

Decoding attended and unattended visual information in visual working memory

Exploring the effects of attention on visual working memory storage

By

Femke R. A. M. Ruijs

Utrecht University

Thesis for the Major Research Project

Neuroscience & Cognition Masters Program

Supervised by: Dr. Surya Gayet PhD

20-02-2023

Abstract

Our visual working memory can store multiple items simultaneously, even if only one is currently behaviourally relevant. This currently relevant memory item is held in VWM in an attended state, while others are held in an unattended state. Some studies suggest that these attended and unattended memory items can be stored in different brain regions, or even in different formats. However, research is inconclusive about where and how these differences exist across the cortex. In this study we aimed to get insight into where and how the brain distributes neural resources to store attended and unattended visual information. Using 7T fMRI, we examined the neural patterns of three participants while they performed a visual working memory task. Using support vector machine classification (SVM) we were able to decode both attended and unattended orientations held in VWM in occipital and parietal regions. We found attended visual information being represented more in the occipital lobe compared to unattended visual information. Additionally, we found some evidence which suggests attended and unattended orientations can have different neural representations during VWM storage. We conclude that while both attended and unattended memory items can be represented in similar brain regions, attended items are represented stronger in the occipital lobe. This could enable higher resolution, visual-like representations for attended memory items which facilitate comparison with concurrent visual input. Finally, we highlight the benefit of using 7T fMRI and many trials per participant in decoding visual information during visual working memory retention.

Introduction

The outside world is complex and ever-changing. As we navigate ourselves through our daily lives, we process countless objects and scenes as well as visualizing them to perform a task or movement. Think of looking for your keys, where you hold their image in your mind while you scan your surroundings for a match. Our visual working memory (VWM) is an essential process that is tasked with retaining this kind, and many other kinds, of visual information in the absence of the direct visual input stemming from the item (Baddeley, 1998; Ungerleider et al., 1998). Using this, we can seamlessly interact with our surroundings, and perform both simple and complex tasks.

Visual working memory can utilise different neural resources across the cortex to efficiently store different kinds of visual information. For example, stored visual items that are currently task-relevant could be stored differently than items that might become relevant later. However, it is unclear what these differences are exactly, as different studies point to different regions and effects. In this study, we aim to understand where visual information is stored in visual working memory, and whether (and how) this differs between items that are currently task-relevant and items that are not.

Locating visual information storage in VWM.

While visual working memory plays an essential role in our functioning, there is much debate about where exactly the brain stores this visual information. Early fMRI and PET studies mostly point to the prefrontal cortex as the main driver of working memory, as activity increases the most here during VWM retention (Wager & Smith, 2003). In some cases, other regions show increased activity as well, such as parietal, occipital and temporal areas. However, there is high variability in which areas are found to be contributing between articles. Visual working memory appears to be a complex process capable of using many brain regions around the cortex in order to function.

A caveat of these studies pointing to the prefrontal cortex for storing visual information is that they are based on locating brain regions by looking at increased overall activity during memory tasks. However, brain regions that have increased activity during working memory tasks do not necessarily contain task-relevant visual information (Riggall & Postle, 2012). In VWM tasks, many processes are active, such as attention, planning and other executive functions, which are often managed in the prefrontal cortex (Eriksson et al., 2015; Funahashi & Andreau, 2013; Haxby et al., 2000). This implies that increased activity in

prefrontal areas during VWM tasks could just reflect WM supporting processes, rather than visual information storage (Haxby et al. 2000).

Conversely, brain regions that do not necessarily become more active during VWM retention can store visual information. Harrison and Tong (2009) found item-specific patterns in V1-4 even if overall activity had returned to baseline. Using fMRI, they were able to decode the orientation a participant was retaining in VWM through neural activity patterns alone. Brain regions can thus efficiently store information in VWM without increasing overall activity, but more by tuning the activity itself to be more informative. In fact, it is suggested that task performance is more linked to specificity within neural patterns, rather than activity strength (Ester et al., 2013). Therefore, it might be more beneficial to look at whether neural activity is more attuned to the visual information, instead of whether it increases overall. This may allow brain regions that store visual information without becoming more active to be uncovered.

Later studies using decoding methods consistently find evidence for visual information storage in early visual cortex (EVC) in VWM as opposed to prefrontal areas (Pratte & Tong, 2014; Riggall & Postle, 2012; Serences, 2009). In fact, the EVC holds visual information in many different VWM tasks, such as motion, colour, orientation, and location (Christophel et al., 2018; Riggall & Postle, 2012; Serences et al., 2009; Sprague et al., 2014). Furthermore, the neural activity generated by holding an item in memory is similar to that of viewing that item, indicating that the images held in EVC by VWM are perception-like in nature (Albers et al., 2013; Serences et al., 2009). This was found using cross-decoding. Here, they trained a decoding model on brain data collected while a participant viewed a stimulus, and tested it to decode the same stimulus held in VWM. Since they were able to decode the stimulus in VWM using this model, they concluded that the visual cortex exhibited similar patterns during perception and VWM retention. So, not only can visual information during VWM storage be found in EVC, the neural activity here is similar to when simply viewing that item.

While visual information is persistently found in EVC in these VWM studies, there is much more variability in which other regions researchers are able to decode visual information. Some studies were able to decode visual information from parietal regions, such as the intraparietal sulcus (Christophel et al., 2012; Ester et al., 2015; Rademaker et al., 2019), and the frontal eye fields (Christophel et al., 2018; Ester et al., 2015). Some are even able to decode information from the frontal cortex, where others only find increased activation but no decoding (Jerde et al., 2012; Sprague et al., 2012). These differences in where visual

information can be decoded indicate the brain does not distribute the same brain resources to each item, instead taking on a more specialized approach.

The differences in where visual information is found during VWM retention appears to be driven by the relevant feature in a specific task (Lee & Baker, 2016). Here, a specific aspect (feature) of a memory item, such as its colour, orientation, or size influences which brain resources are used to store the information. For example, in motion-based tasks, task relevant information can be decoded in the visual motion area (MT), an area that is used to detect motion (Emrich et al., 2013; Riggall & Postle, 2012; Tootell et al., 1995). Similarly, task-relevant information that is normally processed in posterior parietal regions (e.g. colour, orientation) can also be decoded from this region during VWM retention (Christophel et al., 2012; Sprague et al., 2014). Furthermore, while visual information often cannot be decoded from prefrontal areas, non-visual information, such as category or name, can (Lee et al., 2013; Riggall & Postle, 2012). The neural resources used in VWM to retain visual and non-visual information thus seem to differ between features and tasks. This division of resources and labour appears to be based on which information is task-relevant and which brain regions are functionally able to support the retention of that information.

Item-relevance in visual working memory.

When holding a single item in VWM, which features are task-relevant might influence which neural resources are used to hold that information (Lee & Baker, 2016; Riggall & Postle, 2012). However, we often hold multiple items in VWM, like when remembering multiple digits of a phone number or looking for both cheese and milk in a grocery store (Fukada et al., 2010). In the latter example, it is possible for one item to be prioritised over the other in VWM. It is suggested that when searching for multiple items, only one mnemonic template can be ‘active’ and guide visual search (Frătescu et al., 2019; Moorselaar et al., 2014). Thus, when looking for cheese and milk, the template of only one of those items can be actively prioritised in VWM, and you would need to switch in between them to search for both. On the other hand, there is some evidence that multiple mnemonic templates can be active at a given time (Hollingworth & Beck, 2016). In this instance, the templates for both cheese and milk could guide visual search simultaneously. Hence, while multiple items can be held in visual working memory concurrently, it is possible that one item can be held in a more active or prioritised state (Olivers et al., 2011).

Continuing the example of cheese and milk, there are instances where an item is sure to be prioritised over the other in visual working memory. For example, when you choose to

search for cheese first, and then milk. Here, only one mnemonic template is relevant at a time (first cheese, then milk), while the other is simply stored in VWM for later use (or discarded once the item is found). In this case it is not a feature of an item that is more task-relevant, but an item itself. If the brain allocates resources based on task-relevance this raises the question of whether item relevance also influences where the information from that item is stored.

First we need to understand what happens in visual working memory when an item becomes more relevant. If one item becomes more relevant or important in a specific task we will draw more attention to it (Baluch & Itti, 2011). Attention is often viewed as a process that is used as a ‘spotlight’ on the world around us (Carrasco, 2011; Scholl, 2001). For example, drawing your attention on a specific location in the visual field, or a specific feature, like colour or orientation. However, attention does not only work outwards, but also inwards (Lepsien & Nobre, 2007; Serences et al., 2009). Through attention, one can shine a spotlight on internal items or features. Using the cheese/milk example, if a specific item is more relevant (cheese) in a specific moment, we will inwardly draw more attention to it, without forgetting the other (milk). Thus, if there are differences in how neural resources are distributed between items, attention will be a driving force behind this.

(Un)attended visual information storage in VWM.

In a later study, researchers were able to find differences in where attended and unattended visual information is stored in VWM (Christophel et al., 2018). They used a paradigm where participants had to remember two distinct orientations, after which a retro-cue indicated which of them was to be used in an upcoming change detection task. This meant that one memory item became temporarily more task-relevant than the other. Consequently, the participant prioritised this cued item (attended) in VWM over the un-cued item (unattended) (Kerzel & Witzel, 2019). After they completed this first task, a second retro-cue would indicate which of the items would be used in the next task. Since each of the original items could be cued for the second task, this ensured the participant did not discard the previously un-cued (unattended) item (Christophel et al, 2018). By doing so, both attended and unattended memory items could be studied in the first half of the task. Using this paradigm, the researchers found evidence for both attended and unattended information in the IPS and frontal eye fields, but only information about attended items in the EVC. They concluded that the brain likely distributes neural resources that are able to retain information more precisely to attended items, as opposed to unattended items. This could boost the

performance of actions such as visual search, as it might help comparing visual input to the stored mnemonic representations.

However, a later study was able to find evidence for unattended visual information retention in the EVC (Yu et al., 2020). Using a similar paradigm, they were able to decode both attended and unattended memory items in the EVC. They attribute this finding to a difference in analysis method. The study by Christophel et al (2018) trained and tested their decoding model on data from the same attentional condition. This meant that unattended information was decoded by training and testing on unattended memory items. In contrast, Yu et al. (2020) trained their model on the much stronger represented attended memory items, and then tested it on unattended items. This method allowed for weaker unattended signals to be uncovered. In addition, reanalysis of the data by Christophel et al. (2018) showed that when training data on attended items, and testing them on unattended items, unattended items could be decoded from the EVC (Iamshchinina et al., 2021). It is worth mentioning, however, that the study by Yu et al. (2020) did not find evidence for unattended items in the IPS, regardless of whether they trained on attended or unattended memory items. Thus, the analysis method and design can affect whether and where weaker patterns can be uncovered. Although, even when using a more sensitive analysis method, there are still differences in where unattended information is found in VWM.

Additionally, there also appear to be differences in how visual information is stored in VWM, depending on its attentional status. There is some evidence that attended and unattended visual information can be stored in opposite neural codes (van Loon et al., 2018; Yu et al., 2020). For instance, attended and unattended orientations stored in VWM were found to be represented in opposite neural patterns (rotated 90 degrees) (Yu et al., 2020). Similarly, objects held in VWM were represented in opposite representational patterns in object-selective cortex (van Loon et al., 2018). However, this finding is not robust across all brain regions and types of information (Yu et al., 2020). Still, it gives a unique insight in how the brain might utilize different methods in storing attended and unattended information.

In short, there appear to be differences in how and where attended and unattended visual information is stored in VWM. Some studies show that attended and unattended information is stored in opposite neural representations, though this is not the case for all brain regions and features. Additionally, the location of attended and unattended visual information is debated. It is implied that unattended information can only be decoded from EVC when training on stronger represented attended information. However, even with strong

analysis methods, there are still differences in where unattended visual information can be decoded.

Limitations in previous research.

An important note to keep in mind is that an absence of evidence is not evidence for absence. What the studies by Yu et al. (2020) and Iamshchinina et al. (2021) show us is that finding evidence for inherently weaker patterns can depend on your analysis method. Since, when training a decoding model on stronger attended information, unattended visual information could be decoded from the EVC, which was not possible before. Thus, when trying to decode information from VWM, it is important to use a highly sensitive method. Doing so could reveal weaker mnemonic activity patterns, which would give a clearer picture of where and how visual information is stored in VWM.

We note two limitations in previous research which could have influenced whether and where visual information can be found during VWM retention:

1). Most, if not all the previously mentioned decoding studies use 3-Tesla fMRI to collect their data. However, with the rise of ultra-high-field 7-Tesla fMRI in research, advances could be made here. 7T fMRI is better able to capture changes in BOLD-signal and has a better signal-to-noise ratio than 3T fMRI (Triantafyllou et al., 2005; van der Zwaag et al., 2009). Because of this, 7T fMRI analysis increases statistical strength, allowing detection of weaker patterns (Torissi et al., 2018). Consequently, using 7T fMRI for VWM decoding studies could increase sensitivity to weaker patterns, which normally would be a bottleneck for finding unattended or otherwise muted mnemonic representations.

Within the usage of 7T fMRI, it might also be beneficial to use a within-subjects analysis with many trials per participant, as opposed to less trials with many participants. This is because visual working memory response patterns are unique for each individual. Uncovering more muted response patterns can therefore be optimised by using many trials for each participant (Gordon et al., 2017). By doing so, we can get a detailed display of response patterns and might be able to decode information in regions previously impossible. Hence, using less participants but more trials per participant (across multiple sessions) could allow for more specific and detailed estimations of where and how visual information is stored in VWM.

2). Furthermore, many studies make use of regions of interest (ROI's) from which they decode their information (Christophel et al., 2018; Yu et al., 2020). These are previously defined voxels that are used to represent the activation of a specific brain region. As a

consequence, the analysis is limited to the previously defined regions that will be examined. This comes with several limitations. Firstly, it is possible that there are other regions that store visual information in VWM that are not detected as they are not analysed in the first place. Secondly, brain regions can be quite heterogenous, with some voxels containing relevant information and others not. These non-informative regions then introduce non-relevant noise, which decreases the chances of finding an effect. Finally, these ROIs are often defined by taking the n-most active voxels during a measuring delay. This poses two problems. Firstly, higher neural activation is not necessarily correlated with higher visual information decoding (Riggall & Postle, 2012). Therefore, the voxels that contain information might not be selected through this process. Secondly, you run the risk of ‘double-dipping’ which can artificially increase decoding strength and decrease reliability (Kriegeskorte et al., 2009). Hence, when looking at where visual information is stored in visual working memory, using a non-ROI based approach could prove beneficial. Using this method, regions are not predetermined based on previous research and voxel selection methods are less likely to influence results.

Therefore, with 7T fMRI and a non-ROI based approach, the detection of activation patterns across the whole cortex could be enhanced. In visual working memory research, this would increase the chances of finding weaker mnemonic representations in different attentional conditions. It would thus provide a clearer picture of where visual information is stored in visual working memory.

The present study

In this study we aim to use these methods to answer the question; *Where in the cortex is visual orientation information stored during visual working memory retention?* Using different attentional conditions we can expand that question and discover; *Are there differences in where attended and unattended memory items are stored in visual working memory?* If we find both attended and unattended visual information stored in VWM, we could examine this further and answer the question; *Does the brain store attended and unattended information differently?* If we find differences in where and/or how unattended visual information is stored compared to attended, it might be interesting to see how this affects behaviour. So, we also want to examine; *Does visual information being stored in an unattended state affect task performance?*

To answer these questions a paradigm similar to that of Christophel et al. (2018) is used. Here, two items (orientations) are held in visual working memory, and a retro-cue indicates which of the two is more relevant in an upcoming task. This leads to the cued item

getting more attention (attended) than the un-cued item (unattended). After this, another cue indicates which of the items will be relevant in the next task. This can be the same item as before (repeat) or the other (switch). In the repeat condition, a previously attended item is attended again. In the switch condition a previously unattended item becomes attended. Using fMRI to decode neural activity patterns in these conditions, we can decipher where attended and unattended visual information is stored. Furthermore, it is possible to decode where a previously unattended, but now attended item is stored. This allows us to fully understand how attention influences visual working memory storage.

Additionally, we will focus on within-subjects statistics as opposed to group-level statistics. To this end, we had each participant complete many trials, thus increasing the statistical power for uncovering memory representations. The goal is to get a complete and unique insight of where and how visual information is stored in the brain.

For this study, will take on a hypothesis-free approach. That is, we will not search for VWM storage in specific predetermined brain regions, but across the whole cortex. As a result, no regions are excluded, providing us with the broadest description of where visual information is stored during visual working memory retention. Thus, using this method allows our research questions to be answered with as little bias as possible.

Hypotheses and expectations

Where in the cortex is visual orientation information stored during visual working memory retention?

Whilst we do not pre-select our analysis regions, we have some expectations of where we will find attended and unattended visual information. Attended memorised orientations have previously been decoded from the EVC and IPS during visual working memory retention (Christophel et al., 2012; Christophel et al., 2018; Rademaker et al., 2019).

Therefore, we expect to find information about the attended item in the first delay in occipital and parietal regions. Regarding unattended memory items, there is some evidence that indicates unattended memorised orientations can also be stored in the EVC and IPS (Christophel et al., Yu et al., 2020). However, neither study showed evidence for unattended memory items in both regions. Using our more sensitive methods, we expect to find some evidence for unattended memory items in the parietal and the occipital cortex.

Are there differences in where attended and unattended memory items are stored in visual working memory?

Our expectations regarding whether attended and unattended memory items are stored in different brain regions are less concrete. While there is much evidence for attended orientations held in VWM being stored in the EVC and IPS, the support for unattended orientations being stored here is weaker. It is possible that both are represented in both regions, but one more strongly than the other, or in a larger or different sub-region. This provides us with three possibilities regarding the effect of attention on the location of visual information storage.

1: There is no difference in which regions store attended and unattended visual information. Here we would see significant decoding accuracies in the same brain regions for both attended and unattended memory items. We would then also find no statistical difference between the decoding accuracies of either attentional condition. Since previous studies on attended and unattended memory items make use of ROI's, and we will analyse the whole cortex, it is difficult to assess whether this would be consistent with previous studies. Since, even if some studies found no difference between attended and unattended information in their ROI's, it is possible that there were differences outside those ROI's. However, considering unattended memory item representations are weaker than attended ones (van Loon, 2018) and are often not found in the same regions as attended information (Christophel et al., 2018; Yu et al., 2020) this option would be unlikely.

2: Attended memory items are stored in different brain regions than unattended memory items. Here we might observe some overlap in where attended and unattended information is stored, but we would find decoding for attended but not unattended items in one or more area(s). Additionally, we would observe significantly more decoding strength for attended items than unattended items within these regions. This would be in line with previous findings where attended information is found in more regions than unattended information (Christophel et al., 2018; Yu et al., 2020).

3: Unattended memory items are stored in different brain regions than attended memory items. Here we might observe some overlap in attended and unattended information storage, but significant decoding for unattended and not attended items in one or more area(s). Additionally, we would observe significantly more decoding strength for unattended than attended items within these regions. This would be a novel finding, made possible by non-ROI based analysis. It would imply that there are brain regions specifically utilised for storing unattended or currently irrelevant visual information.

Does the brain store attended and unattended information differently?

Some studies suggest that attended and unattended visual information can be represented differently in occipital regions (van Loon, 2018; Yu et al., 2020). Therefore, we expect to see some regions where attended and unattended visual information is stored differently in VWM. For our results, this would mean finding limited or ideally negative cross-decoding between attended and unattended memory items in occipital (or other) regions. Here, negative cross-decoding would indicate that unattended and attended mnemonic representations are opposite from one another.

Does visual information being stored in an unattended state affect task performance?

In addition to our fMRI experiment, we setup a behavioural experiment to study the behavioural effects of storing visual information in an unattended state. This is because while a within-subjects analysis with few participants is ideal for collecting detailed data on neural response patterns, it does not lend itself to making inferences about behavioural effects. This is because our experimental design only results in two behavioural data points per trial. Therefore, we would need more behavioural data than we could collect from the small participant group of our fMRI experiment to make statistical inferences. Therefore, we opted to set up a behavioural experiment with more participants than in our fMRI experiment, so that we could gain insight into whether storing information in an attended or unattended state affects task performance.

If unattended memory items are less strongly represented than attended memory items, we expect task performance to decrease if a participant has to use a memory item that was previously stored unattended, as opposed to attended. This would be in line with a previous study (Christophel et al., 2018). Here, they found that task performance decreased significantly when a target item was previously held in an unattended state, compared to items stored only in an attended state.

Methods

Two experiments were conducted. In experiment 1, each participant completed the experiment in multiple sessions while 7T fMRI and behavioural data was collected. In experiment 2, only behavioural data was collected.

Participants

For experiment 1, we aimed to collect detailed neural responses in each participant, and therefore opted for many trials per participant as opposed to less trials with more participants (Gordon et al., 2017). For experiment 1, three right-handed participants were recruited (one female), aged 23-40. Each participant had been in an fMRI at least once before. One participant received additional monetary compensation for participating. All scanning sessions were performed in the Spinoza Centre for Neuroimaging in Amsterdam.

To investigate the behavioural effects of the experiment, we setup another experiment with fewer trials per participant, but more participants (Experiment 2). For this experiment, 21 people participated (15 female), aged 19-32 (s.d.=3.45). One participant (male, aged 32) was excluded due to extremely poor performance (s.d. of mean error compared to other participants > 3). Most participants completed the experiment in one session, though two had to return for a second session due to hardware malfunctions. Participants received either monetary compensation or ‘participation hours’, a form of credit Psychology students at the University of Utrecht need to collect to complete their bachelor’s degree. This experiment was conducted in the Langeveld Building at the University of Utrecht, The Netherlands. Written informed consent for each participant was acquired before participation.

Procedure

Experiment 1

The stimuli were projected on a 698x393 mm LCD- screen (BOLDscreen 32, Cambridge Research Systems, Rochester, UK) with a resolution of 1920x1080 pixels, RGB colour, a contrast ratio of 1400:1, and a framerate of 120 Hz. The screen was positioned 1850 mm from the participant’s eyes, and visible through a mirror placed on the participant’s head coil. All stimuli, tasks, and conditions were created and presented using MATLAB and Psychtoolbox-3 (version 3.0.17).

For each trial, participants completed two consecutive delayed match-to-sample (DMTS) tasks. During the entire experiment, a fixation dot was visible. Before the start of each trial, this fixation dot turned blue (rgb= [65 105 255]) for 1000ms indicating that a new trial was starting. After another 1000ms, two donut-shaped Gabors with different orientation gratings were shown successively, with a fixation screen in between (800ms each). After this, a donut-shaped mask with a numerical retro-cue in the centre (either 1 or 2) was shown for 800ms, indicating which of the previously shown orientations was to be used in the upcoming task. After a delay of 8000ms, a randomly oriented line appeared. Here, participants had 3000ms to

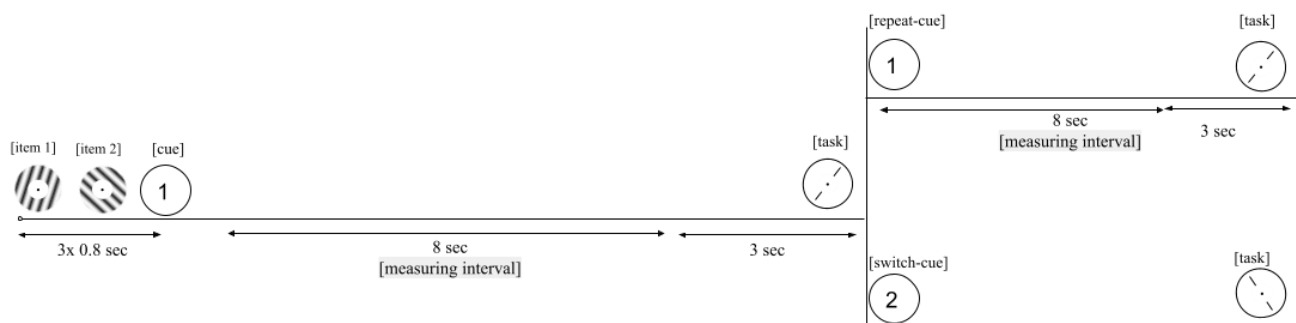
match this line to the remembered orientation of the cued stimulus. This was done using a rotary dial (Curdes HHSC-1x1-DIAL). After the task period ended, the screen ended automatically and the last shown orientation was saved as the response.

Immediately after, a second retro-cue was shown, again indicating which of the two originally shown orientations was to be used in the upcoming task. This could either be the same orientation as in the first delay (repeat trial) or the other orientation (switch trial). The task was repeated for this second cued item following another 8000ms delay. After this, the trial ended. Trial setup can be found in Figure 1.

The experiment contained 27 blocks of 12 trials each. After each block, the participant was shown the mean absolute response error (in degrees) for their first and second response for that block. Between each experimental block, the participant had the opportunity to communicate with the fMRI technician in the other room. Each participant completed all 27 blocks over 3-4 separate scanning sessions. At the beginning of each scanning sessions, participants completed practice trials, where they were shown the correct orientation after each response.

Figure 1.

Trial setup diagram.



Note. Walkthrough of a single trial showing the two conditional possibilities of the second retro-cue. Target stimuli were shown for 800ms each. Each retro-cue was accompanied by a donut-shaped noise mask (not pictured). After the first task, the second retro-cue could either cue the same item as the first retro-cue (repeat trial), or a different one (switch trial).

Experiment 2

Stimuli were projected on a 651x363mm screen, with a resolution of 1920x1080, and a framerate of 60Hz. Stimuli were created and presented using MATLAB 2021b and Psychtoolbox-3 (version 3.0.17) in a darkened room. The participant viewed the screen from a headrest positioned 570mm away. The researcher was present during the entire experiment to observe and answer possible questions.

The trials in this experiment were the same as experiment 1, with two exceptions. Firstly, the delay period between each cue and response was 4000ms (instead of 8000ms). This is because the behavioural experiment did not require measuring a delay-specific BOLD response. Secondly, the rotary dial could only alter the response by 56 degrees per complete dial-rotation, and therefore needed to be rotated more than the rotary dial used in the fMRI experiment.

This behavioural version of the experiment contained ten experimental blocks of 12 trials each, as well as a practice block with immediate response feedback. The participant could take breaks between each block. All other aspects of the trial were remained the same as in experiment 1.

Conditions

Each trial contained 2 retro-cues, followed by a task. Since each retro-cue could cue either of two orientations, this gave us 4 experimental conditions:

- 1: The first shown item was cued in both the first and second delay (1-1).
- 2: The first shown item was cued in the first delay, and the second shown item was cued in the second delay (1-2).
- 3: The second shown item was cued in both the first and the second delay (2-2).
- 4: The second shown item was cued in the first delay, and the first shown item was cued in the second delay (2-1).

In the first delay, the cued item was prioritised and becomes the attended memory item (AMI), while the other was not prioritised but still remembered (unattended memory item – UMI) (Christophel et al., 2018).

In condition 1 and 3, the same item was cued in both delays (repeat trial). As a result, a previously attended memory item became attended again in this second delay (previously attended memory item – PAMI). In condition 2 and 4, one item was attended in the first delay, and the other in the second delay (switch trial). Consequently, a previously unattended memory item in the first delay became attended in the second delay (previously unattended memory item – PUMI). Since in all conditions, the un-cued item in the second delay could no longer become relevant and could be discarded from VWM, it was not included in our analysis.

Stimuli

For both experiments, all stimuli were specified in dva (degrees visual angle) and then converted into pixels based on pixel size and distance from the presentation screen. All

stimuli were presented against a grey (rgb=3x128) background. A bullet point fixation dot was visible in the middle of the screen for the entire duration of the experiment (black outer circle (rgb=3x0, 0.2 dva=0.2) & white inner circle (rgb=3x255, dva=0.1). The only time the fixation dot was not visible was during the presentation of a numerical cue (font=Calibri, font-size=60 pixels).

All other stimuli were presented in a donut-shape around the fixation dot (dva=5-by-5) with a cosine blur around all edges of 0.5 dva. The memory targets were sinusoidal Gabors gratings created using a 0.6 grating contrast, a target size of 1.2 dva (10 cycles per dva) and cosine blur of 0.5 dva, after which the whole Gabor was enlarged to five dva. Masks were created by creating 100x100 pixel brown noise and enlarging to 5 dva.

The two orientations in each trial were constrained to be at least six degrees apart, to ensure they were perceived as different. The previously mentioned four conditions-of-interest were also balanced to be equally present in each block, meaning each condition was presented (12/4 =) 3 times each block.

For experiment 1, target orientations were computed by dividing the 180° orientational space by 81, creating 81 equally spaced item orientations. These item orientations were counterbalanced across the four conditions, as well as the two possible presentation positions (the first or second shown Gabor in the sequence) across the whole experiment.

For experiment 2, the same 81 possible orientations were used. However, since the experiment had fewer trials, it was not possible to counterbalance all 81 orientations for each participant. Therefore, for each participant, 60 of the possible 81 orientations were randomly selected to appear in the experiment. These 60 orientations were then counterbalanced across all four conditions over the whole experiment, but not the possible presentation positions.

Data-processing

Experiment 1

MRI data acquisition: (f)MRI data was acquired on a 7T Philips Achieve scanner. We acquired T2-weighted functional images using a 32-channel head coil with a resolution of 1.688x1.688x1.7 mm. with 57 slices containing 128x128 voxels, which were acquired for each TR lasting 1500ms. A flip angle of 70 degrees was used, and a TE of 22.49 seconds. A gradient echo sequence with a SENSE acceleration factor of 2.4, multiband factor 3 and anterior-posterior encoding was used. A second order B0 based shim was used. This encompassed most of the brain except the cerebellum and parts of the anterior temporal lobes.

Each scan contained 244 TR's, resulting in a total scan time of 366 seconds for each block. At the start of each block, a delay of 16 seconds was included, to return the participants BOLD-signal to baseline before the first trial started. At the end of each block a delay of 8 seconds was included to pick up remaining BOLD-signal from the last trial. Aside from the experimental functional scans, we acquired two top-up functional scans per session, which were recorded using the opposite phase-encoding direction from the functional runs to correct for image distortion.

We acquired a high resolution T1-weighted scan in the first scanning session for two participants (one was already available for the third). This scan was automatically segmented using Freesurfer, after which labels were manually edited to minimize segmentation errors using ITK-SNAP. This gave us a cortical surface model with grey-white segmentations to be used as masks. We also collected several anatomical T1 scans in the same resolution as the functional scans.

fMRI pre-processing: fMRI pre-processing was performed using a pipeline as described in Hendrikx et al. (2022) in AFNI (Cox, 1996). A transformation matrix was constructed, which contained all necessary steps from the raw data to the final co-registered images, reducing the number of interpolation steps reduced to one. We used motion correction on the functional images and co-registered them to the high-resolution T1 scan we acquired.

Functional data motion correction was first applied to two image series that were gathered using opposing gradient encoding directions. Afterwards, the distortion transformation between the average images of these two image series were calculated. Then, brain position transformations between and within functional scans were calculated.

To determine the transformation from the low-resolution functional images to the high-resolution whole-brain T1 image, a low-resolution T1 image with the same position and orientation as the functional images was used. Next, the transformation that co-registers the functional data to the same-space T1 was determined.

The product of all these transformations was then applied to the functional data to transform the data to the same space as the high-resolution T1 scan. This was repeated for each fMRI session, so all scans were in the same high-resolution space. No other spatial or temporal smoothing was applied.

Within-subjects statistical analysis design: For experiment 1, we performed a within-participants statistical analysis. This is because we used a small amount of participants (N=3),

but collected a large amount of data per participant. Therefore, all data selection, analysis and significance testing was done for each participant separately.

fMRI data selection: After pre-processing the fMRI data, we timestamped the moment the cue was presented as the start of the memory delay, as from this point the participant would start prioritising one stimulus over the other in visual working memory. We then timestamped 9 seconds after cue presentation as the end of the measuring delay, as this would encompass the whole period in which the participant was actively retaining one stimulus over the other until response. We then added 4.25 seconds to both timepoints to account for the delay in BOLD-signal to ensure we got the most task-relevant activity. This was repeated for both delays in each trial.

To calculate mean voxel activity for each measuring delay, the measuring delay start and end was rounded to the nearest TR, resulting in six images for each delay. We then calculated the mean over these six images, resulting in a single mean activation image (i.e., one mean per voxel) for each measuring period. These images were then normalised by z-scoring all mean activation images within each run. This was done to reduce the effect of voxel signal differences caused by physical or physiological scanning differences between runs. Consequently, each value indicated the mean activation height for each voxel normalised for each run. This was repeated for all 27 runs, resulting in 648 (324*2) z-scored mean activation images, one for each delay in each trial.

Trial selection: We removed responses where the absolute response error was >2.5 s.d. removed from the mean absolute error for that participant. This was done to ensure only delays where the participant attended the correct item were included, since our goal is to decode this from VWM.

Voxel selection: Since we are using a hypothesis-free approach in locating visual information in the cortex, we are interested in all voxels that might hold visual information. Therefore, a grey-matter mask from the segmentation as described above was used. This included all grey matter voxels for each participant.

Voxel-based analysis: To identify voxels containing visual information, a searchlight method was used (Etzel et al, 2013; Kriegeskorte et al., 2006). Using CoSMoMVPA's *cosmo_searchlight* function in MATLAB we could analyse small voxel neighbourhoods

across the cortex (Oosterhof et al., 2016). Neighbourhoods with a radius of three were created using *cosmo_spherical_neighborhood*. Depending on the participant, this resulted in neighbourhoods containing an average of 75-95 voxels.

Support Vector Machine Classification (csvm): To decode visual information in the cortex, a Support Vector Machine (SVM) was used (Mahmoudi et al., 2012; Noble, 2006). For each condition (AMI-UMI-PAMI-PUMI), trials were divided into orientation bins. They were first divided into horizontal ($<67.5^\circ$ and $\geq 112.5^\circ$) and vertical ($<22.5^\circ$ or $\geq 157^\circ$) trials (Rademaker et al., 2019). We used a cross-validation method where we iteratively took out one trial at a time to be used as a test trial, and all other trials were used as training trials. Training was done using MATLAB's built-in SVM function *fitcsvm*. To ensure the training bins held an equal number of trials in both classes, all items from the smallest bin were used, after which the same amount of trials were randomly selected from the other orientation bin.

Using MATLAB's *predict* function, we tested the model on the test item to get a classification score. This indicates the distance from the decision boundary used to predict in which class the test item belonged. A positive score indicates the item is predicted to be in one class, and a negative score indicates the item is predicted to be in the other class. The absolute height of the score itself indicates how sure the model is that an item belongs or does not belong in a specific class. We multiplied this classification score with the class indicators (e.g., -1 for horizontal and 1 for vertical). By doing so, we obtained distance-to-bound values where a positive number indicated the model correctly predicted which class a test item belonged to, and a negative number indicated the model incorrectly predicted which class it belonged to.

The remaining trials were then divided into diagonal 1 ($<22.5^\circ$ and $\geq 67.5^\circ$), and diagonal 2 ($<112.5^\circ$ and $\geq 157.5^\circ$) bins and the SVM process was repeated.

This cross-validation method for the SVM was executed using CoSMoMVPA's searchlight function. Here, a searchlight sphere was created around each voxel in the grey matter mask. Then, for each searchlight sphere, the SVM analysis was executed using only the data from the voxels in the sphere. The distance-to-bound value of each sphere was then projected to the centre voxel of the sphere. This resulted in a distance-to-bound value for each voxel in each trial. Voxels from the same trial were then converted into a single NIfTI image, resulting in a distance-to-bound image for each trial in each condition.

Cross-decoding: To analyse whether concurrently held information is stored in the same format, we used a similar SVM classifying mechanism as described above to cross-

decode the AMI and UMI. We split the data for the AMI and UMI into orientation bins (horizontal/vertical and diagonal1/diagonal2). Instead of taking a leave-one-out method to train and test the data we trained on one dataset and tested on the other. This means for both orientation classifiers, we trained on the AMI and tested on the UMI (AMI→UMI), and the other way around (UMI→AMI).

Significance testing: To test whether which brain regions contained a significant amount of visual information, CoSMoMVPA's *cosmo_montecarlo_cluster_stat* function was used. This function uses threshold-free cluster enhancement (TFCE) along with Monte Carlo simulations to compute cluster-statistics corrected for multiple comparisons (Oosterhof et al., 2016; Pernet et al., 2015). TFCE computes a score for each voxel which represents cluster-based spatial support (Smith & Nichols, 2009). Consequently, high-value voxels surrounded by high-value voxels get a higher score, as there is much cluster-based support for a high value. On the other hand, high-value voxels surrounded by low-value voxels get a relatively low score, as there is little cluster-based support for a higher value. In fMRI, this type of cluster enhancement is useful as it boosts the strength of extended areas of signal. Thus, the signal of brain regions that hold visual information is enhanced, while that of single voxel-outliers is diminished.

For within-subject analysis, we took all SVM distance-to-bound trial maps as our real data. Then, for each attentional condition (AMI-UMI-PAMI-PUMI), and our cross-decoding maps (AMI→UMI and UMI→AMI), we computed random-effect cluster statistics with a null-hypothesis mean of zero and 10.000 iterations. This entailed randomly flipping the signs of the distance-to-bound maps for each iteration, to preserve spatial smoothness and autocorrelations between voxels. This resulted in a single map containing a two-tailed z-score for each voxel over all trials, which could then be converted into p-values.

Since the *cosmo_montecarlo_cluster_stat* analysis included corrections for multiple comparisons, an alpha of 0.05 was used. Consequently, an absolute z-score of 1.96 ($p<.05$) was used as a significance threshold. Here, any voxel containing a value above 1.96 indicated significant positive decoding for that region, while any voxel containing a value below -1.96 ($p<.05$) indicated significant negative decoding for that region. Though each value above this point is statistically significant, we further distinguish absolute z-scores above 3.3 ($p<.001$). This is done descriptively to indicate where significance is much higher than our threshold of $p<0.05$.

To compare decoding strength between attended and unattended memory items, we used a paired-sample t-test with unknown variance (without correcting for multiple

comparisons). For each voxel, we compared the SVM distance-distance-to-bound values of all trials between the AMI and the UMI. We repeated this for every participant where we found significant decoding of either the AMI or the UMI.

Experiment 2

For the behavioural experiment, we first removed all switch errors and outliers. Switch errors occur when the participant reports the un-cued item as opposed to the cued item, ‘switching’ them. These errors were removed by selecting responses where the response was closer to the un-cued item than the cued item, and the difference between the un-cued and cued orientation was larger than 60. After this, outliers were removed by removing all responses that were >2.5 s.d. from the mean.

The data was then separated into three conditions. The first delay, which contained all responses from the first task in each trial. The repeat delay, which contained all responses in the second task, where the cued item was the same in both the first and second task. The switch delay, which contained all responses in the second task, where the cued item was different than in the first task. Using one-way ANOVA, we tested whether there were significant differences between the absolute error means in these three conditions. If found, they were then further analysed using Tukey-Kramer post-hoc tests.

Results

Experiment 1

Participant 1:

In the first delay, we were able to decode the AMI across the occipital and parietal lobe ($z > 3.3$, $p < .001$). We did not find significant decoding for the UMI. So, for this participant, we found information about attended memory items in the occipital and parietal lobe, but we did not find unattended information being stored anywhere. Using paired-sample t-testing, we found that 59.6% of voxels that contained significant information about the AMI, contained significantly more information about the AMI than the UMI in the occipital and parietal lobe ($t(316) > 1.97$, $p < .05$ (two-tailed)) (Figure 7). This implies that for this participant, these brain regions store more information about attended than unattended memory items when held in VWM. There was no significant cross-decoding between the AMI and the UMI in either direction. In fact, we found negative cross-decoding in occipital and frontal regions when training on the UMI and testing on the AMI ($z < -1.96$, $p < .05$) (Figure 3). In other words,

when training our model on unattended representations, and testing them on attended representations, the model performed significantly below chance. This could indicate that attended and unattended information are represented differently in VWM.

In the second delay, the PAMI could be decoded from the parietal and frontal lobe, as well as part of the occipital lobe ($z > 1.96$, $p < .05$ (two-tailed)). The PUMI could be decoded from the occipital lobe ($z > 1.96$, $p < .05$ (two-tailed)) (Figure 2). Thus, in the second delay, we found information about previously attended (and currently attended) memory items in occipital, parietal, and frontal regions. We also found information for previously unattended (but now attended) the occipital lobe.

Participant 2:

In the first delay, we could not decode the AMI or UMI. Thus, for this participant, we were not able to find information about either attended or unattended items held in VWM. We found no significant cross-decoding between the AMI and UMI in either direction. So, when training on one attentional condition and testing on the other, we were also unable to find any information about either memory item.

In the second delay, we found significant decoding for the PAMI in the occipital, parietal, and frontal lobe ($z > 3.3$, $p < .001$) and significant decoding for the PUMI in the frontal lobe ($z > 1.96$, $p < .05$) (Figure 4). This means that for this participant, we found information about previously attended memory items in multiple brain regions, and information about previously unattended memory items in the frontal lobe.

Participant 3:

In the first delay, we found significant decoding of the AMI in the occipital and parietal lobe ($z > 3.3$, $p < .001$ (two-tailed)), as well as the frontal lobe ($z > 1.96$, $p < .05$ (two-tailed)). We also found significant decoding for the UMI in the parietal and occipital lobe ($z > 1.96$, $p < .05$ (two-tailed)) (Figure 5). We thus found information about both attended and unattended memory items stored in VWM in occipital and parietal regions. We also found information about attended orientations in frontal regions.

Further analysis showed overlap between voxels that held significant information about the AMI and the UMI. Here, 21.6% of voxels that held information about the AMI, also held information about the UMI, and 25.9% of voxels that held information about the UMI also held information about the AMI. These voxels could be found in both the parietal and occipital lobe (Figure 6).

We found that 34% of voxels that contained significant information about the AMI, also contained significantly more information about the AMI than the UMI. These voxels could be found mostly in occipital, but also parietal and frontal regions ($t(319) > 1.97$, $p < .05$). On the other hand, 25.4% of voxels that contained significant information about the UMI also contained significantly more information about the UMI than the AMI. These voxels could be found in both occipital and parietal regions ($t(319) > 1.97$, $p < .05$) (Figure 7). This means that there are brain regions that store more attended information than unattended, as well as regions that store more unattended information than attended. Only 5.53% of voxels that stored information about both the AMI and the UMI, stored significantly more information about the AMI than the UMI, and 0.63% stored significantly more information about the UMI than the AMI. Consequently, brain regions that stored information about both attended and unattended information, rarely stored more information about one than the other. Finally we found no significant cross-decoding between the AMI and the UMI, in either direction. Hence, when training on one condition and testing on the other, we found no regions that stored information about one condition in the same patterns as the other.

In the second delay we found significant decoding for the PAMI only in the occipital lobe ($z > 3.3$, $p < .001$ (two-tailed)). We found no significant evidence for the PUMI in any brain region. Thus, in the occipital lobe, we found information about attended information in the second delay if the item was previously attended, but not when it was previously unattended.

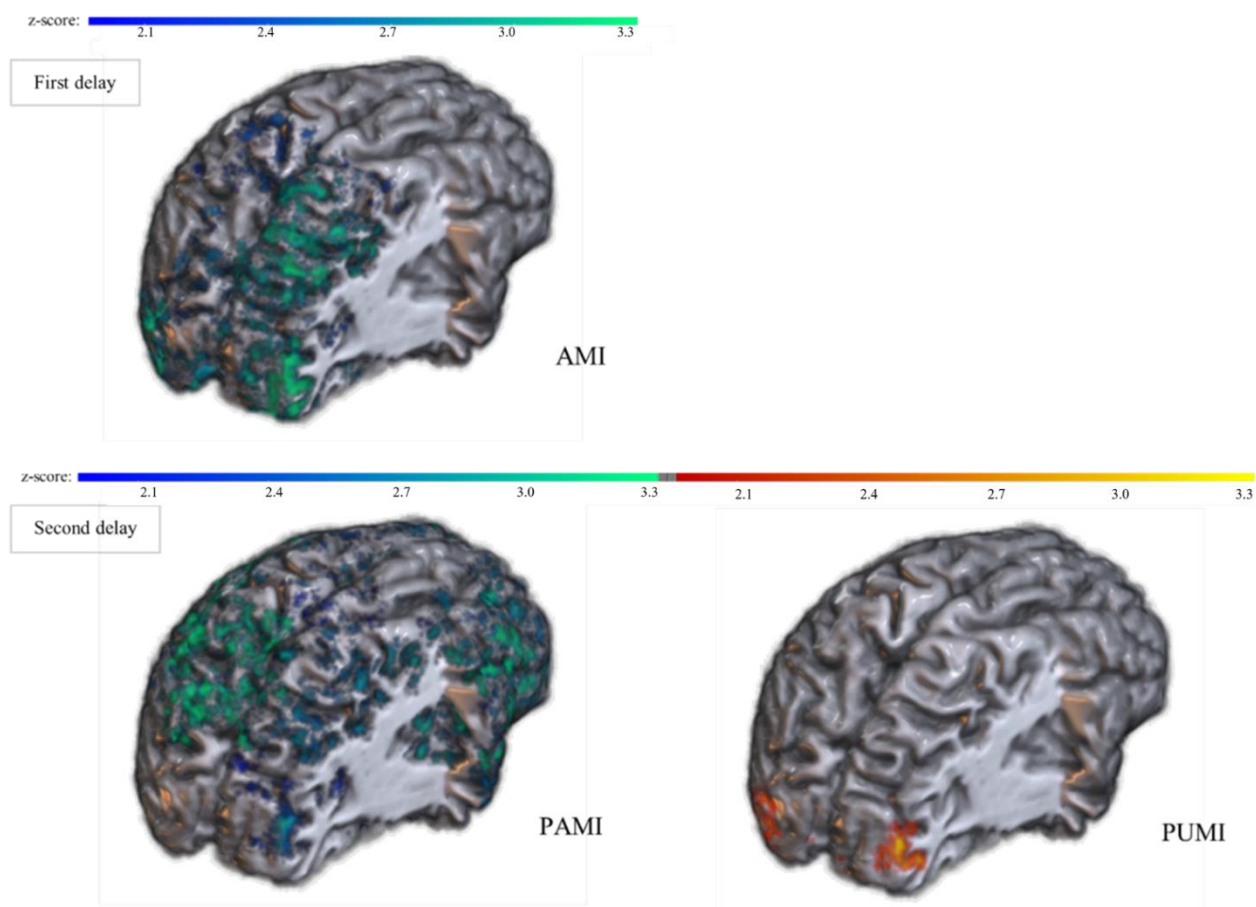
Experiment 2

To understand how memory items being stored in an unattended state affected task performance in the second delay, we conducted a group-level behavioural analysis of our behavioural experiment. A one-way ANOVA was conducted to compare mean absolute response error between the first task, second task (repeat), and second task (switch). There was a significant difference between groups ($F(2,60) = 5.03$, $p = .010$) (Figure 8). A Tukey-Kramer post-hoc analysis revealed mean absolute error was significantly higher in switch delays ($M = 11.34$, $SD = 2.88$ degrees) than in the first delay ($M = 8.94$, $SD = 2.19$ degrees) ($p = .007$, 95% $CI = [-4.25, -0.56]$). There was no significant difference between repeat delays ($M = 9.82$, $SD = 2.34$) and the first delay ($p = .486$) or switch delay ($p = .125$). This means that task performance in the second delay in comparison to the first delay decreased if an item was previously unattended (switch trial), as opposed to previously attended (repeat trial).

However, there was no significant difference in task performance between tasks concerning previously attended and unattended memory items in the second delay.

Figure 2.

Significant decoding of orientations held in visual working memory for Participant 1, seperated by delay and memory condition.




Note. Z-scores indicating significant SVM decoding of orientations calculated using TFCE & using a whole-brain correction for multiple comparison. Z-scores convert to range [$p=.05 - p<.001$].

Figure 3.

Above-chance negative cross-decoding between the AMI and the UMI in the first delay for Participant 1.

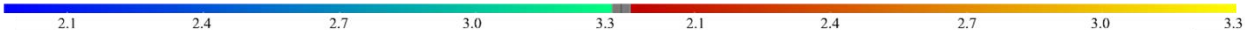
z-score:  -2.1 -2.4 -2.7 -3.0 -3.3




Note. Z-scores indicating significant SVM decoding of orientations calculated by training on the UMI and testing on the AMI. Significance testing done using TFCE & utilising a whole-brain correction for multiple comparison. Z-scores convert to range $[p=.05 - p<.001]$.

Figure 4.


Significant decoding of orientations held in visual working memory in the second delay for Participant 2, separated by condition.

z-score:  2.1 2.4 2.7 3.0 3.3 2.1 2.4 2.7 3.0 3.3

Second delay



PAMI

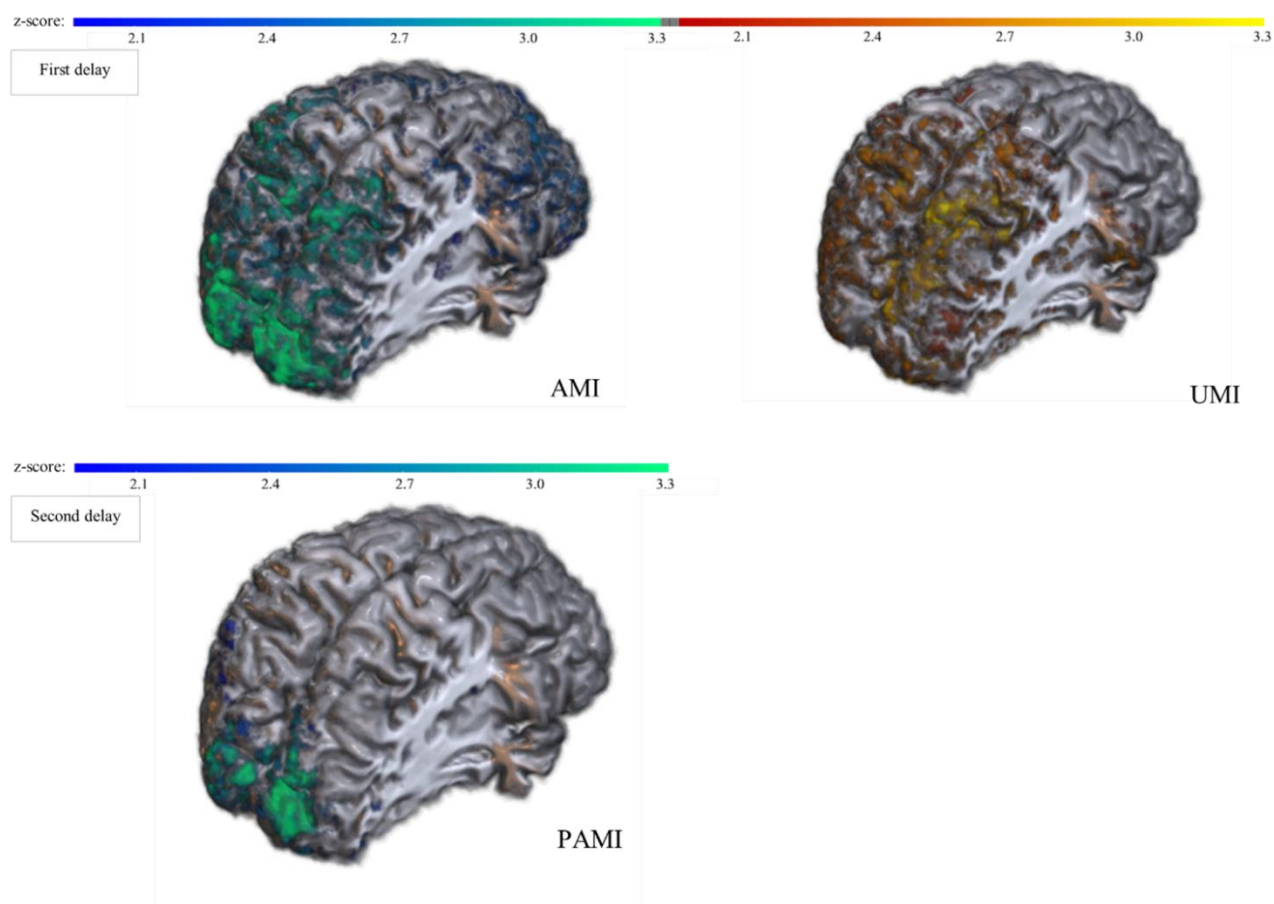


PUMI

Note. Z-scores indicating significant SVM decoding of orientations calculated using TFCE & using a whole-brain correction for multiple comparison. Z-scores convert to range $[p=.05 - p<.001]$.

Figure 5.


Significant decoding of orientations held in visual working memory for Participant 1, separated by delay and memory condition.



Note. Z-scores indicating significant SVM decoding of orientations calculated using TFCE & using a whole-brain correction for multiple comparison. Z-scores convert to range [$p=.05 - p<.001$].

Figure 6.

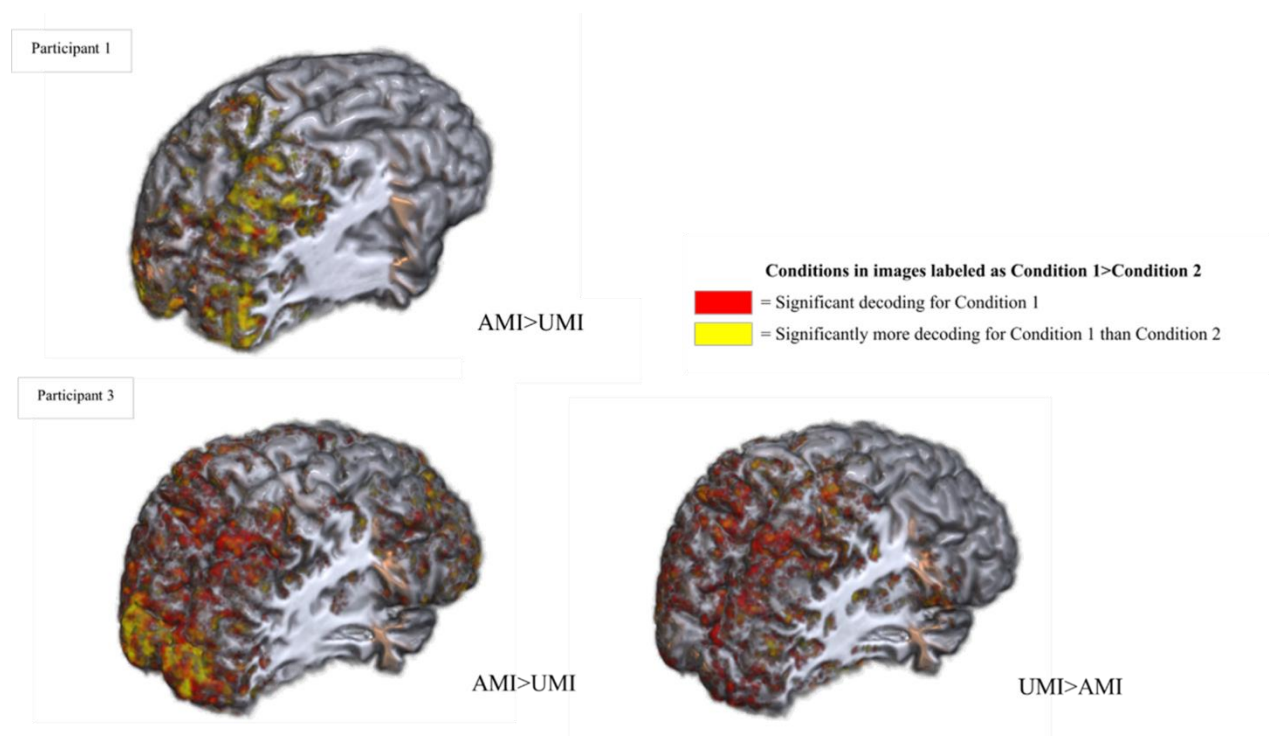
Overlap between decoding for attended and unattended memory items in the first delay for participant 3.



Note. Z-scores for SVM decoding for either attentional condition calculated using TFCE and whole-brain correction for multiple comparison ($z>1.96, p<.05$ (two-tailed)).

Figure 7.

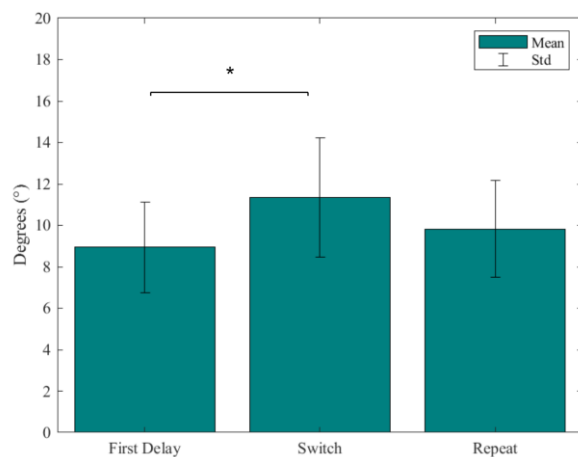
Brain regions containing significantly more information about AMI>UMI and UMI>AMI, separated by participant.



Note. Significant SVM decoding calculated using TFCE and whole-brain corrections for multiple comparisons ($z > 1.96$, $p < .05$ (two-tailed)). Participant 1: Regions containing significantly more information about the AMI than the UMI highlighted ($t(315) > 1.97$, $p < .05$ (two-tailed)). Participant 3: Regions containing significantly more information about AMI>UMI and UMI>AMI highlighted ($t(319) > 1.97$, $p < .05$ (two-tailed)).

Figure 8.

Mean response error and standard deviation for first and second delay (repeat & switch) from Experiment 2.



Note. * = $p < .05$.

Exploratory analysis 1

One possible caveat in these findings is that we used a binary classification model to decode visual information. This meant we had to bin different orientations into horizontal/vertical/diagonal1/diagonal2 groups, thereby discarding potentially informative orientation information. It is possible that binning different orientations into the same relatively large group (45 degrees per bin), subtle signal differences between the different orientations within each bin increased noise, and decreased decoding accuracy. Furthermore, we could only train on half the data in each SVM (horizontal/vertical or diagonal1/diagonal2 trials).

Some studies make use of an inverted encoding model (IEM) which bypasses these issues (Rademaker et al., 2019; Yu et al, 2020). This model uses linear regression to train a model to predict voxel activity by assigning weights to different orientation channels. This model can then be inverted to decode novel stimuli using the channel weights. Since this model allows for continuous input, the different orientations do not need to be binned, and can be used as-is. Consequently, the output value for the test item can be quantified on a continuous (0-180 degree) scale, as opposed to a classification between two bins. Additionally, all data can be used to train and test on each trial, increasing model strength. Thus, using an IEM model could have increased decoding strength.

To explore how using an IEM model instead of an SVM model could have impacted our results, we added an exploratory analysis. Here, we used IEM to decode visual information in one participant (participant 3) so we could compare it to our SVM analysis. We used 9 IEM basis functions to calculate weights for 180 channels, one for each degree in the 180 degree orientation space. We rounded each experimental orientation to the nearest integer and trained and tested the model on each trial using a leave-one-out method. We did this using the same searchlight method as our SVM analysis, training and testing on data from a voxel sphere with a radius of 3, and then projecting the channel responses on the centre voxel. We repeated this process for each voxel in the grey matter mask we also used for our SVM analysis.

We then used a fidelity method as described in Rademaker et al. (2019), to calculate the amount of information found for the memorised orientation. This was done by shifting all orientation channel responses to the centre of 0 degrees, and multiplying them with the function $\cosd(\text{abs}(1:180)*2-180)$. Multiplying the channel responses with this function incrementally flips the sign of channel responses at orientations further away from the viewed orientation. This means that high channel responses at the viewed orientation and low channel responses far from the viewed orientations produce a positive value, while low channel responses at the viewed orientation and high channel responses at orientations far from the viewed orientations produce a negative value. Positive values thus indicate neural support for the

correct orientation or lack of evidence for the incorrect orientation, both of which support orientation decoding. Finally, we took the mean of all these new fidelity values for each voxel over all trials. This resulted in a map of fidelity values for each voxel over all trials. Here a positive value indicates that a voxel sphere consistently showed evidence for the viewed orientation, and not for other orientations. We did not perform additional statistical testing, as this analysis was done mainly to get an idea of analysis differences, and not to infer actual reconstruction effects.

Since we only calculated the mean fidelity scores over all trials for each condition for the IEM, we opted to compare it to the mean scores of the SVM's signed distance-to-bound scores, instead of the z-scored results from the TFCE statistical analysis. This is because the mean SVM distance-to-bound scores are a measure of the mean decoding strength of our SVM model, and the mean fidelity scores are a measure of the mean decoding strength of our IEM model. Therefore, both values reflect their respective model's estimate of decoding strength without being tested for significance. This allows for the best comparison of results.

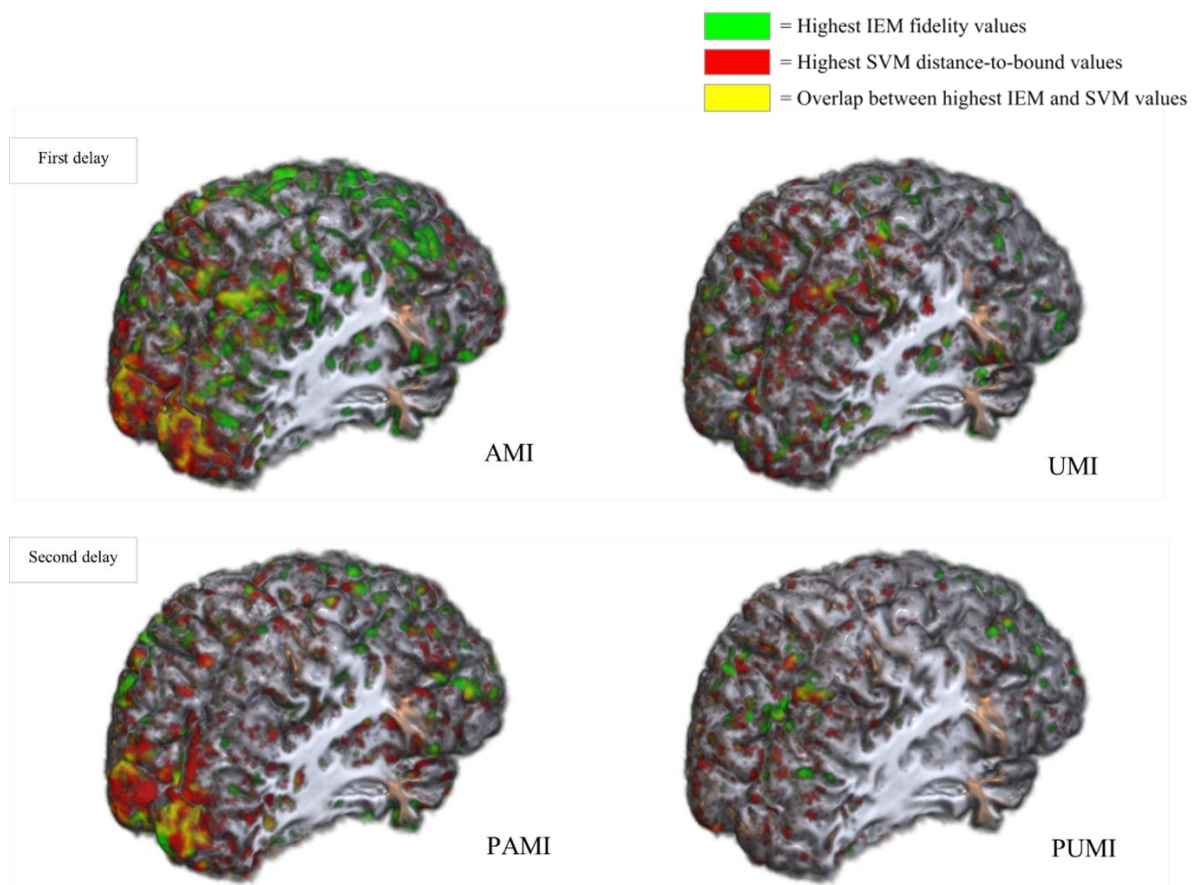
Since there is no official cutoff for these non-standardised measures, we looked at the highest means of both analysis methods (fidelity score/distance-to-bound score).

Results exploratory analysis 1: For each attentional condition (AMI-UMI-PAMI-PUMI) we observed overlap between high mean IEM fidelity and SVM distance-to-bound values across occipital and parietal regions (Figure 9). Here, the highest mean-value regions either overlap or are spatially close to each other. Regarding the frontal lobe, there seems to be much more variability, with little to no overlap or spatially close regions.

So, at least for the occipital and parietal lobe, both IEM and SVM analysis results in high value scores in similar regions. In the frontal lobe, however, IEM and SVM values point to different sub-regions. Since these IEM and SVM scores are not-significance tested it is unclear whether these differences are statistically significant, so one should be careful in making inferences based on this data. The takeaway from this exploratory analysis is that both SVM and IEM seem to point to similar brain regions in the occipital and parietal lobe, and as such could likely both be used to decode visual information here without losing much decoding strength.

Figure 9. (Exploratory analysis)

Highest IEM fidelity and SVM distance per attentional condition for participant 3.



Note. Overlap between highest means for the AMI (IEM>.09, SVM>.25). Overlap between highest means for the UMI (IEM>.09, SVM>.25). Overlap between highest means for the PAMI (IEM>.10, SVM>.25). Overlap between highest means for the PUMI (IEM>.07, SVM>.27).

Exploratory analysis 2

In this study, we use signed distance-to-bound values to quantify the decoding strength of our SVM model across the cortex, and then used this as a measure for visual information storage. We assume that regions with high distance-to-bound values indicate visual information storage, but we do not have a direct way of examining whether this is true. One way to verify this is by correlating the distance-to-bound values in the regions where visual information is found with task performance. Since, if the distance-to-bound values in these regions correlate with visual information storage, you would expect them to also correlate with task performance. This is because you need to correctly evaluate and remember the original stimulus in order to correctly match it in our match-to-sample task.

In our second exploratory analysis, we examine in which brain regions distance-to-bound values actually correlate with task performance. This is done mainly as a sanity check, since it is unlikely that a small group of <100 voxels will significantly impact task performance. Especially since we learned from our main results that visual information is stored in many areas across the cortex. Therefore, task performance overall is more likely to correlate with the weighed product of many different areas containing visual information, as opposed to individual isolated areas. Still, you would expect to observe some correlation, however small, in areas where we decoded visual information.

To explore the correlation between distance to bound values and task performance, we analysed two sets of data. We only included participants and conditions where we were able to decode visual information, since we were interested in seeing whether the decoding in regions where we found visual information actually correlated with behaviour.

Attended memory items: Here, we took all distance to bound values from our AMI condition, and correlated them with the responses in the first trial task. For this, we could use all AMI distance-to-bound values and all first delay responses. This is because the task in this first delay pertains the attended memory item in the first delay.

Unattended memory items: Here, we took the distance to bound values from our UMI condition that were later used in the switch trial second task. Then, we also took all responses from the switch trial second task. This is because in a switch trial, the previously unattended memory item (UMI) is used in the task. Thus, if there are brain regions where the distance-to-bound values for the UMI indicate visual information storage, you would expect the values of these regions to correlate with task performance in the second half of these switch trials.

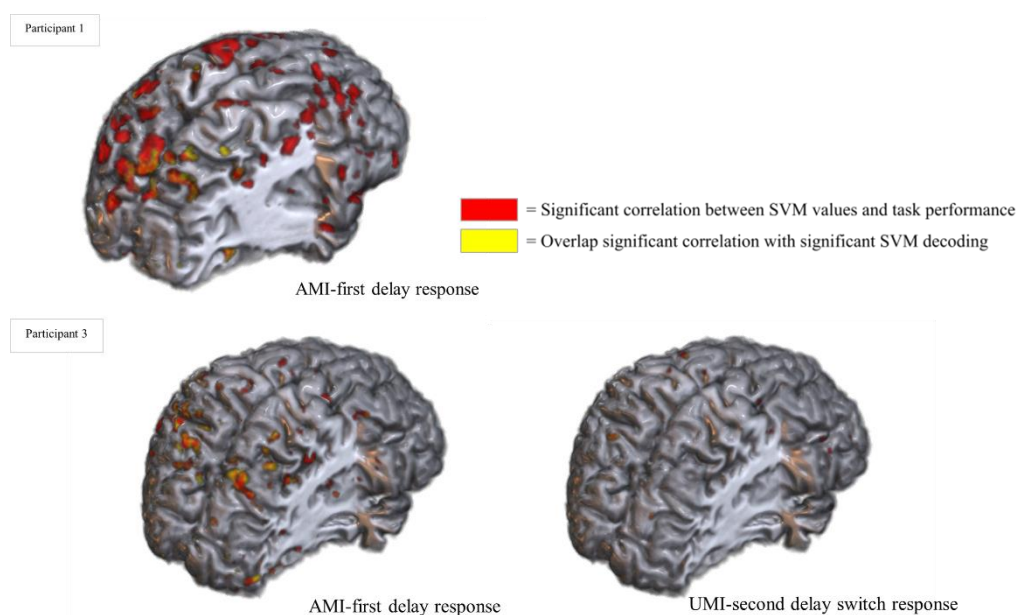
For this analysis, we used a searchlight with a radius of 3 to correlate the distance-to-bound values in each searchlight sphere with the absolute errors from the corresponding response trials.

Again, we projected the result of this analysis to the centre voxel of each sphere. For both conditions (attended and unattended), we used Pearson's pairwise linear correlation to quantify the relationship between the signed distance-to-bound values and the absolute errors from the same trials. Here, a negative correlation coefficient indicates a positive effect on task performance, as errors decrease as performance increases. We used an alpha of 0.05 and no correction for multiple comparisons. This gave us a map with correlation coefficients, and one containing p-values for the correlation analysis for each voxel in each attentional condition (attended/unattended).

Results exploratory analysis 2: We calculated the correlation between distance-to-bound values and task performance for the participants where we were able to decode the AMI (and UMI) (ptc 1&3). For the attended memory items we see regions where distance-to-bound values significantly negatively correlate with response errors in occipital, parietal, and frontal regions (Figure 10). Interestingly, if we single out the areas that actually contain visual information according to our main analysis, mainly regions in the parietal lobe remain. Thus, according to this metric, mainly parietal, and some occipital regions containing visual information about attended visual information also exhibit activity correlating with task performance. For unattended memory items, we see small regions in the parietal lobe where distance-to-bound values negatively correlate with response errors. Most of these regions also contain visual information about unattended memory items according to our main analysis. These results provide some indication that the regions we describe in our main analysis as containing visual information indeed do, as they negatively correlate with response errors (and thus positively with task performance).

Figure 10. (Exploratory analysis)

Above-chance correlations between SVM distance-to-bound values and absolute response errors.



Note.

Significant negative Pearson's pairwise linear correlation coefficients ($p < .05$) (not corrected for multiple comparisons). Significant SVM decoding calculated using TFCE and whole-brain corrections for multiple comparisons ($z > 1.96$, $p < .05$ (two-tailed)).

Discussion

In this study we aimed to understand where visual information is stored during visual working memory retention, depending on the attentional status of the working memory content. We examined where visual information during VWM storage could be decoded for attended, unattended, previously attended and previously unattended memory items. Then, we examined where significant differences between attended and unattended memory storage could be found. Finally we wanted to see how this relates to behaviour. We examined 7T fMRI data from three participants and found differing evidence for each attention condition.

Before we discuss our fMRI results, a few consequences of our approach should be mentioned. For experiment 1, we opted for a within-subjects analysis with many trials per participant. This gave us the opportunity to get a detailed assessment of neural patterns within each participant with high statistical accuracy. Consequently, it allowed us to uncover and decode weaker patterns, as we have large amounts of data on those patterns for each participant. However, using few participants doesn't give us the statistical strength for population level inference. Thus, while our results can give a unique insight into possible neural mechanisms, they cannot be used to make conclusions about VWM visual information storage on a population level.

Where in the cortex is visual orientation information stored during visual working memory retention?

To answer this question, we used support vector machine (SVM) classification to decode visual information across the entire cortex. We did this for attended and unattended visual information in the first experimental delay, and attended information that was either previously attended or previously unattended in the second experimental delay. This allowed us to decode with high accuracy where visual information is stored during VWM retention.

In the first delay, we were able to decode attended visual information in occipital and parietal lobes in two participants (ptc 1&3). For one of those participants we also found information about attended orientations in the frontal lobe (ptc 1). We were not able to find any evidence for attended visual information in the remaining participant (ptc 2). The finding that visual information of an attended memory item can be found in occipital and parietal regions is in line with previous studies and our expectations (Christophel et al., 2018; Christophel et al., 2012; Ester et al., 2015; Yu et al., 2020). Though we did not find this same

result for each participant, it confirms that attended visual information can be stored in more than one region in visual working memory.

We were also able to decode unattended visual information in the first delay in one participant (ptc 3). For this participant, we found evidence for unattended visual information in both occipital and parietal regions. Depending on the analysis method, information about unattended memory items have previously been decoded from occipital and parietal regions (Christophel et al., 2018; Iamshchinina et al., 2021; Yu et al., 2020). However, this has only been possible when training the decoding model on attended memory items. This is because attended memory items are represented stronger in neural patterns, so training on them and then testing on unattended representations improves decoding sensitivity for unattended mnemonic representations (van Loon, 2018). The fact that we were able to decode unattended memory items in the occipital lobe while also training on unattended items, is a novel finding. This finding adds two things. Firstly, it confirms the benefit of using a method utilising 7T fMRI with many trials per participant, as it allowed us to detect and analyse weaker mnemonic patterns. Consequently, it highlights why it is important not to confuse null results with the absence of information within VWM research, as it is possible those results are a consequence of insensitive research methods. Secondly, it adds support that unattended information can be represented in the EVC, something that was previously thought to be reserved for attended visual information. It is possible that unattended visual information is represented in similar regions as attended information, though encoded in weaker neural patterns.

In the second delay, we found information about attended memory items that were previously attended in the occipital lobe for each participant. We were also able to decode visual information in the parietal and frontal lobe in two participants (ptc 1&2). We were only able to decode visual information in a small region of the parietal lobe in one participant (ptc 3). Since this memory item has been attended throughout the trial, it is unsurprising to find evidence for it in occipital and parietal regions, as attended memory items in the first delay could also be decoded here.

In the second delay, we either found information about attended memory items that were previously unattended in small regions in the frontal lobe (ptc 1), occipital lobe (ptc 2), or not at all (ptc 3). Finding little to no evidence for attended memory items that were previously unattended is interesting. This is because the item is currently attended, and therefore should be in the same attentional state as attended memory items in the first delay. Since we found evidence for these first delay attended items in the occipital and parietal lobe,

one would expect to find information for previously unattended (but currently attended) memory items in these regions as well. While some loss of precision should be expected, finding this little evidence for this previously unattended items could indicate that there is a non-reversible loss of voxel representational strength in the unattended state.

For some participants, we also found decoding in frontal regions for attended visual information in the first and/or second delay. While this could mean visual information is stored there, it is also possible this reflects possible preparatory motor or other higher-order processes (Haxby et al. 2000; Lee et al., 2013). Since this paradigm relies on items to become task-relevant, removing motor or decision processes usually found in the frontal lobe is problematic. Since, without a task, it is difficult to make something task-relevant. Additionally, we found that the locations in which we were able to decode visual information in the frontal lobe depended on the analysis method. When comparing our SVM analysis with an IEM in our exploratory analysis, we found considerably less overlap in regions containing visual information in the frontal lobe, compared the occipital and parietal lobe. The lack of generalisation between analysis methods for the frontal lobe could mean different things. Firstly, it could mean that the frontal lobe equips different strategies for storing visual information. For example, more categorical memory storage would be better detected by the SVM model, while more continuous mnemonic representations would be better detected by the IEM model. Secondly, it is possible that it is a result of spatial distortions in the frontal cortex caused by the air-filled cavities and bone structures located near this lobe (Juchem et al., 2010). Since there are several factors that could have influenced information decoding in the frontal lobe, it is unclear whether these positive results actually reflect visual information storage.

Are there differences in where attended and unattended memory items are stored in visual working memory?

To answer this question we compared the decoding model output from attended and unattended memory items collected in our fMRI experiment. We only found evidence for attended and unattended visual information in one participant (ptc 3). Therefore, we will focus on this participant's results for this question.

We found 21-26% overlap in regions where attended and unattended visual information was stored in VWM in the occipital and parietal lobe in the first delay. Only 1-6% of these regions contained more information about one memory item over the other. In

other words, there is some overlap in where attended and unattended information is stored in VWM, and these regions generally do not appear to store more information about one item over the other.

Within each lobe that contained information about attended memory items (occipital, parietal, & frontal), we found areas that contained more information about attended than unattended memory items. This was mostly visible in the occipital lobe, where a large region of the medial occipital lobe (where the EVC is located) contained more information about attended than unattended memory items. This is in line with previous research which suggests that the brain allocates more neural resources from the EVC to attended visual information (Christophel et al., 2018). The EVC is able to store and retain visual information in a detailed, perception like nature (Albers et al., 2013; Serences et al., 2009). Storing attended visual items in this detailed state could then aid in tasks like search. This is because a detailed mnemonic representation of the searched object better matches the simultaneous visual input from the environment. Our study thus provides additional support for the hypothesis that the brain distributes more occipital resources to attended memory items as opposed to unattended memory items.

Additionally, within each lobe that contained information about unattended memory items (occipital & parietal) we found areas that contained more information about unattended memory items than attended memory items. While these regions could be found in both the occipital and parietal lobe, they were more sparse in the occipital than the parietal lobe. This is also consistent with the aforementioned theory that the brain allocates more resources from occipital regions to attended than unattended memory items, while the parietal cortex is more insensitive to attentional effects (Christophel et al., 2018). The parietal cortex might thus be tasked with storing both attended and unattended visual information relatively equally, while attended visual information is prioritised in the occipital lobe. We found some regions within the occipital lobe where unattended memory items are represented more strongly than attended items. However, these regions were smaller and located more laterally than the occipital regions that store attended information over unattended. Thus, while the brain stores unattended memory items in occipital regions, these regions are smaller and more laterally located than the regions that store attended memory items.

Previous studies have already succeeded in decoding attended and unattended visual information from occipital and parietal regions (Christophel et al., 2018; van Loon et al., 2018; Yu et al., 2020). This study adds that within these information storage regions, there are areas that contain more information about the AMI than the UMI, but also areas that contain

more information about the UMI than the AMI. Therefore, it seems the brain is capable of distributing separate neural resources based on whether an item is attended or not.

Does the brain store attended and unattended information differently?

To answer this question, we used cross-decoding on our fMRI data. Here, we trained our classifier on one attentional condition (attended/unattended) and tested it on the other. We found no significant cross-decoding between the two conditions in any participant. In fact, for one participant (ptc 1) we found negative cross-decoding for attended memory items when training on unattended memory items in occipital and frontal regions. We used a binary classification method with orientational bins directly opposite one another in the orientation space. Therefore, negative cross-decoding between attentional conditions could mean that the voxel patterns for the orientation bins are opposite for attended and unattended memory items. Here, the neural representation of unattended memory items would be rotated about 90 degrees relative to attended items. This is in accordance with previous studies (van Loon et al., 2018; Yu et al., 2020). Here, it was that found attended and unattended memorised orientations are represented in opposite patterns in the EVC (Yu et al., 2020) and object selective cortex (van Loon et al., 2018). This priority-based remapping may be utilised to store attended and unattended memory items concurrently, while reducing the interference of unattended memory items on attended ones. This allows the accessibility for this attended item to be optimised without being disrupted by currently irrelevant information.

However, since the attended and unattended memory items were constrained to be at least 6 degrees apart, it is also possible that this caused an imbalance in our classifier model. Since, by constraining the two memory items to be 6 degrees apart, it increased the chance of the unattended memory item to be in a different bin than the attended memory item. This might have affected pattern recognition in the training stage, biasing it towards opposite neural patterns, potentially causing negative cross-decoding in the testing stage. This could be prevented by not constraining the possible orientations of the attended and unattended memory items. Nevertheless, detecting opposite neural patterns for attended and unattended memory items is an interesting finding that warrants further research into how and where these patterns differ.

Does visual information being stored in an unattended state affect task performance?

We found that unattended memory items are less strongly and less widespread represented than attended items. This lower-resolution storage of unattended memory items

also translated to behaviour. In our behavioural experiment (Experiment 2), we found that response errors increased when an item was previously unattended, as opposed to the first attended response. This was not the case for previously attended items. Hence, a memory item having previously been stored in an unattended state negatively impacted response, while having been stored in an attended state did not. We did not find a significant difference in response errors between previously attended and previously unattended items in the second trial response. This deviates from a similar study, where researchers did find a significant difference between previously attended and previously unattended response trials (Christophel et al., 2018). It is possible that this is a result of our relatively small sample size (N=20 as opposed to N=89) and the fact we did not utilise every possible orientation for every participant. We also only had 120 trials per participant, as opposed to 192 in the study by Christophel et al. Using a larger sample group could result in a significant difference between switch and repeat trial performance. Still, the finding that task performance in the second delay decreases for items held in an unattended state, but not an attended state compared to the first response supports the theory that there is a loss of information in the mnemonic representation of unattended memory items.

Limitations & recommendations for future research

Our results give a unique insight into possible neural mechanisms that underlie attended and unattended memory storage. Using 7T fMRI and within-participant statistics with many trials per participant, we were able to decode visual information in many brain areas across the cortex. In one participant, this was even possible for low-strength unattended memory items. This means that in future studies, the use of 7T fMRI could prove useful in finding group-level effects, as this method seems to lend itself to detecting weaker activation patterns.

However, we did not find consistent results across all participant for any of our attentional conditions. Within each condition, there was high variability in where visual information could be decoded between participants. This could be caused by (a combination of) three things.

- 1: It is possible that our analysis method was still not sensitive enough. For example, we were not only able to find evidence for unattended memory items stored in visual working memory for two participants. It is possible that the neural patterns in which unattended memory items are encoded were too weak to be fully detected in these participants. Using more trials could improve this, by giving us more data to train the classifier models on.

2: it is possible that our statistical significance testing method was too strict. We analysed our data using an SVM classifier, and then tested for significance while correcting for multiple comparisons over the entire cortex. This means that while our results are highly accurate, smaller effects could have been deemed insignificant, since we applied a whole brain correction. This could be improved by limiting the searchlight method to parietal and occipital regions. This preserves some of the benefit of non-ROI based research, while increasing statistical power.

3: Finally, it is also possible that some null-results are the result of different neural strategies for retaining visual information in VWM. For example, in the first delay, participants had to memorise two different orientations. We trained and tested our model on the assumption that these orientations would be memorised separately. Hence, we assumed participants would memorise item 1, and item 2 separately, and on a continuous scale. However, it is also possible that some participants stored the unattended item relative to the attended item. Here, instead of storing the orientation of the unattended item itself, they might have stored how far away (in degrees) it was from the attended item. This would lead to the unattended item not being able to be decoded using our analysis method. Therefore, individual mnemonic strategy differences could have affected our results.

In short, there are several aspects that could have influenced our results. For further research it might be beneficial to restrict the analysis to the occipital and parietal lobes. Furthermore, it might be useful to repeat this experiment on more participants. This would allow us to gain more insight into whether our results generalise across a larger participant group.

In conclusion, using 7T fMRI we were able to decode visual information in different attentional formats. We decoded attended information in occipital, parietal, and frontal regions. In addition, we were able to decode unattended visual information in occipital and parietal regions. Since we applied a strict whole brain correction for our results, this study provides further evidence that unattended memory items can be represented in similar regions as attended memory items. Furthermore, we found that there are brain regions that are specifically used to store either attended or unattended visual information. Within these results, we found that attended memory items were more widely represented in the occipital lobe than unattended memory items. This could be done to facilitate higher-resolution representations for attended information, as this information is likely more behaviourally relevant. Finally, there is some evidence that attended and unattended visual information

might be represented differently from one another. This could be an efficient way of separating concurrently stored memory items in VWM, as each has its own neural code. Different mnemonic neural codes and memory storage locations could be mechanisms used by the brain to efficiently distribute neural resources between attended and unattended visual memory items.

Contribution statement:

Thesis written by F. Ruijs (author).

Author was supervised by Dr. S. Gayet Phd., who conceived the idea for this thesis and supervised author throughout the project as part of author's main research project for the Neuroscience and Cognition master programme at Utrecht University. All experiment setup, data collection, and analysis was supervised.

fMRI scanning conducted by van Ackooij, Msc., Dr. B. Harvey, and Dr. S. Gayet, M. at Spinoza Centre for Neuroimaging. fMRI pre-processing conducted by M. van Ackooij, Msc. and Dr. B. Harvey.

IEM analysis script for exploratory analysis 1 taken from Rademaker, Chunharas, and Serences (2020) (open access) and appended for searchlight use by author.

All other analysis scripts written and executed by author.

References

- Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & De Lange, F. P. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, 23(15), 1427-1431.
- Baddeley, A. (1998). Working memory. *Comptes Rendus de l'Académie des Sciences-Series III-Sciences de la Vie*, 321(2-3), 167-173.
- Baluch, F., & Itti, L. (2011). Mechanisms of top-down attention. *Trends in neurosciences*, 34(4), 210-224.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision research*, 51(13), 1484-1525.
- Christophel, T. B., Hebart, M. N., & Haynes, J. D. (2012). Decoding the contents of visual short-term memory from human visual and parietal cortex. *Journal of Neuroscience*, 32(38), 12983-12989.
- Christophel, T. B., Iamshchinina, P., Yan, C., Allefeld, C., & Haynes, J. D. (2018). Cortical specialization for attended versus unattended working memory. *Nature neuroscience*, 21(4), 494-496.
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, 29(3), 162-173.
- Emrich, S. M., Riggall, A. C., LaRocque, J. J., & Postle, B. R. (2013). Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *Journal of Neuroscience*, 33(15), 6516-6523.
- Eriksson, J., Vogel, E. K., Lansner, A., Bergström, F., & Nyberg, L. (2015). Neurocognitive architecture of working memory. *Neuron*, 88(1), 33-46.
- Ester, E. F., Anderson, D. E., Serences, J. T., & Awh, E. (2013). A neural measure of precision in visual working memory. *Journal of cognitive neuroscience*, 25(5), 754-761.
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*, 87(4), 893-905.
- Etzel, J. A., Zacks, J. M., & Braver, T. S. (2013). Searchlight analysis: promise, pitfalls, and potential. *Neuroimage*, 78, 261-269.
- Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current opinion in neurobiology*, 20(2), 177-182.

- Frătescu, M., Van Moorselaar, D., & Mathôt, S. (2019). Can you have multiple attentional templates? Large-scale replications of Van Moorselaar, Theeuwes, and Olivers (2014) and Hollingworth and Beck (2016). *Attention, Perception, & Psychophysics*, *81*, 2700-2709.
- Funahashi, S., & Andreau, J. M. (2013). Prefrontal cortex and neural mechanisms of executive function. *Journal of Physiology-Paris*, *107*(6), 471-482
- Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., ... & Dosenbach, N. U. (2017). Precision functional mapping of individual human brains. *Neuron*, *95*(4), 791-807.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*, 632-635.
- Hollingworth, A., & Beck, V. M. (2016). Memory-based attention capture when multiple items are maintained in visual working memory. *Journal of experimental psychology: human perception and performance*, *42*(7), 911.
- Haxby, J. V., Petit, L., Ungerleider, L. G., & Courtney, S. M. (2000). Distinguishing the functional roles of multiple regions in distributed neural systems for visual working memory. *Neuroimage*, *11*(2), 145-156.
- Hendriks, E., Paul, J. M., van Ackooij, M., van der Stoep, N., & Harvey, B. M. (2022). Visual timing-tuned responses in human association cortices and response dynamics in early visual cortex. *Nature Communications*, *13*(1), 3952.
- Iamshchinina, P., Christophel, T. B., Gayet, S., & Rademaker, R. L. (2021). Essential considerations for exploring visual working memory storage in the human brain. *Visual Cognition*, *29*(7), 425-436.
- Jerde, T. A., Merriam, E. P., Riggall, A. C., Hedges, J. H., & Curtis, C. E. (2012). Prioritized maps of space in human frontoparietal cortex. *Journal of Neuroscience*, *32*(48), 17382-17390.
- Juchem, C., Nixon, T. W., McIntyre, S., Rothman, D. L., & de Graaf, R. A. (2010). Magnetic field homogenization of the human prefrontal cortex with a set of localized electrical coils. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *63*(1), 171-180.
- Kerzel, D., & Witzel, C. (2019). The allocation of resources in visual working memory and multiple attentional templates. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(5), 645.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain

- mapping. *Proceedings of the National Academy of Sciences*, 103(10), 3863-3868.
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: the dangers of double dipping. *Nature neuroscience*, 12(5), 535-540.
- Lee, S. H., & Baker, C. I. (2016). Multi-voxel decoding and the topography of maintained information during visual working memory. *Frontiers in systems neuroscience*, 10, 2.
- Lee, S. H., Kravitz, D. J., & Baker, C. I. (2013). Goal-dependent dissociation of visual and prefrontal cortices during working memory. *Nature neuroscience*, 16(8), 997-999.
- Lepsien, J., & Nobre, A. C. (2007). Attentional modulation of object representations in working memory. *Cerebral cortex*, 17(9), 2072-2083.
- Mahmoudi, A., Takerkart, S., Regragui, F., Boussaoud, D., & Brovelli, A. (2012). Multivoxel pattern analysis for fMRI data: a review. *Computational and mathematical methods in medicine*, 2012.
- Noble, W. S. (2006). What is a support vector machine?. *Nature biotechnology*, 24(12), 1565-1567.
- Olivers, C. N., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in cognitive sciences*, 15(7), 327-334.
- Oosterhof, N. N., Connolly, A. C., & Haxby, J. V. (2016). CoSMoMVPA: multi-modal multivariate pattern analysis of neuroimaging data in Matlab/GNU Octave. *Frontiers in neuroinformatics*, 10, 27.
- Pratte, M. S., & Tong, F. (2014). Spatial specificity of working memory representations in the early visual cortex. *Journal of vision*, 14(3), 22-22.
- Pernet, C. R., Latinus, M., Nichols, T. E., & Rousselet, G. A. (2015). Cluster-based computational methods for mass univariate analyses of event-related brain potentials/fields: A simulation study. *Journal of neuroscience methods*, 250, 85-93.
- Rademaker, R. L., Chunharas, C., & Serences, J. T. (2019). Coexisting representations of sensory and mnemonic information in human visual cortex. *Nature neuroscience*, 22(8), 1336-1344.
- Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *Journal of Neuroscience*, 32(38), 12990-12998.
- Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition*, 80(1-2), 1-46.

- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological science*, *20*(2), 207-214.
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, *44*(1), 83-98.
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, *24*(18), 2174-2180.
- Torrìsi, S., Chen, G., Glen, D., Bandettini, P. A., Baker, C. I., Reynolds, R., ... & Ernst, M. (2018). Statistical power comparisons at 3T and 7T with a GO/NOGO task. *NeuroImage*, *175*, 100-110.
- Triantafyllou, C., Hoge, R. D., Krueger, G., Wiggins, C. J., Potthast, A., Wiggins, G. C., & Wald, L. L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage*, *26*(1), 243-250.
- Ungerleider, L. G., Courtney, S. M., & Haxby, J. V. (1998). A neural system for human visual working memory. *Proceedings of the National Academy of Sciences*, *95*(3), 883-890.
- van der Zwaag, W., Francis, S., Head, K., Peters, A., Gowland, P., Morris, P., & Bowtell, R. (2009). fMRI at 1.5, 3 and 7 T: characterising BOLD signal changes. *Neuroimage*, *47*(4), 1425-1434.
- van Loon, A. M., Olmos-Solis, K., Fahrenfort, J. J., & Olivers, C. N. (2018). Current and future goals are represented in opposite patterns in object-selective cortex. *ELife*, *7*.
- van Moorselaar, D., Theeuwes, J., & Olivers, C. N. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention?. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450.
- Wager, T. D., & Smith, E. E. (2003). Neuroimaging studies of working memory. *Cognitive, Affective, & Behavioral Neuroscience*, *3*(4), 255-274.
- Yu, Q., Teng, C., & Postle, B. R. (2020). Different states of priority recruit different neural representations in visual working memory. *PLoS biology*, *18*(6), e3000769.