece of evidence about Ethan, the family Doctor, that he thought was trivial is actually a crucial clue, and he now started to see its connection with the other evidence. Then everything falls into place. “The murderer must be Ethan!” Dave shouts to himself.

Experiences like this are often mysterious: the agents who are struck with an insight after sleep did not engage in conscious thinking during sleep, but somehow they are able to come up with an insight, and their understanding of the subject has thereby improved. What are the mysterious nonconscious processes that enable us to improve our understanding of something without consciously thinking about it? Psychology and neuroscience studies on memory have in fact demystified these processes to some extent by identifying the important role of sleep for the active consolidation of memory contents (Klinzing et al., 2019). Memory consolidation is the process where new memory that is prone to decay gets stabilized and integrated into long term memory through active replay during sleep or sometimes also during waking rests (Squire et al., 2015; Wamsley, 2022). During this process, some parts of the encoded memory contents are highlighted while others are forgotten, and the memory contents also go through further processing where the memory system recognizes inferential or associative relations between disparate contents, as well as extracts general patterns from detailed information encoded in the memory contents.

Apart from a handful of exceptions (Aronowitz, 2019; Miyazono & Tooming, 2023), few in the philosophical literature have considered the epistemic role of memory consolidation. The focus of this paper is therefore on the role of memory consolidation in improving our understanding. I will argue that one of the functions of memory is to improve our understanding, thanks to the memory consolidation process. Memory consolidation improves our understanding through processes such as integrating piecemeal information to recognize distant connections, as well as extracting general patterns from disparate contents. These processes improve our understanding by generating structured representations of relationships that are central to our epistemic goals into our consciousness.

This view thereby allows me to provide a new argument for the view that memory plays an active role in our cognitive economy that is importantly different from the current arguments for this view. A traditional view of memory is that it serves the function of storing and preserving previously observed information (Burge, 1997; Goldman, 2009). Let us call this the "storehouse" view of memory. More recently, many philosophers have taken seriously the more active role of memory demonstrated by empirical research and rejected the "storehouse" view. They instead argue that memory plays an active role in generating new epistemic values, especially new epistemic justification (Lackey 2005; Michaelian 2011; Fernández, 2016; Bernecker & Grundmann, 2019).

The focus of my account, however, is not on epistemic justification but on understanding. While I am open to the possibility that the memory consolidation process that I argue to improve understanding also generates other epistemic values, including justification. My account will mostly focus on understanding for the following two reasons: first, the previous literature that focuses on justification, and especially on propositional justification, overlooks an important way in which memory can be epistemically generative, which is through further processing previously acquired information, instead of incorporating new information that was not previously available to normal cognition. Second, understanding better captures an important function of memory, whereas justification (or some other epistemic values, such as knowledge) sets the target too wide. Memory functions to improve understanding, and that explains why memory also generates justifications for the specific kind of beliefs that improve understanding, instead of justification for just any random beliefs.

Below in section 2, I identify several important features of understanding, as well as how representations can provide or improve our understanding through making accessible some underlying relationships that are crucial given our cognitive and practical goal. Then in section 3, I introduce what memory consolidation is as well as its underlying mechanisms. In section 4, I explain how these mechanisms enable memory consolidation to improve the investigator's understanding. In particular, I identify two processes that enable memory consolidation to generate structured representations that highlight relationships between different components: the process of extracting general patterns, as well as integration of piecemeal information. I also provide evidence that shows the memory consolidation generates representations in a way that is sensitive to the agent's goals. In section 5, I discuss the implication of my account for the current literature on the role of memory in generating epistemic values. In particular, I argue that my account highlights a way in which memory generates new epistemic values that has been overlooked by the current literature. While current accounts focus on memory generating new epistemic values through making use of information not accessible to normal cognition, my account focuses on memory generating new epistemic values through further processing information that was already accessible to normal cognition, and this process does not always require the assistance of encapsulated information. I then provide an analysis of why previous accounts have overlooked this other way in which memory generates new epistemic values: this is likely because these accounts narrowly focus on propositional justification. My account instead expands the scope of epistemic values under consideration and highlights understanding as one of them.

2. Understanding and the function of memory

What is understanding? The understanding literature typically characterizes understanding as requiring the agent to see how a wide range of things can fit together (Riggs, 2003; Grimm, 2011). It requires “grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information” (Kvanvig 2003, 192). Understanding neuroscience requires understanding how different networks in the brain work with each other. Understanding how an engine works requires understanding how different parts of the engine work together. Understanding is therefore often *structured* in the following sense: it requires more than simply remembering isolated pieces of information, but also seeing how different parts fit together in virtue of the relationships between them.

Another feature of understanding is that the object of understanding depends on the context of the agent’s cognitive or practical goal (Bengson, 2015; Potochnik, 2020). As a result, the object of understanding can even vary for two agents who are investigating the same subject matter, but with different goals in mind. For example, Bengson argues that understanding in different contexts requires being aware of features that are central in the given context. Bengson writes:

> For example, what is central in a culinary context need not be identical to what is central in a botanical context, since the former context calls for all and only what is needed to characterize what the target is botanically, whereas the latter context calls for all and only what is needed to characterize what the target is culinarily (Bengson, 2015, 20)

In other words, if we were to study basil in a culinary context, then features such as its aroma and flavor are central to our purpose; if were to study basil in a botanical context, then features such as how much watering it needs and what temperature it grows best in are central to our purpose. A plausible explanation for this difference in the recognized features required for understanding in the two different contexts is that in the culinary context, the object of understanding are the culinary features of basil, while in the botanical context, the object of understanding are its attributes as a plant. Let us call this feature of understanding "goal-sensitivity".

I will take being structured and sensitive to the agent's goals to be two central features of understanding. Admittedly, understanding applies to a wide range of things in science, math, arts, as well in our daily life. We can understand a person, a language, a style of dance, a field in mathematics or science, etc, and there might be a diversity of different kinds of understanding towards different objects (Baumberger et al., 2016). I do not take being structured and sensitive to the agent's goals to apply to every kind of understanding. However, it suffices for my argument to claim that these two features are often taken to characterize notions of understanding that are often used in our daily life and the sciences, among other domains. Moreover, I will argue that our memory system improves the kinds of understanding that are captured by these two features, even if it does not improve other kinds of understanding.

How do we acquire or improve our understanding of a subject or a system? This is often possible through our acceptance of a representation of the subject or system. The representation can take the form of a graph, a schema, a model, a theory, etc (Elgin, 2016). A representation can provide investigators with improved understanding of the subject that it represents by highlighting or making accessible underlying relationships between different components of the subject or system, and also select and prioritize relationships that are crucial given the background goal of the investigator.

For instance, Elgin argues that scientific theories improve our understanding of scientific principles by providing us with idealized representations of real world scenarios. For example, the ideal-gas law provides scientists with understanding of the behavior of gas under many conditions by positing the concept of ideal gas, i.e. a body of gas composed by randomly moving particles that do not interact with each other. By omitting the effect of compounding factors, such as interparticle interactions, the ideal-gas law highlights the relationship between the pressure, volume and temperature of a body of gas, when the effect of interparticle interactions on these properties of the gas is insignificant. Scientific theories can therefore prioritize and make accessible certain relationships by providing a structured representation.

Relatedly, what representation is suitable for improving understanding also depends on the goal. For example, if our goal is to study the relationship between the pressure, volume and temperature of a body of gas under conditions with no significant interparticle interactions, then the ideal gas provides a suitable representation that can improve our understanding of the relationship between these properties of a body of gas. However, if our goal is instead to study the interparticle interactions in the same body of gas (despite the interactions being weak), then the ideal gas would not be a suitable representation, since it omits the exact features and relationships crucial for our investigation.[36]

The claim that I will make in this paper is that one of the functions of the human memory system is to improve our understanding of subjects and systems represented in our initial experience. Our memory system performs this function by generating new representations that are both structured and goal-sensitive. These new representations make accessible and highlight important underlying relationships between different components of the target, and prioritize relationships that matter for the cognitive or practical goal of the agent.

Finally, a caveat about the notion of understanding that I adopt here: there is significant disagreement in the understanding literature about whether understanding can be false. While some

[36] Elgin (2016) herself is not explicitly committed to the goal-sensitivity of the object of understanding, but others take the suitability of scientific modeling to be sensitive to the aim of the investigation (e.g. Potochnick (2020)). See Potochnik (2020) for a discussion.

theorists think that understanding only needs to be "true enough" for the purpose of one's investigation (Zagzebski, 2001; Elgin, 2009, 2017; Riggs, 2009), others think that being true is necessary for understanding (Grimm, 2006; Pritchard, 2010; Hills, 2016).[37] I will not take a side in this debate about the factivity of understanding, but it suffices for me to say that my account is compatible with both. In fact, in my discussion in Section 3, I will present empirical evidence that memory generates representations that improve *true* understanding. Having explained what I take understanding to be and how a representation can improve understanding, I will turn to clarifying what notion of function I will be working with.

2.1. The functions of memory

What does it mean to say that memory has a certain function? Philosophers draw a distinction between two different notions of function in the discussion of biological functions. The notion of *causal role function* refers to what the faculty or process does for us. More specifically, it refers to how it causally contributes to the activities of the larger biological system(s) that it is a part of (Cummins 1975, 1983). Meanwhile, the notion of *etiological function* focuses on giving an evolutionary explanation for why a certain trait is selected and retained in the organism. In other words, the etiological function of a trait refers to what the trait performed for our ancestors, such that it was selected and retained in evolution (Wright, 1973; Millikan, 1984). My account will focus on the causal role function of memory. Specifically, I will argue that the human memory system causally contributes to our larger reasoning system by improving our understanding of what is represented in our initial experiences. As I will argue below, this is possible thanks to the fact that our memory

[37] Some of those who claim that understanding is factive are also committed to the view that understanding is a kind of knowledge (e.g. Grimm, 2006; Sliwa, 2015). I will not take a stance on whether understanding is in fact a species of knowledge, but it suffices to say the cases that I will discuss in Section 3 can very plausibly be interpreted as memory generating knowledge.

system systematically generates representations with the right features such that they can improve understanding. While the etiological function of memory is not the focus of my account in this paper, I think it is also quite plausible that the trait of improving understanding demonstrated by our memory system was selected in our evolutionary history because it facilitated effective reasoning and decision making for our ancestors. In the next two sections, I will explain the underlying mechanisms of a specific stage of memory, i.e. memory consolidation, to support my claim about the causal role function of memory.

Importantly, I only argue that improving understanding is one of the many functions of memory. It is certainly not its only function, and my account is not necessarily in competition with existing accounts that argue that memory has other functions. The traditional "storehouse" view that the function of memory is to fully preserve all of the information stored in memory has been challenged by the recent literature informed by empirical research. The most prominent group of opponents to the storehouse view are the simulationists. Simulationist argue that the function, or one of the functions, of episodic memory in particular is to simulate hypothetical and counterfactual scenarios for the service of generating useful information for future reasoning (De Brigard, 2014; Michaelian, 2016; Schacter & Addis, 2007). For example, De Brigard (2014) argues that one of the functions of episodic memory is to simulate hypothetical scenarios of what could have been, which explains why we sometimes have false memories of what happened. These cases are not occurrences of malfunctioning of our memory system, but rather its normal functioning to generate information about hypothetical scenarios to facilitate our reasoning about future events.

My account is certainly not in contention with these simulationist accounts. First, while simulationists focus on the episodic reconstruction process of episodic memory, my account focuses on a very different memory process, i.e. memory consolidation, which has not received much attention in the philosophy literature. It is certainly possible that different memory processes play different functional roles. Second, it is also possible that mental simulation also generates understanding, albeit possibly through a very different way than how memory consolidation generates understanding. For example, some have argued that moral understanding can be improved by first-person experience of relevant situations (e.g. first-personal experiences of being harmed by implicit bias can provide one with understanding of implicit bias) (Sliwa, 2017). Mental simulation might then preserve and highlight these first-person experiences and thereby provide understanding. Since the type of understanding, and how understanding is generated in this process is potentially very different from how memory consolidation improves understanding, I will not explore this line of thought here. However, it suffices for me to say that if mental simulation does improve understanding, then the function of memory highlighted by simulationists is not only compatible with my account, but in fact also supports my account that improving understanding is a function of memory.

Another view that challenges the "storehouse" view from a different angle is from Aronowitz (2019), who is one of the few in the philosophy literature that does in fact take memory consolidation seriously. Aronowitz argues that the function of memory is to facilitate optimal retrieval in future reasoning and decision making. This is possible thanks to the fact that through memory consolidation, our memory system builds simplified and goal-sensitive representations based on the contents initially encoded in memory, akin to how scientific models are built based on data.

My account is in agreement with Aronowitz's account that the representations generated by memory consolidation are akin to scientific models, but I go beyond Aronowitz's account in further fleshing out the epistemically generative role of memory. In particular, I am interested in identifying what kind of epistemic value memory generates, which I argue to be understanding. I agree with Aronowitz, however, insofar as I think that the analogy with scientific modeling provides a fruitful way to help us understand how memory consolidation generates understanding.

3. What is memory consolidation

In this section, I will introduce memory consolidation, the memory process between memory encoding and retrieval. Empirical evidence has demonstrated that our retrieved memory content is often different from our initial experience, and that the generation of contents distinct from the initial experience can happen at the time of memory encoding, memory consolidation, or retrieval (Intraub & Richardson, 1989; Frankland & Bontempi, 2005; Alba & Hasher 1983; Roediger & McDermott, 1995). However, for the purpose of my argument, I will focus on the generation of contents different from the initial experience in the stage of memory consolidation in particular, and explain the cognitive and neurobiological mechanisms that underlie memory consolidation.

Let us start with an intuitive example: if someone asks you what do kids' birthday parties that you have been to typically look like, you would probably say that kids' birthday parties typically involve a cake, candles, balloons, gifts, and loud kids. This more general representation is formed based on your experience at specific birthday parties. However, you probably would not be able to provide more detailed information of these specific birthday parties, such as whether the cake at a specific party was a chocolate cake, and what gifts the birthday kid received at another party. How, then, do you form the more general representations about kids' birthday parties that you have been to? This is possible thanks to the memory consolidation, during which our memory system generates new representations through further processing previously acquired information.

Studies have shown that during sleep (and especially in the stages of slow-wave sleep and rapid-eye movement (REM) sleep) and waking rests, representations of our recent memories in the hippocampus and the neocortex are reactivated. Through replay between the hippocampus and the neocortex, our memory representations gradually transfer from the hippocampus to the neocortex, and the representation in the neocortex also becomes more integrated into the existing cognitive schema (Born, 2010; Sirota & Buzsaki, 2005; Klinzing et al., 2019). Through transferring the representation from the hippocampus to the neocortex, the memory representation becomes more independent of the episodic context, as the context details are forgotten. The content of the representation also changes in accordance with the synaptic changes within the neocortex. These changes include strengthening some synaptic connections, while weakening others, and also building new connections. The resulting memory representations after consolidation are therefore "more-general, schema-like knowledge representations that are well embedded in pre-existing knowledge networks" (Klinzing et al., 2019).

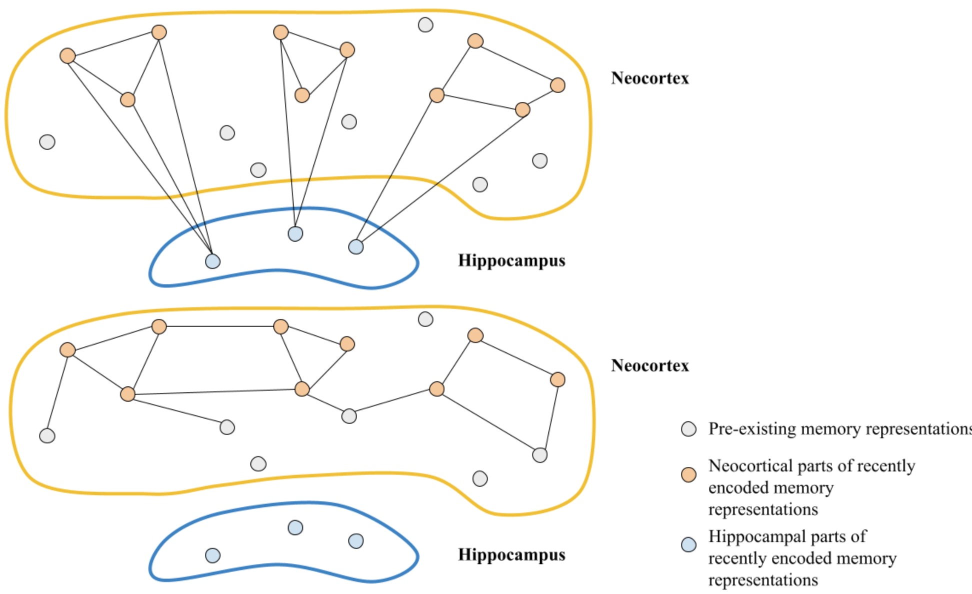

Figure 3.1: An illustration of reactivation of hippocampus and neocortex during memory consolidation.

Therefore, this process enables your memory system to form more general representations, such as what the past kids' birthday parties you have been to were like. This representation is stripped of many details specific to a particular memory episode but not shared across different episodes, as they are lost in the consolidation process. The general representation also gets updated as you acquire new experiences. If you recently went to a few kids' birthday parties where there was no balloon because the hosts do not like plastic wastes, then your general representation of kids' birthday parties gets updated, and the feature of having balloons becomes weakened.

Here, I would like to clarify the scope of memory representations that the process of memory consolidation applies to. Traditionally, psychologists draw a distinction between two kinds of long-term declarative memory: episodic and semantic memory (Tulving, 2002). Episodic memory is involved in remembering specific episodes of personally experienced events. For example, when I recall being at a specific kids' birthday party to celebrate my youngest cousin's birthday, I am retrieving an episodic memory. Semantic memory, on the other hand, is involved in remembering more abstract facts. For example, when I recall that guests at kids' birthday parties are usually expected to bring gifts, or that quinceañera is a celebration of a girl's 15th birthday I am recalling a piece of semantic memory.

Both episodic and semantic memory can be the inputs of memory consolidation. First, from the above example of consolidating different memory episodes of specific kids' birthday parties, it is quite obvious that episodic memory can be the input of memory consolidation. In fact, memory consolidation is closely related to the process of semanticization of episodic memory, i.e. the process through which specific episodic memory loses its contextual details and turns into more abstract and general forms of semantic memory (Aronowitz, 2023). Moreover, semantic memories can also be inputs to the memory consolidation process. For example, if I learned many facts about the customs of quinceñera in different Latin American countries, through memory consolidation I can then form the more general summary that quinceñera is a celebration of girls' transition from childhood to adulthood.[38]

[38] My account also does not assume that there is a binary distinction between episodic and semantic memory. The output of memory consolidation can preserve elements of both episodic and semantic memory, such as a generic memory of the experience of being at kids' birthday parties and being surrounded by loud kids. It is therefore compatible with recent challenges to the episodic/semantic memory distinction (Boyle, 2022; Aronowitz, 2023).

Moreover, in addition to studies that show that declarative memory (including episodic and semantic memory) goes through a consolidation process, there is also evidence that non-declarative memory (including procedural memory about how to perform skills) also goes through a consolidation process, and that the consolidation process also provides benefit for skill acquisition (Wang et al., 2021). Individuals are better at performing some skills that they learned after memory consolidation during sleep or wakeful rests. It is therefore plausible that memory consolidation of both declarative and nondeclarative memory improves understanding. However, I am only concerned with epistemic instead of practical understanding in this paper, and nondeclarative memory is typically involved in improving practical understanding, e.g. understanding how to play backhands in tennis, or understanding how to play a piece on the violin. Therefore, I will only focus on memory consolidation of declarative memory in this paper.

Having briefly illustrated the neurobiological mechanisms that underlie memory consolidation, as well as the kinds of memory representations that this process applies to, in the next section I will discuss several important processes that happen during memory consolidation that are crucial for getting representations that improve understanding.

4. How does memory consolidation improve understanding

Now that I have introduced what memory consolidation is, I will provide my argument that improving understanding is one of the functions of memory. Recall that in section 2, I argued that an investigator can improve their understanding of a subject or a system by committing to a representation that is both structured and goal-sensitive. It is structured in the sense it highlights important underlying relationships between different components of the target, and thereby reveals how things “hang together”. It is goal-sensitive in that the relationships that it highlights are the ones that matter for the cognitive or practical goal of the agent. In this section, I will argue that memory consolidation generates representations that have both features. Memory consolidation generates structured representations thanks to two mechanisms that I will identify below: (1) extraction of general patterns from disparate contents, and (2) integration of piecemeal information. Both of these two mechanisms generate new representations through further processing of previously acquired information. Moreover, I will draw on empirical evidence to argue that the representations generated through consolidation are appropriate for the agent's goals thanks to the fact that the consolidation process itself is adaptive to the agent's goals.

4.1. Extracting general patterns

One kind of process that occurs during consolidation involves extracting more general and qualitative representations based on the detailed contents initially encoded in memory (Klinzing et al. 2019; Gilboa & Marlatte, 2017; Tse et al. 2007). These representations that highlight more general and qualitative features can then improve our understanding of the represented subjects by revealing underlying connections between disparate contents in virtue of their shared common features. I start by considering one study that provides some evidence that memory consolidation generates representations that highlight general features that individuals were not aware of before (Wagner et al., 2004).

In the experiment, subjects are tested twice on several math problems with several hours in between. They are divided into a "sleep" group that gets retested after a night's sleep and two "no sleep" groups that either stay up overnight and get tested again in the morning, or first get tested in the morning, stay awake during the day, and then get tested again in the evening on the same day.

The problems that the subjects are asked to solve are a kind of number reduction task (NRT). In each string there are three digits (1,4,9), and they are processed through the following rules: (1) two identical digits produce the same digit (i.e. 1 and 1 result in 1); (2) two non-identical digits produce the third remaining digit (i.e. 1 and 4 result in 9). Subjects are asked to calculate the final result for each string as fast as possible.

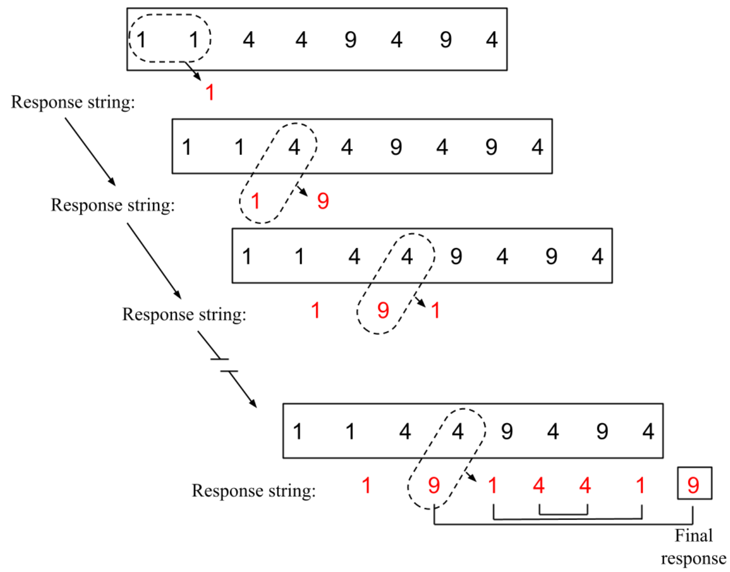

Figure 3.2: Illustration of the number reduction task from Walker et al. (2004). The string to be processed in this case is "11449494", and the response string is "1914419", with the last digit "9" being the final response.

While not announced to the subjects, there is an underlying pattern to the strings that the subjects are presented with: the strings are designed in such a way that the last three numbers ("419" in the example) in the response string mirror the previous three responses ("914" in the example). Once subjects notice this pattern by themselves, they can just report the second letter in the response string as the final answer, instead of processing all the way to the last digit. Discovering this pattern will thus speed up the calculation dramatically.

The reason for having two "no sleep" groups is to control for the effect of fatigue on their performance on the retest, and the "no sleep" group that stays up during the day (instead of during the evening) in fact self-reports similar level of fatigue as the "sleep" group when the two groups are retested. Therefore, the difference between the performance of participants in the "sleep" group and the “no sleep” group that stays awake during the day can only be explained in terms of factors other than exhaustion, making sleep-time memory consolidation a very probable difference maker.

Results showed that whether the participants went to sleep or not contributes hugely to how likely the participants will discover the underlying pattern. Compared to the two “no sleep” groups (including the one that stayed awake through the night, and the other that stayed awake throughout the day), participants in the group that goes to sleep after doing the problems for the first time is much more likely to have the insight about the underlying pattern when they perform the task again the second day, which thereby allows them to solve the task much faster.

Given the important role of sleep for memory consolidation, a plausible explanation for the difference between the performance of the sleep and no sleep groups is that the memory consolidation process during sleep made a huge difference. While as I have mentioned before, memory consolidation occurs both during sleep and wakeful rests, some studies have shown that sleep plays an especially beneficial role in consolidating memory that involves associated contents instead of individual items (Studte et al., 2015). Moreover, it is unclear whether the “no sleep” groups in fact have the opportunity to take rests that are free from external stimuli.[39] These two conditions together suggest that the difference between the sleep group and the no-sleep group can plausibly be explained by the effect of memory consolidation during sleep.

Therefore, it is likely that during sleep, the memory system of the subjects in the sleep group undergoes reactivation and generates new, abstract representations based on features extracted from the more detailed features and properties in the original observation. In particular, the new abstract representation extracts a pattern shared by the different response strings: that they all share the property of having the last three digits mirror the three digits before them.[40] When the subjects in the sleep group are retested on the task the second day, they retrieve the abstract representation generated

[39] Wagner et al. (2004) does not provide data about what the “no-sleep” groups do in their awake time before the retest.
[40] I discussed this case and the role of gist memory in creative thinking in more detail in Chapter 2.

through memory consolidation, and bring it into conscious cognition, which thereby allows them to have the insight for how to solve the problems much faster.

In this case, memory consolidation improves the participants' understanding of the math problem by generating representations that highlight underlying general and abstract patterns shared by disparate contents. These patterns would not have been easily accessible to the participants without memory consolidation due to the differences in the particular digits that make up each answer string.

This case of memory improving understanding is not just an accident. In fact more generally, extraction of abstract and general features is built into the cognitive architecture of the memory consolidation process (Frankland & Bontempi, 2005; Alba & Hasher, 1983), making it especially suitable for making accessible and highlighting these features that were initially not accessible to the individuals. For example, results from Durrant et al. (2011) shows that memory consolidation during slow-wave sleep, in particular, improves individuals' ability to extract statistical patterns from sequences of auditory stimuli. Individuals are much better at recognizing the statistical patterns after sleep, compared to if they were tested immediately after they listened to the stimuli. Therefore, memory consolidation *systematically* generates mental representations that make general and abstract features accessible. These features then enable memory consolidation to generate structured representations of how things "hang together" by prioritizing underlying commonalities between disparate contents that was not initially accessible to the investigator.

4.2. Integrating piecemeal information

Memory consolidation also involves another process that generates new representations by further processing previously acquired information: integration of piecemeal information with each other, or integration of newly acquired pieces of information into previously generated schema. This process can then generate representations of relationships between contents that individuals were not initially aware of (Ellenbogen et al., 2007; Cai et al., 2009; Lau et al., 2010). This process can thereby improve understanding when the relationship between distant contents is central to the cognitive purpose of the investigator.

Let me start by presenting a study that shows the role of memory consolidation in reorganizing and integrating piecemeal information that was acquired in a disorganized way. In a study by Ellenbogen and collaborators (Ellenbogen et al., 2007), subjects are shown some objects in pairs and instructed to learn the relation between each pair (i.e. which one should be selected over the other). Overall, they learn the relations schematized as “A>B”, “B>C”, “C>D”, “D>E”, and “E>F” (“>” meaning “selected over”, and the letters represent the objects), though not presented in this order to avoid revealing the underlying overall hierarchy. They are then separated into three groups to receive a follow-up test either after 20 minutes, or after 12 hours of awake time, or after 12 hours including sleep. The follow-up test involves comparisons between items that have one or two degrees of separation (e.g. comparing A and C, which involves one degree of separation, and comparing A and D, which involves two degrees of separation). Results show that both of the two groups that get tested after 12 hours performed better than the group that gets tested shortly after (i.e. in 20 minutes) with respect to the non-immediate comparison. Moreover, the group that goes to sleep during the 12 hour performed better than the group that stayed awake for the 12 hours with respect to comparisons that require two degrees of separation (e.g. comparing A and D, or B and E), even though the two groups performed equally well on comparisons that require one degree of separation (e.g. comparing A and C, or B and D).

A plausible explanation of these results is that the memory consolidation process during both sleep and wakeful rests enable individuals to recognize new connections between contents that were not presented together. Moreover, given the existing evidence that sleep in particular plays an especially crucial role in consolidating relational memory (Studte et al., 2015), we can further explain why the sleep group performs better than the group that was tested after 12 hours of awake time with respect to relations between more distant items: memory consolidation process during sleep enables individuals to recognize connections between even very distant contents by piecing together different relations that they learned at different times.

What might plausibly explain why memory consolidation can integrate information so effectively is the fact that contextual details are often forgotten during the consolidation process. As mentioned above, during sleep, this is possible through the transferral of information from the hippocampus (where contextual details are stored) to the neocortex. Psychologists have argued that mental context often shifts gradually throughout time (Howard et al., 2015).[41] Evidence for this theory includes the contiguity effect: we tend to recall items that were studied together in immediate succession as well. This is likely because neighboring items in the study stage have more similar contexts then items that were learned further apart in time in the study stage. In the case of the Ellenbogen et al. (2007) experiment, it is likely difficult for subjects who were tested shortly after the learning stage to compare distant items accurately, because it requires piecing together information that they learned in different contexts. For example, since they did not learn “A>B” and “B>C” in immediate succession, they likely associate “A>B” with one context, and “B>C” with a different one, and the contextual difference made if difficult for them to see that putting the two together and applying the transitivity rule generates “A>C”. Comparing items that require two degrees of separation (e.g. A and D) are even harder than comparing those with one degree of separation, because it requires piecing together three pieces of information (“A>B”, “B>C”, “C>D”), each of which is learned in a different context.

Therefore, discarding contextual details in fact comes with the surprising benefit of facilitating integration of information acquired in different contexts. In this case, the memory system

[41] There is also evidence that the context can shift abruptly when surprising events occur, or when there are changes in the subject’s goal (DuBrow et al., 2017). However, this is not the case throughout the study period in the Ellenbogen et al. (2007) experiment.

goes through an inference-like process and applies the transitivity rule to different pieces of information. This was not possible before discarding contextual details, because each of the individual pieces of information was associated with a different context, and so it was not possible to apply the transitivity rule to process them together.

Apart from this study that provides evidence for the role of memory consolidation in piecing together different newly acquired information, several studies have also shown that memory consolidation during sleep is crucial for integration of newly learned information into one's existing cognitive schema (Dumay & Gaskell, 2007; Tamminen et al., 2010). In particular, they showed that memory consolidation improves integration of newly learned words into one's lexicon. While the participants in some of these experiments do not have a certain cognitive purpose (e.g. in the Ellebogen et al. (2007) experiment), the same mechanism similarly happens in many contexts where investigators do in fact have a cognitive purpose, and relationships between different items are in fact crucial to the investigators' cognitive endeavors. For example, if I need to find out how to drive from Boston to Washington D.C. for an upcoming move, and suppose I learned in one occasion that I can drive to New York City from Boston on route 95, and in another occasion that I can drive to Washington D.C. from New York City on route 95 as well, then piecing together these two pieces of information will help me find out how drive from Boston to D.C. immediately. Similarly in science, piecing together data acquired from different experiments enables us to generate schemas that help us have a better grasp of the general subject, and integrating new data into the existing schema enables the existing schema to generalize to explain a wider range of phenomena. In these cases where the relationship between different items are crucial for the investigator's cognitive purpose, memory consolidation improves the investigator's understanding of the relevant subject by generating representations that make these otherwise elusive relationships accessible to the investigator.

4.3. Goal-sensitivity of memory consolidation

I have demonstrated that two different mechanisms of memory consolidation enable it to generate structured representations that prioritize relationships between different components of the subject or system that were not accessible to the investigator before. I will now explain how memory consolidation generates these representations in a way that is sensitive to the goals of the agent.

First, many psychologists have argued that the consolidation is highly sensitive to the agent's goals in the sense that memories that are tagged as useful for the agent's goals are processed and strengthened through replay during memory consolidation, while those memories that were not tagged as useful will likely be lost during consolidation (McNamara et al., 2014; Richards & Frankland, 2017; Cowan et al., 2021). Research has shown that subjects only show significant increase in brain activities that correspond to memory consolidation during sleep if they have been told during the initial learning stage that the information they are learning will be used for a future task (Wilhelm et al., 2011). In this study from Wilhelm and collaborators (2011), participants first participated in a card game to learn the location of some card pairs. One group of the participants are then told that they will be tested on the card locations at a later time, while another group is not given any information and do not expect a later test on what they learned. Both groups then go through a sleep period, and results show that the group that expects to be tested later displays more slow oscillation activity as well as sleep spindle count during post learning slow-wave sleep (SWS), and both of these two measurements correlate with increased consolidation. This result therefore shows that our memory system tends to consolidate memories that are important for our goals. A study on rats also shows that getting rewards enhances later reactivation of the memory during memory consolidation (Singer & Frank, 2009). The recording study from Singer and Frank shows that after rats go through spatial learning with rewards, their hippocampal cells show much more activity sharp wave ripples (SWRs) activities, which occur during memory consolidation, compared to unrewarded spatial learning. Given the close relationship between reward and goal, this result also shows that goal-relevant memories are replayed more during consolidation.

I grant that most of the experimental evidence only directly shows that the agent's goals influence which memory episode(s) gets processed and strengthened through replay, but does not directly show whether the agent's goal also affects what kind of information *within* the episode(s) is preserved through the processing during replay. However, some have suggested that the agent's goals influence the latter in a similar way as how it influences the former: within a memory episode (or across a few memory episodes), memory consolidation highlights information that is especially crucial given the agent's goals, while forgetting the information considered redundant and unuseful for the goals (Cowan et al., 2021). Consider an intuitive example from Cowan et al. (2021): suppose on your way to work, you ran towards the subway entrance to catch the train, and subsequently fell on the stairs as you were running downstairs at the entrance. Certainly some form of this memory will be retained in consolidation, because it is relevant to your goal of catching the train, as well as your goal of getting to where you need to be in one piece without hurting yourself. But more specifically within this episode, you will likely only remember the association between running downstairs and falling, but not the scenery you saw as you were running towards the subway entrance. This is because the former is relevant to your goal of catching the train in time and in one piece, while the latter is not. Cowan et al. (2021) argue that just as the episode itself is tagged as useful and preserved in memory consolidation, the more specific association between running downstairs and falling gets tagged and is subsequently preserved in memory consolidation, while the scenery you saw in the episode is lost.

So to sum up, our memory system systematically generates representations that make accessible and highlight relationships between different components of the subject or system that is represented. Moreover, the relationships that these representations selectively highlight are the ones that are crucial for the agent's background goal. Given the causal role that our memory system plays, the empirical evidence therefore supports my claim that one of the functions of our memory system is to improve understanding.

5. Varieties of epistemic generation in memory

I would now like to compare my account with the approach of some existing accounts in philosophy that also argue that memory generates new epistemic values. I will argue that many of the existing accounts focus on the role of memory in generating new epistemic values through a very different process: through memory accessing encapsulated information that is not accessible to normal cognition. While I agree that memory does generate new epistemic values through this means, my account differs from them by focusing on a very different way in which memory generates new epistemic values: through further processing information that was already available to normal cognition. These two ways in which memory can generate new epistemic values are not mutually exclusive: generating new epistemic values through further processing is compatible with the use of additional information that facilitates the processing, but the use of additional information is not required for the former. More importantly, the role of memory in generating epistemic values through further processing existing information has been overlooked by most of the existing literature on the epistemic role of memory, but it reflects one of the crucial functions of memory.

5.1. Epistemic generation in memory and the inaccessibility condition

In the past two decades, many who have argued that memory plays an epistemically generative role have focused on the role of memory in generating *justifications*, instead of understanding (Lackey, 2005; Michaelian, 2011; Fernández, 2016; Bernecker & Grundmann, 2019). Following the convention in the literature, I will use "generativism" to refer to the specific view that memory generates epistemic justification (instead of other epistemic values, such as understanding, knowledge, etc). I will consider a recent account from Miyazono and Tooming (2024b) that argues in favor of this view that memory generates new epistemic justification. Apart from offering an objection specific to this account, I will highlight the strategy shared by this account and a wider family of accounts. In particular, they tend to make use of the access that memory has to additional information encapsulated from normal cognition.

In their recent paper, Miyazono and Tooming (2024b) argue that memory generates new epistemic justification by drawing on the fact that memory has access to certain principles that are encapsulated from normal cognition.[42] I will first point out a problem in their argument specifically in which process (memory or core cognition) they attribute the epistemically generative role to. Then, I will illustrate a difference in focus between my account and a wider family of accounts that the argument from Miyazono and Tooming (2024b) belongs to.

Miyazono and Tooming's argument relies on what they call the "inaccessibility condition", which they also used in a similar argument that they developed for the claim that imagination also generates new justification (Miyazono & Tooming, 2024a). According to this condition, memory (or imagination) generates new justification thanks to the fact that it has access to principles encapsulated from normal cognition. In the case of imagination, they argue that these principles include those used in core cognition, intuitive physics, mental simulation, but in the case of memory, they focus on core cognition in particular.

Core cognition is a set of mental systems that emerge early in human development. It contains principles that support some basic intuitions about the world and are encapsulated from normal cognition, such as objects move as bounded wholes, through continuous paths, etc. Infants as young as a few months old have been demonstrated to have cognitive capacities that support the

[42] Like several others in this literature, Miyazono and Tooming draw a distinction between propositional and doxastic justification. This distinction might have influenced the focus of their account, and other accounts that also argue that memory generates new justification. I will introduce this distinction later.

claim that they have core cognition (Kinzler & Spelke, 2007). Miyazono and Tooming argue that memory can generate justifications for beliefs that agents cannot generate through normal cognition, thanks to the fact that memory has access to principles in core cognition in its processing of initially encoded information.

One kind of case that Miyazono and Tooming draw on is event completion. Sometimes our retrieved memory contents about past events involve details that were not part of the initial experience, but were instead filled in in accordance with principles stored in our core cognition. Miyazono and Tooming draw on studies that show that subjects who only observed fragments of an event that contains a causal relation recall having observed the entire event. For instance, you might remember the entire trajectory of a great goal that you scored in a soccer game, even though you only actually saw the beginning and the end of the trajectory. This happens because our core cognition system stores a principle about spatialtemporal continuity. According to this principle, if something is a cohesive object, then it will move along a continuous trajectory. Given that you observed the beginning and the end of the trajectory, the core cognition system then fills out the trajectory in between and generates the new content, i.e. the experience of seeing the entire trajectory of the goal. Miyazono and Tooming use this type of case to show that our retrieved memory content contains new content that subjects did not possess at the time of observation, and therefore our memory system generates new knowledge.

My objection to Miyazono and Tooming's argument is the following: it is unclear whether *memory* plays any significant role in the generation of the new justification, or whether core cognition on its own generates the new justification. In particular, it is unclear whether the memory system is involved at all in the generation of the new content that fills out the event. Empirical evidence shows that core cognition operates immediately and automatically on early-stage visual representations (Spelke, 2022). For this reason, core cognition likely already forms the representations that fill out the trajectory in between before the initial experience is encoded in long-term memory. If we were to ask the soccer player immediately after they saw the ball went into the goal, whether they saw the full trajectory of the ball, or if they only saw the beginning and the end of the ball, they would most likely say that they saw the entire trajectory. It is therefore unlikely that memory is involved at all in the generation of the new content, or in the generation of new justification.[43]

But more importantly, beyond this specific objection to Miyazono and Tooming, I would like to point out that the more general strategy that Miyazono and Tooming adopt focuses on a very different way in which memory generates new epistemic value, compared to the way in which memory is generative that my account focuses on. The more general strategy that Miyazono and Tooming adopt is to rely on the inaccessibility condition to argue for the epistemically generative role of memory, and I grant that some variations might be immune to my objection specific to core cognition. In their argument for the role of imagination (instead of memory) in generating new epistemic value, Miyazono and Tooming also appeal to the inaccessibility condition, and they offer a more general argument. In particular, they argue that imagination generates new justifications thanks to the fact that it has access to encapsulated principles in core cognition, intuitive physics, mental simulation, among others (Miyazono & Tooming, 2024a). Therefore, even if their specific argument for generativism about memory that relies on core cognition fails, it is still possible to devise similar arguments for generativism about memory that rely on the inaccessibility condition, but appeal to encapsulated principles in other modules.

For example, Michaelian (2011) provides a generativist argument based on the phenomenon of boundary extension. According to evidence from Intraub et al. (1992), humans tend to recall seeing pictures from a wider angle than they in fact did. For example, if they saw a picture of a house but close-up enough that some edges extend beyond the picture, they tend to complete the edges in

[43] Certainly to report that they saw the whole trajectory even immediately after they scored requires the use of working memory, but I am assuming that Miyazono and Tooming want to argue that generating new propositions involve memory in a more substantial sense, especially through involving long term memory .

their recalling. The completion of the shapes and edges of objects likely relies on information encapsulated in visual perception. Therefore, at least under one interpretation, Michaelian's generativist argument likely relies on the inaccessibility condition as well.

I agree with these accounts that memory's special access to certain information (compared to normal cognition) enables it to generate new epistemic values (including justifications). However, as I have argued with my account, there is another way in which memory can generate epistemic values: through further processing information that was already available to conscious cognition. This way in which memory generates new epistemic values has been overlooked by the literature on the epistemic role of memory so far. Importantly, to further process existing information already accessible to conscious cognition and thereby generate new epistemic values, memory does not necessarily need to access encapsulated information to assist the further processing of information.

Of course, encapsulated information *can* be involved in the processing of existing information (already available to normal cognition) as well. For example, results from Durrant et al. (2011) have shown that subjects become better at extracting statistical patterns from auditory inputs through memory consolidation. In this case, the processing of auditory inputs through memory consolidation likely involves rules about extracting statistical patterns that are encapsulated in perception. However, many of the cases that I have presented above of memory generating understanding through further processing existing information do not involve memory accessing additional information. In both the number reduction task from Wagner et al. (2004) and the inferential relations task from Ellenbogen et al. (2007), memory consolidation helped agents generate new representations through processing their previous acquired information better without accessing additional encapsulated information. In the case of the number reduction task, agents already had the implicit knowledge that if they saw the collection of answer strings that they were presented with, then there is a shared symmetrical pattern among the strings. This information is not encapsulated in a separate module, e.g. core cognition or visual perception, and inaccessible to normal cognition.

Similarly, in the case of the inferential relations task, the transitivity rule is also not encapsulated from normal cognition. Instead, agents already knew the transitivity rule, but simply could not apply it to the contents due to the contents being presented in a disorganized way. In these two cases, memory consolidation simply helped agents process the existing information already available to normal cognition, and helped them recognize underlying patterns as well as distant connections. This role that memory consolidation plays in processing existing information enables it to improve the agent's understanding, and to do that in neither of these two cases does memory need to appeal to additional information to facilitate the kinds of processing in these two tasks. In the next section, I consider a potential response from Miyaznono and Tooming, as well as many other generativists, motivated by the differences in the kinds of epistemic value that our respective accounts focus on.

5.2. Justification, understanding and generation of different epistemic values

I grant that some of the generativists might not think that my account shows that memory generates new epistemic values *by their measure*. This is because many generativists have a bias in favor of *propositional* instead of *doxastic* justification. Many in the literature on the epistemic role of memory draw a distinction between *propositional* and *doxastic* justification (Bernecker, 2009; Bernecker & Grundmann, 2019; Miyazono & Tooming, 2024b). The broader literature on justification in epistemology usually draws the distinction in the following way: P is *propositionally justified* for agent S if S is in possession of good reasons for believing in P. S's belief that P is *doxastically justified* if S in fact properly bases their belief in P on good reasons (Kvanvig, 2003; Feldman, 2002; Pollock and Cruz, 1999). A proposition that P can be *propositionally justified* for agent S, but S does not have doxastic justification for it. This happens, for example, when P is supported by all the evidence that S has, but S fails to form the belief that P in a way that is properly based on the evidence.

Miyazono and Tooming (2024b), among others, explicitly set propositional justification as their focus, and Bernecker (2009) claims that theories that show that memory generates "merely" doxastic justification does not succeed in showing that memory is "robustly" epistemically generative. Given their focus on propositional justification, it is quite plausible that they would not think of the role of memory in further processing information that was already available to the agent (without also accessing additional new information) as epistemically generative, because they might not think that memory generates new propositional justification through this process. To understand this worry consider some of the cases that I presented above as cases of memory improving understanding, but some might not take to be cases of memory generating propositional justification. For example, in the inferential relations task from Ellenbogen et al. (2007), prior to memory consolidation, the subjects already had the beliefs that A>B, B>C, and that the transitivity rule is true. They merely didn't put all of them together due to the fact that the different pieces of information were presented in different contexts. In this case, one might think that memory consolidation did not generate a new reason, but merely facilitated the process through which the agent formed a new belief fittingly based on the reasons. One might think that memory therefore "merely" generates doxastic justification, but not propositional justification.[44]

Relatedly, they might think that the role of memory in further processing information is *only* epistemically generative insofar as it makes use of additional information that was not already accessible to normal cognition, if they think that these are the cases where memory does generate new propositional justification. They might grant that memory generates propositional justification in the case of extracting patterns from auditory inputs in Durrant et al. (2011), since memory makes use

[44] There is the possibility that if one adopts certain views about what propositional justification is that takes the agent's capacities into account, we can still interpret memory as generating new propositional justification in this case. In particular, according to the account from Turri (2010), P is propositionally justified for agent S if S has the cognitive capacity to form the belief that P through the use of such capacity. One might think that memory generates new propositional justification in this case because the agent did not have the capacity to form the belief that "A>C" before. However, adopting this strategy would be assuming Turri's specific view of what propositional justification is as well as certain views about what capacity is, and I do not assume everyone in the literature would be on board with these assumptions.

of encapsulated principles about summarizing statistical patterns, which provides new reasons to the agent that they did not have before. I grant that memory is not “robustly” epistemically generative by their measure, if they are only interested in arguing for the claim that memory generates new propositional justification. I agree that there are good reasons for aiming for the claim about propositional justification: the claim that memory generates new propositional justification instead of doxastic justification or understanding is certainly more surprising.

However, the primary focus of my account is not on propositional justification, but rather on understanding, and there are good reasons for that as well: how well agents make use of the reasons that they have to form beliefs based on them is just as crucial for their epistemic lifes as how they acquire new reasons, and many actors affect the performance of non-ideal agents like us with respect to the former. The mere fact that we possess certain reasons rarely guarantee that we will make proper use of them. When the evidence is presented to us in a disorganized way, or when less useful features are made more salient, we tend to perform worse at making use of the reasons we have. The fact that memory improves agents’ performance in making proper use of the reasons they have therefore shows that memory plays an important epistemically generative role, regardless of whether memory also generates new reasons.

Finally, one might wonder: why does my account focus on the role of memory in improving understanding, instead of generating other epistemic values, such as epistemic justification broadly considered (including both doxastic and propositional justification), or knowledge? The exact consolidation process of generating structured and goal-sensitive representations that I take to improve understanding can also be used to argue that memory generates doxastic (and sometimes propositional) justification, as well as knowledge. I have already demonstrated above that we can plausibly interpret the consolidation process as generating doxastic (as well as in some cases, propositional) justifications. Moreover, all of the cases that I have presented involve generating *true* justified beliefs, which makes the claim that memory generates new knowledge highly plausible. In fact, I am open to the possibility that the consolidation process generates a variety of different epistemic values. However, focusing on understanding among one of the generated epistemic values has an explanatory advantage by capturing one of the functions of memory. While memory does generate justifications for many different kinds of beliefs, many of the justifications that memory generates are the specific type of justifications that are clustered around beliefs that improve understanding. They are beliefs about relations between different components of a certain subject or system, and these relations are the ones that matter for the goals of the agent. Similarly, lots of knowledge generated by memory is the special kind of knowledge that improves understanding.[45] The fact that memory generates these specific kinds of justifications and knowledge is determined by the cognitive architecture of the memory system. Therefore, an advantage of highlighting understanding as one of the epistemic values that memory generates is that it captures an important function of memory, i.e. to improve the agent's understanding of what is represented in their initial experience. This fact can then explain why memory also generates the specific type of knowledge that improves understanding, or the specific type of justification for understanding-conducive beliefs.[46]

6. Conclusion

In this paper, I have argued that one of the functions of the human memory system is to improve our understanding of subjects and systems represented in our initial experiences. Our memory system fulfills this function through the memory consolidation, which further processes information that we

[45] Those who think that understanding is reducible to knowledge might also think that lots of knowledge generated by memory is the special kind of knowledge that *constitutes* understanding.

[46] Finally, while I won't pursue this argument here, the similar argument also runs for knowledge (in addition to justification). The empirical evidence that I have provided likely also shows that memory generates knowledge, since the cases involve memory generating justified true belief, but there is equally the danger of casting the web too wide: memory in fact systematically generates the special type of knowledge that is conducive to understanding, or just is understanding for those who think that understanding is reducible to knowledge.

acquired before. Memory consolidation generates structured representations that make accessible or highlight hidden relationships that are crucial for our cognitive and practical goals. These representations are sometimes difficult to generate with conscious cognition but are made possible through memory consolidation, thanks to some of the processes that happen during consolidation, including integration of piecemeal information acquired in different contexts to highlight relationships between distant contents, as well as extracting general patterns from disparate memory contents. By reorienting the discussion about memory as a generative source of epistemic value towards its role in relation to understanding, I have diverged from the current generativist literature that focuses on the role of memory in generating epistemic justification, and more specifically propositional justification. By instead focusing on understanding, I highlight a crucial function of memory that has been overlooked by the current literature.

Bibliography

Abadie, M., Waroquier, L., & Terrier, P. (2013). Gist memory in the unconscious-thought effect. *Psychological Science*, 24(7), 1253-1259.

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, 93, 203–231.

Albrecht, A. R., & Scholl, B. J. (2010). Perceptually averaging in a continuous visual world: Extracting statistical summary representations over time. *Psychological Science*, *21*(4), 560-567.

Anderson, B. and Druker, M. (2013). Attention improves perceptual quality. *Psychonomic Bulletin & Review* 20(1):120-7.

Andrews‐Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self‐generated thought: Component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, 1316(1), 29-52.

Anton-Erxleben, K., & Carrasco, M. (2013). Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nature Reviews Neuroscience*, *14*(3), 188-200.

Aronowitz, S. (2019). Memory is a modeling system. *Mind & Language*, *34*(4), 483-502.

Aronowitz, S. (2023). Semanticization Challenges the Episodic-Semantic Distinction. *The British Journal for Philosophy of Science*

Baird, B., Smallwood, J., Mrazek, M. D., Kam, J. W., Franklin, M. S., & Schooler, J. W. (2012). Inspired by distraction: Mind wandering facilitates creative incubation. *Psychological science*, 23(10), 1117-1122.

Barron, F. (1955). The disposition toward originality. *The Journal of Abnormal and Social Psychology, 51,* 478–485.

Baumberger, C., Beisbart, C., & Brun, G. (2016). What is understanding? An overview of recent debates in epistemology and philosophy of science. In Grimm S., Baumberger C., & Sabine A. (Eds.), *Explaining understanding* (pp. 1-34). Routledge.

Beaty, R.E., & Silvia, P.J. (2012). Why do ideas get more creative across time? An executive interpretation of the serial order effect in divergent thinking tasks. *Psychology of Aesthetics, Creativity, and the Arts*, 6(4), 309–319. https://doi.org/10.1037/a0029171

Beaty, R. E., Silvia, P. J., Nusbaum, E. C., Jauk, E., & Benedek, M. (2014). The roles of associative and executive processes in creative cognition. *Memory & cognition*, 42, 1186-1197.

Beaty, R. E., & Kenett, Y. N. (2023). Associative thinking at the core of creativity. *Trends in Cognitive Sciences*.

Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766.

Benedek, M., Jauk, E., Sommer, M., Arendasy, M., & Neubauer, A.C. (2014). Intelligence, creativity, and cognitive control: the common and differential involvement of executive functions in intelligence and creativity. *Intelligence*, 46, 73–83. https://doi.org/10.1016/j.intell.2014.05.007

Benedek, M., Kenett, Y.N., Umdasch, K., Anaki, D., Faust, M., & Neubauer, A.C. (2017). How semantic memory structure and intelligence contribute to creative thought: a network science approach. *Thinking & Reasoning*, 23(2), 158–183. https://doi.org/10.1080/13546783.2016.1278034

Benedek, M., & Neubauer, A.C. (2013). Revisiting Mednick's model on creativity-related differences in associative hierarchies. Evidence for a common path to uncommon thought. *The Journal of Creative Behavior*, 47(4), 273–289. https://doi.org/10.1002/jocb.35

Bengson, J. (2015). A noetic theory of understanding and intuition as sense-maker. *Inquiry*, *58*(7-8), 633-668.

Bermúdez, J. P. (2017). Do we reflect while performing skillful actions? Automaticity, control, and the perils of distraction. *Philosophical Psychology, 30*(7), 896–924. https://doi.org/10.1080/09515089.2017.1325457

Bernecker, S. (2009). *Memory: A philosophical study*. OUP Oxford.

Bernecker, S., & Grundmann, T. (2019). Knowledge from forgetting. *Philosophy and Phenomenological Research*, *98*(3), 525-540.

Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Psychology Press.

Born, J. (2010). Slow-wave sleep and the consolidation of long-term memory. *The World Journal of Biological Psychiatry*, *11*(sup1), 16-21.

Brainerd, C. J., & Reyna, V. F. (2001). Fuzzy-trace theory: dual processes in memory, reasoning, and cognitive neuroscience. *Advances in child development and behavior*, *28*, 41-100.

Brosowsky, N. P., Murray, S., Schooler, J. W., & Seli, P. (2021). Attention need not always apply: Mind wandering impedes explicit but not implicit sequence learning. *Cognition*,*209*,104530.

Burge, T. (1997). Interlocution, perception, and memory. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *86*(1), 21-47.

Cai, D. J., Mednick, S. A., Harrison, E. M., Kanady, J. C., & Mednick, S. C. (2009). REM, not incubation, improves creativity by priming associative networks. *Proceedings of the National Academy of Sciences*, *106*(25), 10130-10134.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, 51(13), 1484– 1525.

Carruthers, P. (2020). Mechanisms for constrained stochasticity. *Synthese*, *197*(10), 4455-4473.

Carson, S. H., Peterson, J. B., & Higgins, D. M. (2003). Decreased latent inhibition is associated with increased creative achievement in high-functioning individuals. *Journal of personality and social psychology*, *85*(3), 499.

Chong, S. C., & Evans, K. K. (2011). Distributed versus focused attention (count vs estimate). *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(6), 634-638.

Chong, S.C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics* 67, 1–13. https://doi.org/10.3758/BF03195009

Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences*, 106(21), 8719-8724.

Christoff, K., Irving, Z. C., Fox, K. C., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews. Neuroscience*, 17(11), 718–731.

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407-1409.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental psychology: Human perception and performance*, *21*(1), 109.

Chun, M. M., Golomb, J. D., and Turk-Browne, N.B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*; 62:73-101.

Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in cognitive sciences*, *16*(8), 411-417.

Cowan, E. T., Schapiro, A. C., Dunsmoor, J. E., & Murty, V. P. (2021). Memory consolidation as an adaptive process. *Psychonomic Bulletin & Review*, 1-15.

Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, 72(741), 765.

Cummins, R. (1983). *The nature of psychological explanation*. Cambridge, MA: MIT Press.

De Brigard, F. (2012). The role of attention in conscious recollection. *Frontiers in Psychology*, *3*, 29.

De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, *191*, 155-185.

Diekelmann, S., Wilhelm, I., & Born, J. (2009). The whats and whens of sleep-dependent memory consolidation. *Sleep medicine reviews*, *13*(5), 309-321.

Dixon, M. L., Fox, K. C., & Christoff, K. (2014). A framework for understanding the relationship between externally and internally directed cognition. *Neuropsychologia*, *62*, 321-330.

Drosopoulos, S., Schulze, C., Fischer, S., & Born, J. (2007). Sleep's function in the spontaneous recovery and consolidation of memories. *Journal of Experimental Psychology: General*, *136*(2), 169.

DuBrow, S., Rouhani, N., Niv, Y., & Norman, K. A. (2017). Does mental context drift or shift?. *Current opinion in behavioral sciences*, *17*, 141-146.

Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science*, *18*(1), 35-39.

Durrant, S. J., Taylor, C., Cairney, S., & Lewis, P. A. (2011). Sleep-dependent consolidation of statistical learning. *Neuropsychologia*, *49*(5), 1322-1331.

Elgin, C. (2009). Is understanding factive?. *Epistemic value*, 322-330.

Elgin, C. Z. (2016). Exemplification in understanding. In Grimm S., Baumberger C., & Sabine A. (Eds.), *Explaining understanding* (pp. 1-34). Routledge.

Elgin, C. (2017). *True enough*. MIT press.

Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723-7728.

Feldman, R. (2002). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.

Fernández, J. (2016). Epistemic generation in memory. *Philosophy and Phenomenological Research*, 92(3), 620-644.

Fiebelkorn, I. C., Kastner, S. (2020). Functional specialization in the attention network. *Annual Review of Psychology*. Jan 4;71: 221-249.

Friedman, R. S., Fishbach, A., Förster, J., & Werth, L. (2003). Attentional priming effects on creativity. *Creativity research journal*, *15*(2-3), 277-286.

Finke, R. A., Ward, T. B., & Smith, S. M. (1996). *Creative cognition: Theory, research, and applications*. MIT press.

Fischer, S., Drosopoulos, S., Tsen, J., & Born, J. (2006). Implicit learning–explicit knowing: a role for sleep in memory system interaction. *Journal of cognitive neuroscience*, *18*(3), 311-319.

Fogel, S., Ray, L., Fang, Z., Silverbrook, M., Naci, L., & Owen, A. M. (2022). While you were sleeping: Evidence for high-level executive processing of an auditory narrative during sleep. Consciousness and Cognition, 100, 103306. https://doi.org/10.1016/j.concog.2022.103306

Fortney, M. (2018). The centre and periphery of conscious thought. *Journal of Consciousness Studies*, *25*(3-4), 112-136.

Fortney, M. (2020). Unity in the scientific study of intellectual attention. *Canadian Journal of Philosophy*, *50*(4), 444-459.

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature reviews neuroscience*, *6*(2), 119-130.

Friedman, R. S., Fishbach, A., Förster, J., & Werth, L. (2003). Attentional priming effects on creativity. *Creativity research journal*, 15(2-3), 277-286.

Frise, M. (2023). "Epistemological problems of memory." In *Stanford Encyclopedia of Philosophy*, March 27. https://plato.stanford.edu/entries/memory-episprob/.

Frith, E., Kane, M. J., Welhaf, M. S., Christensen, A. P., Silvia, P. J., & Beaty, R. E. (2021). Keeping creativity under control: Contributions of attention control and fluid intelligence to divergent thinking. *Creativity Research Journal*, 33(2), 138-157.

Fujino, M., Ueda, Y., Mizuhara, H., Saiki, J., & Nomura, M. (2018). Open monitoring meditation reduces the involvement of brain regions related to memory function. *Scientific reports*, *8*(1), 9968. https://doi.org/10.1038/s41598-018-28274-4

Gabora, L. (2018). The neural basis and evolution of divergent and convergent thought. In O. Vartanian & R. Jung (Eds.) *The Cambridge Handbook of the Neuroscience of Creativity* (pp. 58-70). Cambridge MA: Cambridge University Press. https://doi.org/10.1017/9781316556238.005

Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive science*, *7*(2), 155-170.

Gilboa, A., & Marlatte, H. (2017). Neurobiology of schemas and schema-mediated memory. *Trends in cognitive sciences*, *21*(8), 618-631.

Goldman, A. (2009), "Internalism, Externalism, and the Architecture of Justification", *Journal of Philosophy*, 106(6): 309–338.

Gopnik, A. (2009). *The philosophical baby: What children's minds tell us about truth, love, and the meaning of life.* Bodley Head.

Greenberg, G. (2023). The iconic-symbolic spectrum. *Philosophical Review*, *132*(4), 579-627.

Grimm, S. R. (2006). Is understanding a species of knowledge?. *The British Journal for the Philosophy of Science*.

Grimm, S. (2011). “Understanding”. In Bernecker, S., & Pritchard, D. (Eds.), *The Routledge companion to epistemology* (pp. 84–94). London: Routledge.

Hills, A. (2016). Understanding why. *Noûs*, *50*(4), 661-688.

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological review*, *119*(2), 431.

Howard, M. W., Shankar, K. H., Aue, W. R., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological review*, *122*(1), 24.

Intraub, H., Bender, R. S., & Mangels, J. A. (1992). Looking at pictures but remembering scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(1), 180.

Intraub, H., & Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 179.

Irving, Z. C. (2016). Mind-wandering is unguided attention: Accounting for the “purposeful” wanderer. *Philosophical Studies*, *173*, 547-571.

Irving, Z. C. (2021). Drifting and directed minds: The significance of mind-wandering for mental agency. *The Journal of Philosophy*, 118(11), 614-644.

Irving, Z. C., & Glasser, A. (2020). Mind‐wandering: A philosophical guide. *Philosophy Compass*, 15(1), e12644.

Irving, Z. C., McGrath, C., Flynn, L., Glasser, A., & Mills, C. (2022). The shower effect: Mind wandering facilitates creative incubation during moderately engaging activities. *Psychology of Aesthetics, Creativity, and the Arts*.

Jenkin, Z. (2020). The epistemic role of core cognition. *Philosophical Review*, *129*(2), 251-298.

Jenkin, Z. (2022). Crossmodal basing. *Mind*, *131*(524), 1163-1194.

Kam, J. W., Irving, Z. C., Mills, C., Patel, S., Gopnik, A., & Knight, R. T. (2021). Distinct electrophysiological signatures of task-unrelated and dynamic thoughts. *Proceedings of the National Academy of Sciences*, *118*(4), e2011796118.

Kenett, Y. N., Anaki, D., & Faust, M. (2014). Investigating the structure of semantic networks in low and high creative persons. *Frontiers in human neuroscience*, *8*, 407.

Kenett, Y. N., Beaty, R. E., Silvia, P. J., Anaki, D., & Faust, M. (2016). Structure and flexibility: Investigating the relation between the structure of the mental lexicon, fluid intelligence, and creative achievement. *Psychology of Aesthetics, Creativity, and the Arts*, *10*(4), 377.

Kentridge, R. W. (2011). Attention without awareness. In C. Mole, D. Smithies & W. Wu (eds.), *Attention: Philosophical and Psychological Essays*. Oxford University Press. pp. 228.

Kinzler, K. D., & Spelke, E. S. (2007). Core systems in human cognition. *Progress in brain research*, *164*, 257-264.

Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature neuroscience*, *22*(10), 1598-1610.

Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge university Press.

Kvanvig, J. L. (2009). 'The Value of Understanding', in A. Haddock, A. Millar, and D. Pritchard (eds), *Epistemic Value* (Oxford, 2009), https://doi.org/10.1093/acprof:oso/9780199231188.003.0005.

Lackey, J. (2005), "Memory as a Generative Epistemic Source", *Philosophy and Phenomenological Research*, 70(3): 636–658. doi:10.1111/j.1933-1592.2005.tb00418.x

Langland-Hassan, P. (2020). *Explaining imagination* (p. 336). Oxford University Press.

Lau, H., Tucker, M. A., & Fishbein, W. (2010). Daytime napping: Effects on human direct associative and relational memory. *Neurobiology of Learning and Memory*, 93(4), 554–560.

Levy, Y. (2022). The most general mental act. In *Mental Action and The Conscious Mind*, edited by Michael Brent and Lisa Miracchi, Routledge.

Lewis, P. A., & Durrant, S. J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends in cognitive sciences*, *15*(8), 343-351.

Li, F.F., VanRullen, R., Koch, C. and Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences USA* 99, 9596 – 9601.

Li, Y., Kenett, Y. N., Hu, W., & Beaty, R. E. (2021). Flexible semantic network structure supports the production of creative metaphor. *Creativity Research Journal*, *33*(3), 209-223.

Lutz, N. D., Diekelmann, S., Hinse-Stern, P., Born, J., & Rauss, K. (2017). Sleep supports the slow abstraction of gist from visual perceptual memories. *Scientific reports*, *7*(1), 42950.

Lutz, A., Slagter, H. A., Dunne, J. D. & Davidson, R. J. (2008). Cognitive-emotional interactions: Attention regulation and monitoring in meditation. *Trends. Cogn. Sci.* 12, 163–169.

Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Noûs*, *50*(3), 629-658.

Marchi, F. (2020). *The Attentional Shaping of Perceptual Experience*. Studies in Brain and Mind, Springer. https://doi.org/10.1007/978-3-030-33558-8.

McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature neuroscience*, *17*(12), 1658-1660.

McVay, J. C., & Kane, M. J. (2009). Conducting the train of thought: working memory capacity, goal neglect, and mind wandering in an executive-control task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(1), 196

Mednick, S. (1962). The associative basis of the creative process. *Psychological review*, 69(3), 220.

Michaelian, K. (2016). *Mental time travel: Episodic memory and our knowledge of the personal past*. MIT Press.

Michaelian, K. (2011), Generative Memory. *Philosophical Psychology*, 24(3): 323–342. doi:10.1080/09515089.2011.559623

Millikan, R. G. (1984). *Language, thought and other biological categories*. Cambridge, MA: The MIT Press.

Miyazono, K., & Tooming, U. (2023). On the putative epistemic generativity of memory and imagination. In *Philosophical Perspectives on Memory and Imagination*. Taylor & Francis.

Miyazono, K., & Tooming, U. (2024a). Imagination as a generative source of justification. *Noûs*, *58*(2), 386-408.

Miyazono, K., & Tooming, U. (2024b). Prospects for epistemic generationism about memory. *Philosophy and the Mind Sciences, 5.*

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science (New York, N.Y.)*, *229*(4715), 782–784. https://doi.org/10.1126/science.4023713

Mudumba, R., & Srinivasan, N. (2021, July 27 preprint). Broad scope of attention results in better temporal resolution: Evidence from a temporal order judgment study.

Mumford, M. D., Mobley, M. I., Reiter‐Palmon, R., Uhlman, C. E., & Doares, L. M. (1991). Process analytic models of creative capacities. *Creativity research journal*, *4*(2), 91-122.

Murray, N., Sujan, H., Hirt, E. R., & Sujan, M. (1990). The influence of mood on categorization: A cognitive flexibility interpretation. *Journal of Personality and Social Psychology*, *59*(3), 411.

Nanay, B. (2016). *Aesthetics as Philosophy of Perception*. Oxford, United Kingdom: Oxford University Press UK.

Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2013). Object-based attention without awareness. *Psychological science*, *24*(6), 836–843. https://doi.org/10.1177/0956797612461449

Nusbaum, E. C., & Silvia, P. J. (2011). Are intelligence and creativity really so different?: Fluid intelligence, executive processes, and strategy use in divergent thinking. *Intelligence*, 39(1), 36-45.

Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, *155*, 23-36.

Olivers, C. N., and Nieuwenhuis, S. (2005). The beneficial effect of concurrent task-irrelevant mental activity on temporal attention. *Psychological Science* 16, 265–269.

Olson, J. A., Nahas, J., Chmoulevitch, D., Cropper, S. J., & Webb, M. E. (2021). Naming unrelated words predicts creativity. *Proceedings of the National Academy of Sciences*, *118*(25), e2022340118.

Phillips, I. (2012). Attention to the passage of time. *Philosophical Perspectives*, *26*, 277-308.

Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge*. Rowman & Littlefield Publishers.

Posner, M. I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Sciences*, *91*(16), 7398-7403.

Potochnik, A. (2020). Idealization and many aims. *Philosophy of science*, 87(5), 933-943.

Prettyman, A. (2023). What is diffuse attention? *Mind and Language*.

Prettyman, A. (2014). *Attention and Conscious Perception*. Doctoral dissertation: University of Toronto.

Pritchard, D. (2010). Knowledge and Understanding. In D. Pritchard, A. Millar, & A. Haddock (Eds.), *The Nature and Value of Knowledge: Three Investigations*. Oxford University Press.

Rångtell, F. H., Karamchedu, S., Andersson, P., van Egmond, L., Hultgren, T., Broman, J. E., ... & Benedict, C. (2017). Learning performance is linked to procedural memory consolidation across both sleep and wakefulness. *Scientific reports*, *7*(1), 10234.

Rawlings, D. (1985). Psychoticism, creativity and dichotic shadowing. *Personality and individual differences*, *6*(6), 737-742.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849–860.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and individual Differences*, 7(1), 1-75.

Richards, B. A., & Frankland, P. W. (2017). The persistence and transience of memory. *Neuron*, *94*(6), 1071-1084.

Riggs, W.D. (2003). Understanding 'virtue' and the virtue of understanding. In *Intellectual virtues: Perspectives from ethics and epistemology*, eds. M. DePaul and L. Zagzebski, 203–226. New York, NY: Oxford University Press.

Riggs, W. (2009). Understanding, knowledge, and the meno requirement. *Epistemic value*, *336*.

Robin, J., & Moscovitch, M. (2017). Details, gist and schema: hippocampal–neocortical interactions underlying recent and remote episodic and spatial memory. *Current opinion in behavioral sciences*, 17, 114-123.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of experimental psychology: Learning, Memory, and Cognition*, *21*(4), 803.

Salvaggio, M. (2018). The justification of reconstructive and reproductive memory beliefs. *Philosophical Studies*, *175*, 649-663.

Samanta, A., van Rongen, L. S., Rossato, J. I., Jacobse, J., Schoenfeld, R., & Genzel, L. (2021). Sleep Leads to Brain-Wide Neural Changes Independent of Allocentric and Egocentric Spatial Training in Humans and Rats. *Cerebral Cortex*, Volume 31(11), p. 4970-4985. https://doi.org/10.1093/cercor/bhab135

Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 773-786.

Schwitzgebel, E. (2013). A dispositional approach to attitudes: Thinking outside of the belief box. In *New essays on belief: Constitution, content and structure* (pp. 75-99). London: Palgrave Macmillan UK.

Silvia, P. J., & Beaty, R. E. (2012). Making creative metaphors: The importance of fluid intelligence for creative thought. *Intelligence*, 40(4), 343-351.

Singer, A. C., & Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, *64*(6), 910-921.

Sirota, A., & Buzsáki, G. (2005). Interaction between neocortical and hippocampal networks via slow oscillations. *Thalamus & related systems*, *3*(4), 245-259.

Sliwa, P. (2015). Understanding and knowing. In *Proceedings of the Aristotelian Society* (Vol. 115, No. 1, pp. 57-74). Oxford, UK: Oxford University Press.

Sliwa, P. (2017). Moral understanding as knowing right from wrong. *Ethics*, *127*(3), 521-552.

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

Smallwood, J., & Schooler, J. W. (2006). The restless mind. *Psychological bulletin*, *132*(6), 946.

Smallwood, J., Tipper, C., Brown, K., Baird, B., Engen, H., Michaels, J. R., ... & Schooler, J. W. (2013). Escaping the here and now: evidence for a role of the default mode network in perceptually decoupled thought. *Neuroimage*, *69*, 120-125.

Smith, S. M., & Blankenship, S. E. (1989). Incubation effects. *Bulletin of the Psychonomic Society*, *27*(4), 311-314.

Smith, S. M., & Blankenship, S. E. (1991). Incubation and the persistence of fixation in problem solving. *The American journal of psychology*, 61-87.

Spelke, E. S. (2022). *What babies know: Core Knowledge and Composition volume 1* (Vol. 1). Oxford University Press.

Squire, L. R., Genzel, L., Wixted, J. T., & Morris, R. G. (2015). Memory consolidation. *Cold Spring Harbor perspectives in biology*, *7*(8), a021766.

Srinivasan, N., Srivastava, P., Lohani, M., and Baijal, S. (2009). Focused and distributed attention. *Progress in Brain Research* 176:87-100.

Sripada, C. S. (2018). An exploration/exploitation trade-off between mind wandering and goal-directed thinking. *The Oxford handbook of spontaneous thought: Mind-wandering, creativity, and dreaming*, 23-34.

Stein, M. I. (1953). Creativity and culture. *The journal of psychology*, *36*(2), 311-322.

Stokes, D. (2021). On perceptual expertise. *Mind & Language*, *36*(2), 241-263.

Stokes, D. (2017). Imagination and Creativity. In A. Kind (Ed.), *The Routledge Handbook of Philosophy of Imagination*. essay, Routledge, Taylor & Francis Group.

Stokes, D. (2014). The role of imagination in creativity. *The philosophy of creativity: New essays*, 157-184.

Stokes, D. (2007). Incubated cognition and creativity. *Journal of Consciousness Studies*, *14*(3), 83-100.

Studte, S., Bridger, E., & Mecklinger, A. (2015). Nap sleep preserves associative but not item memory performance. *Neurobiology of learning and memory*, *120*, 84-93.

Taatgen, N. A., Juvina, I., Schipper, M., Jelmer P., Borst, J. and Martens, S. (2009). Too much control can hurt: A threaded cognition model of the attentional blink. *Cognitive psychology* 59, no. 1: 1-29.

Tamminen, J., Payne, J. D., Stickgold, R., Wamsley, E. J., & Gaskell, M. G. (2010). Sleep spindle activity is associated with the integration of new memories and existing knowledge. *Journal of Neuroscience*, *30*(43), 14356-14360.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and brain sciences*, 24(4), 629-640.

Thakral, P. P., Devitt, A. L., Brashier, N. M., & Schacter, D. L. (2021). Linking creativity and false memory: Common consequences of a flexible memory system. *Cognition*, *217*, 104905.

Thomson, D. R., Ralph, B.C., Besner, D., and Smilek D. (2015). The more your mind wanders, the smaller your attentional blink: An individual differences study. *Quarterly Journal of Experimental Psychology*; 68(1):181-91.

Treisman A. (2006). How the deployment of attention determines what we see. *Visual cognition*, *14*(4-8), 411–443.

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., ... & Morris, R. G. (2007). Schemas and memory consolidation. *Science*, *316*(5821), 76-82.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual review of psychology*, *53*(1), 1-25.

Turri, J. (2010). On the relationship between propositional and doxastic justification. *Philosophy and phenomenological research*, *80*(2), 312-326.

Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Sleep inspires insight. *Nature*, *427*(6972), 352-355.

Wamsley, E. J. (2022). Offline memory consolidation during waking rest. *Nature Reviews Psychology*, *1*(8), 441-453.

Wang, S. Y., Baker, K. C., Culbreth, J. L., Tracy, O., Arora, M., Liu, T., ... & Wamsley, E. J. (2021). 'Sleep-dependent' memory consolidation? Brief periods of post-training rest and sleep provide an equivalent benefit for both declarative and procedural memory. *Learning & Memory*, 28(6), 195-203.

Watzl, S. (2023). What attention is. The priority structure account. *WIREs Cognitive Science*, 14(1), e1632. https://doi.org/10.1002/wcs.1632

Watzl, S. (2017). *Structuring Mind: The Nature of Attention and how it Shapes Consciousness*. Oxford: Oxford University Press.

Wegbreit, E., Suzuki, S., Grabowecky, M., Kounios, J., & Beeman, M. (2012). Visual attention modulates insight versus analytic solving of verbal problems. *The journal of problem solving*, 4(2), 94.

Wilde, O. (1995). *Lady Windermere's fan: Salome; A woman of no importance; An ideal husband; The importance of being earnest*. Oxford University Press.

Wilhelm, I., Diekelmann, S., Molzow, I., Ayoub, A., Mölle, M., & Born, J. (2011). Sleep selectively enhances memory expected to be of future relevance. *Journal of Neuroscience*, *31*(5), 1563-1569.

Wright, L. (1973). Functions. *Philosophical Review*, *82*, 139–168.

Wu, W. (2014). *Attention*. Routledge.

Yeshurun, Y., & Levy, L. (2003). Transient spatial attention degrades temporal resolution. *Psychological Science*, *14*(3), 225–231.

Zabelina, D., Saporta, A., & Beeman, M. (2016). Flexible or leaky attention in creative people? Distinct patterns of attention for different types of creative thinking. *Memory & cognition*, 44, 488-498.

Zabelina, D.L. (2018). Attention and Creativity. In R. E. Jung & O. Vartanian (Eds.), *The Cambridge Handbook of the Neuroscience of Creativity* (pp. 161–179). Cambridge: Cambridge University Press. https://doi.org/10.1017/9781316556238.010

Zagzebski, L. (2001). Recovering Understanding. In M. Steup (Ed), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (pp. 235-252). Oxford University Press.

List of figures

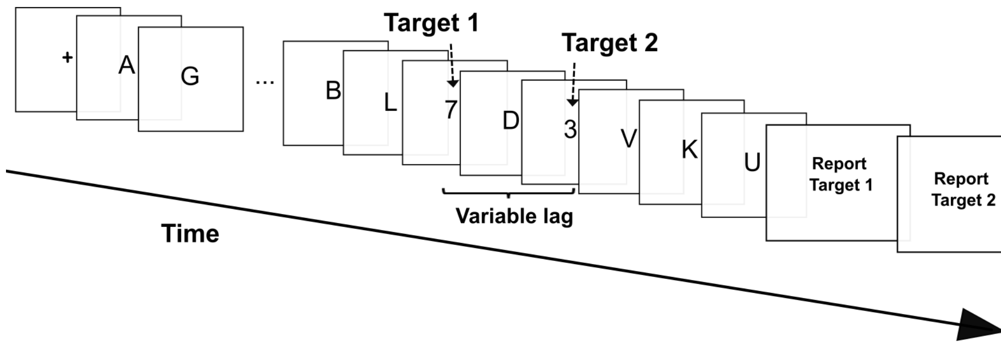

Figure 1.1: RSVP task from Olivers and Nieuwenhuis (2005).

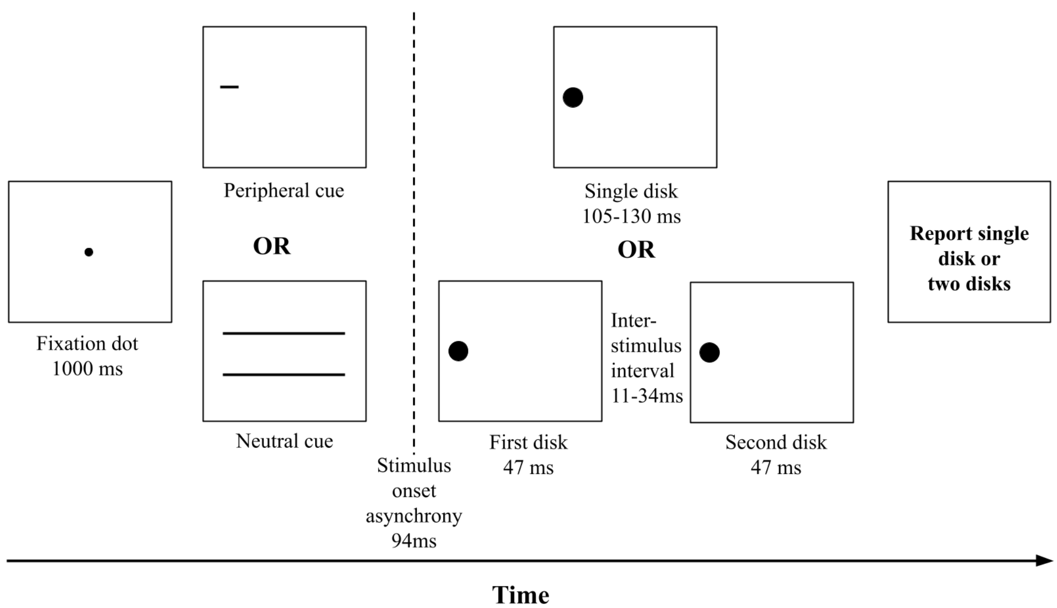

Figure 1.2: Temporal resolution task from Y. Yeshurun and L. Levy (2003).

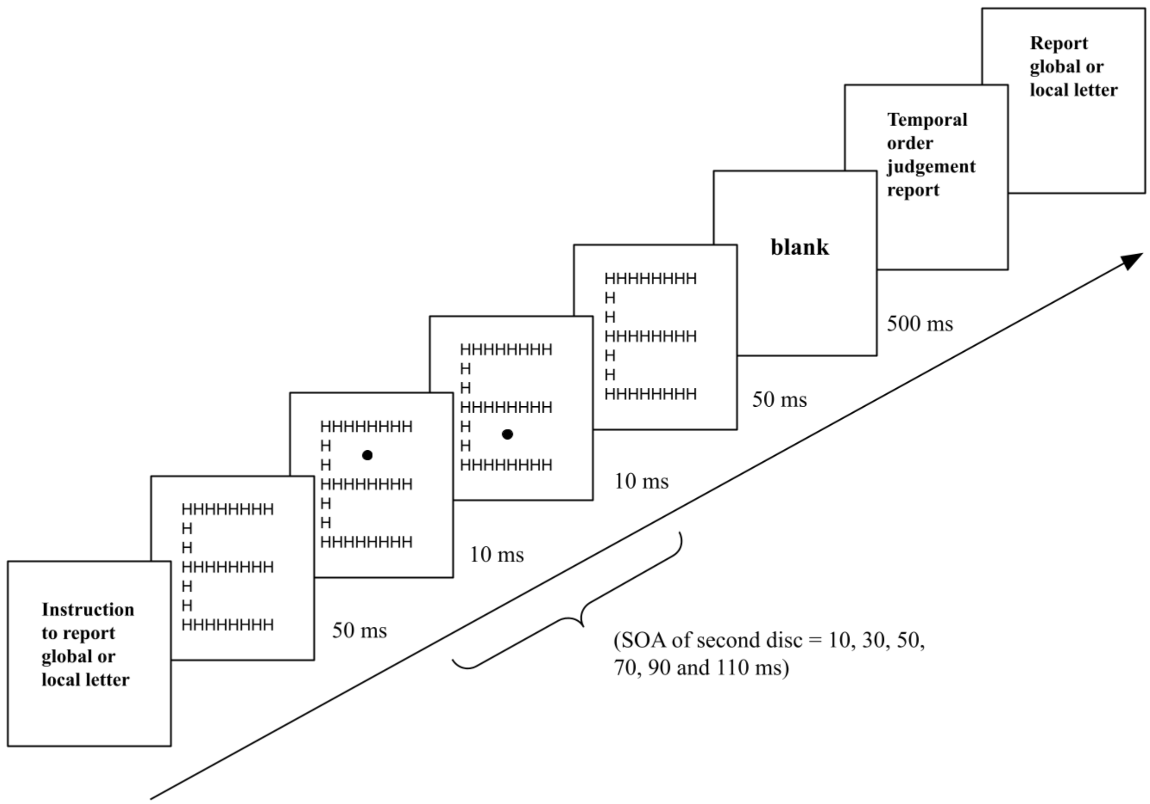

Figure 1.3: Task from Mudumba and Srinivasan (2021).

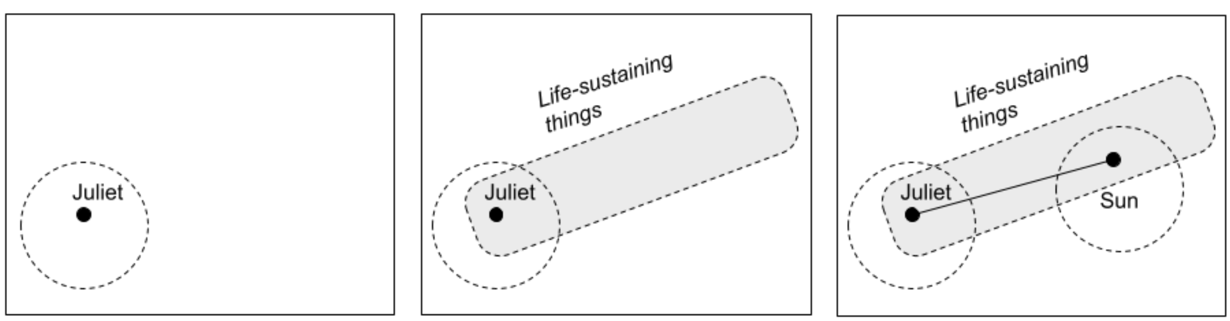

Figure 2.1: Using abstract features represented in gist to guide memory search.

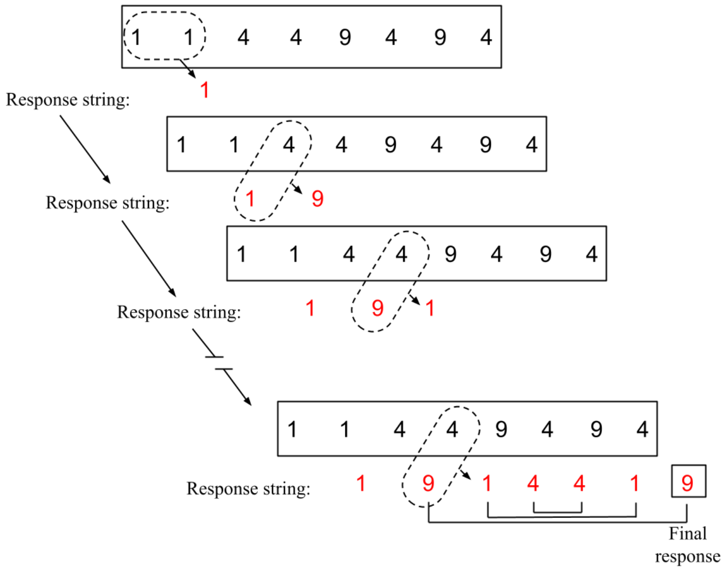

Fig. 2.2: Illustration of the number reduction task from Wagner et al. (2004)

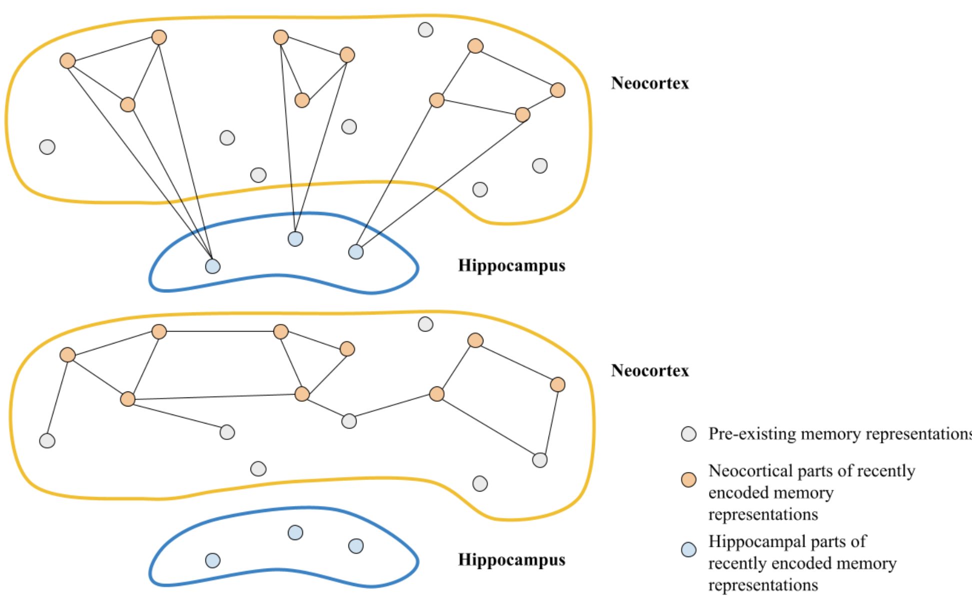

Figure 3.1: An illustration of reactivation of hippocampus and neocortex during memory consolidation.

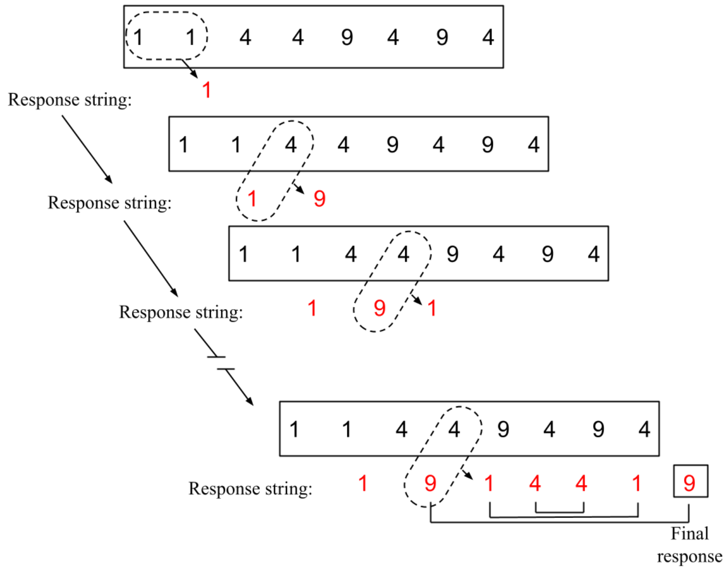

Figure 3.2: Illustration of the number reduction task from Walker et al. (2004). The string to be processed in this case is "11449494", and the response string is "1914419", with the last digit "9" being the final response.