Featural and spatial interference with functionally active and passive items in working memory

Sophia A. Wilhelm [12]*, Elkan G. Akyürek [1]*, Yuanyuan Weng [1], Jelmer P. Borst [3], & Robbert Havekes [2]

[1] Experimental Psychology, Department of Psychology, University of Groningen
[2] Neurobiology Expertise Group, Groningen Institute of Evolutionary Life Sciences, University of Groningen
[3] Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen

* Shared first authors

Running head: Featural and spatial interference in working memory

05/02/2024

Word count (excluding abstract, methods, references): 3054

Address correspondence to:
Sophia Wilhelm, Department of Experimental Psychology, University of Groningen
Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands, s.a.sophia.wilhelm@rug.nl, +31 (0)50 3633635

Elkan Akyürek, Department of Experimental Psychology, University of Groningen
Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands, e.g.akyurek@rug.nl, +31 (0)50 3636406

Abstract

We hypothesised that functionally active and passive states in working memory might be related to different neural mechanisms. Memoranda in active states might be maintained by persistent neural firing, whereas memoranda in passive states might be maintained through short-term synaptic plasticity. We reasoned that this might make these items differentially susceptible to interference during maintenance, and that passive items in particular might be more robust. To test this hypothesis, we gave our participants a working memory task in which one item was prioritised (active) by always probing it first, while the other item was deprioritised (passive) by always probing it second. In three experiments, we presented an interfering task during memory maintenance, in which the stimuli matched either the feature dimension of the memory items (colour or orientation), their spatial location, or both. We found that both active and passive memory items were affected by interference, to a similar extent, with overall performance being closely matched. However, we observed a qualitative difference in the nature of their representations. Specifically, featural interference seemed to produce reduced precision for active items, whereas it resulted in higher guess rates for passive items. The existence of this qualitative difference suggests that functionally active and passive working memory states may indeed map onto different neural mechanisms, even though on average they seem to be equally robust to interference.

Keywords: working memory, activity-silent working memory, interference

Introduction

Working memory (WM) is the cognitive system that allows us to store and manipulate information that is not currently available to the senses (Baddeley & Hitch, 1974). Initial evidence of neurophysiological studies on working memory have proposed ongoing neuronal activity as a neuronal mechanism of storing information in WM (Goldman-Rakic, 1995). However, in recent years this view that the underlying mechanism of WM is exclusively dependent on ongoing neuronal firing has been questioned (Stokes, 2015). Indeed, evidence has accumulated which shows that intact WM behaviour can be observed even in the absence of persistent neuronal firing (LaRocque et al., 2012; Muhle-Karbe et al., 2021; Rose et al., 2016; Wolff et al., 2017), and even in situations where an interference task was presented to disrupt ongoing (alpha) activity (Oberauer & Awh, 2022). Thus, evidence is accumulating making a strong case for maintenance mechanisms of working memory that do not exclusively rely on persistent firing.

It has been shown that memory items that are stored in this activity-silent or -quiescent state of WM can change back into an active state according to task-based cues. For example, in a classic two-item retro-cueing paradigm, items that are known to be relevant only later in the trial first enter an activity-quiescent state but are transformed back into an active state once a response to the prioritised (or cued) item has been made (LaRocque et al., 2012; Wolff et al., 2017). Furthermore, it has also been shown that two items can adaptively be moved back between active and silent states on a single-trial basis, based on which item was cued to be relevant on the current trial (Muhle-Karbe et al., 2021). These results show that in which state an item is stored in WM is highly plastic and brings about the question of what the functional characteristics of these different states could be.

Considering the possible neural mechanisms underlying these silent states can offer further insights into this question. Theoretical accounts have previously proposed short-term synaptic plasticity (STSP) as a potential mechanism underlying the storage of the silent traces (Kozachkov et al., 2022; Miller et al., 2018; Stokes, 2015). Thus, after initial encoding items could be stored in rapid changes in synaptic weights within an assembly of neurons. Computational evidence for these ideas comes from Mongillo et al. (2008), who implemented a model where residual calcium from the initial encoding phase of WM results in short-term changes to the synaptic weights, which produces a synaptic trace of the encoded memory item that can be readout at a later time point. Furthermore, the inclusion of such STSP mechanisms in an artificial recurrent neural network was shown to be necessary to produce firing patterns similar to those observed in primates during WM maintenance, as well as to attain robustness against network degradation (Kozachkov et al., 2022). Finally, a functional spiking-neuron model with a calcium-based STSP implementation developed by Pals and colleagues (2020) was able to reproduce the patterns of working memory maintenance previously seen in EEG recordings of human participants (Wolff et al., 2017). Thus, it is proposed that STSP could be the underlying mechanism of activity-quiescent working memory states.

With accumulating evidence for these distinct states in working memory, it will be important to understand what the function of these states in relation to successful working memory storage

could be. What are the advantages of having different ways to store working memory content in the brain, and in which situation is each mechanism important? While it has previously been argued that a low energy or quiescent state of working memory maintenance would be more energy efficient (Stokes, 2015), one could also wonder whether there could be functional benefits of these distinct stages. One hypothesis could be that having a synaptic trace of information could provide robustness and possibly help to protect items from interference of ongoing sensory input compared to information that is uniquely stored in ongoing firing (Kozachkov et al., 2022; Lundqvist et al., 2018; Miller et al., 2018).

In the current set of experiments, we set out to test the hypothesis that activity-quiescent states of working memory would protect items from upcoming interference, as measured from task performance. Based on the findings by Muhle-Karbe et al. (2021), we assumed that an item that is immediately task relevant would be stored in a neuronally active state, whereas items that would only be relevant at a later point in time, would be stored in an activity-quiescent state. It could then be argued that a passive item that is stored in a quiescent state, relying on a synaptic trace, would be less likely to be interfered with by incoming sensory information, compared to an active item that is stored in ongoing neuronal activity.

To test this hypothesis, we implemented a block design, where one memory item was cued to be probed first (prioritised item), whereas another item was cued to be probed second (deprioritised item). During the maintenance of these items, a distracting visual discrimination task was presented. We manipulated the overlap between this task and the prioritised and deprioritised memory items, matching either the feature (colour or orientation), the spatial location, or both. Although spatial location was not a relevant property of the memory items, it might nonetheless be part of its low-level representation in visual areas subserving WM. Thus, to cover both more task-related as well as sensory aspects of the representation, we included both featural and spatial interference.

Our results show that behavioural performance on the memory task was significantly impacted by the interference task, but that the magnitude of this effect was not different depending on the priority status of the items. Nevertheless, featural interference produced a different pattern of interference for prioritised items than for deprioritised items: For the former, the precision of the response was reduced, whereas for the latter, the guess rate was increased. This suggests that functionally active and passive items may indeed rely on different neural mechanisms, and that although these respond to a similar degree to interference, they do so in different ways.

Materials and Methods

Participants

In total, 92 participants (mean age: 20.7; 66% female) were either recruited as first year psychology students participating in the experiment for course credit, or through the paid-participant pool of the department where they received monetary reward in exchange for participation. The study was conducted in accordance with the Declaration of Helsinki (2008) and

was approved by the Ethical Committee of the Faculty of Behavioural and Social Sciences at the University of Groningen (Study Code = PSY-2122-S-0206).

Apparatus and stimuli

The experiment was programmed using OpenSesame, a freely available software tool for designing experiments (Mathôt et al., 2012). Stimuli were presented using a 19-inch CRT screen, at a resolution of 1,280 x 1,024 pixels with a refresh rate of 100Hz. Participants were seated approximately 60 cm away from the screen in a sound-attenuated testing chamber with dimmed lights. A grey background (RGB: 128) was maintained throughout the experiment. The experimental stimuli covered a circular area of 6.69° degrees of visual angle on the screen. Gabor patches were presented at a contrast of 80% relative to the background, with a spatial frequency of 0.05. Colours were presented on the CIElab colour space (L*=70, a = 0, b = 0, radius = 38).

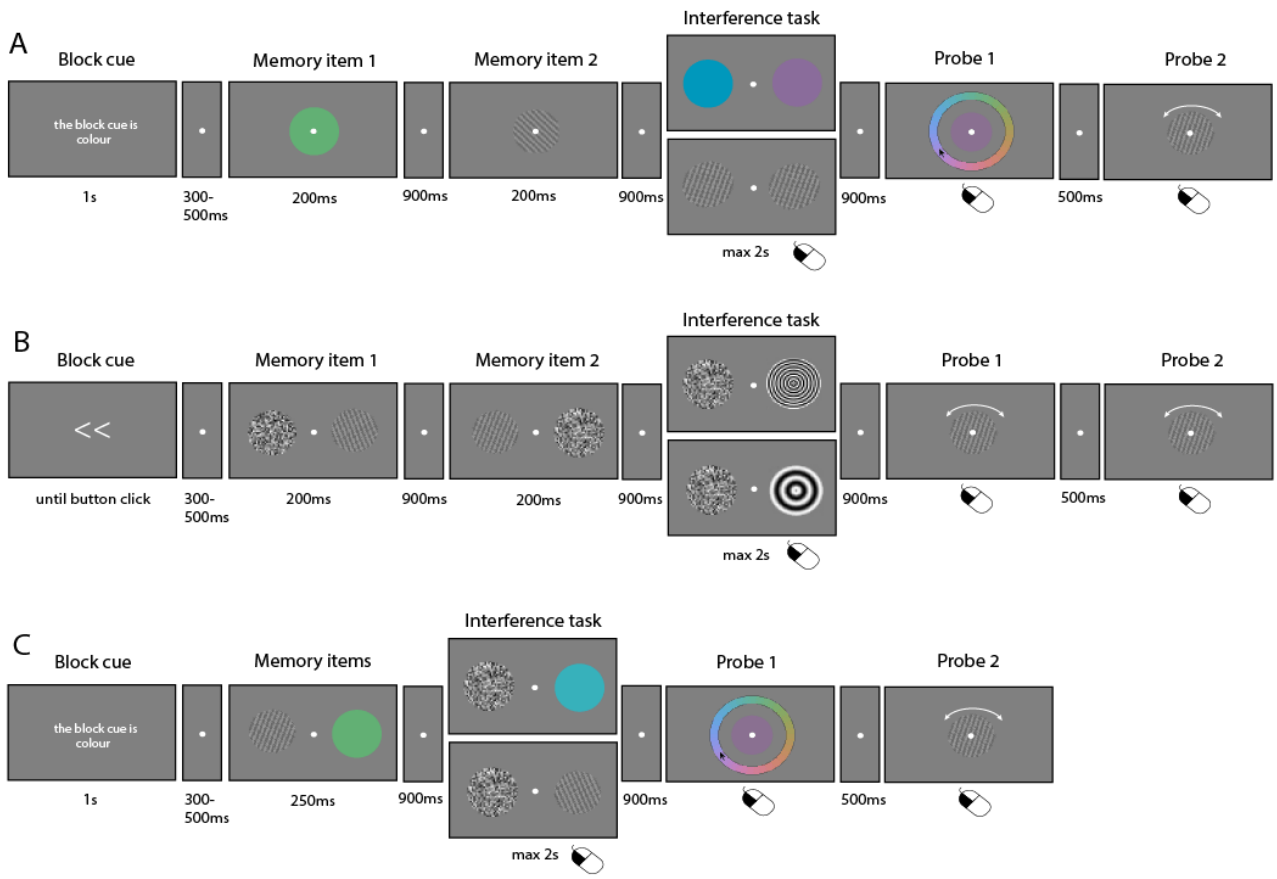

Figure 1: Overview of the trial structure for each experiment. **A**: Experiment 1, **B**: Experiment 2, **C**: Experiment 3.

Experiment 1

In experiment 1 (N= 30), we wanted to test if memory items in a prioritised (cued) state were more prone to featural interference compared to items in a deprioritised (uncued) state. Participants received a block cue, either stating ‘colour’ or ‘orientation’ which was indicative of the response order for the two memory items at the end of each trial. On any given trial, participants

were always presented with one colour as a memory item and one orientation as a memory item (Figure 1A).

For the interference task in this experiment, participants were presented with two items presented laterally on the screen at the same time. The interference items could either consist of two coloured circles or two Gabor patches. The task was to judge if the two items presented on the screen were the same (i.e., the same colour or same orientation) or different from each other. At the end of a trial, participants had to respond to the memory items by selecting the correct colour from a colour wheel for the coloured memory item, and by rotating a Gabor patch in the correct orientation that matched the orientation memory item.

The logic behind the interference manipulation was as follows: On a given trial, if the feature of the prioritized item matched with feature of the interference task, we would call this an interference trial as there is a chance of feature-based interference. If the feature of the prioritized item did not match with the feature of the interference task; we call this a no interference trial for the prioritised item. The same logic applies to the deprioritised item.

Participants completed 320 experimental trials divided into 16 blocks, which took approximately 75 minutes to complete. All manipulations were randomised and counterbalanced.

Experiment 2

The research question of experiment 2 (N=28) was identical to experiment 1, but rather than interfering with memory items at the featural level, the manipulations were now location-specific. Again, we used a block design where we now cued the location of the item that would be probed first at the end of the trial. On any given trial, the participant would then see two memory items presented laterally at an equal distance of 5.93 degree of visual angle from the fixation dot (Figure 1B). The memory items were presented one at a time with a noise patch being presented in the empty location. The noise patch consisted of a Gaussian noise patch. The memory items were Gabor patches where participants again had the task to remember their orientation. The interference task was presented at the same locations as the memory items. For the interference task, in one location the participant would see a black-and-white 'bull's eye' stimulus, and the task was to judge if the spatial frequency of this item was high or low. The corresponding spatial frequencies were 0.5 (low spatial frequency) and 1.4 (high spatial frequency). The 'empty' location would be filled with a noise patch again. At the end of the trial, the participant had to respond to both memory items in the order indicated by the block cue. The response required them to orient a Gabor patch to the remembered orientation.

As in the previous experiment, the logic of interference or no interference conditions remains the same but it is now not dictated by a match of feature dimension between a memory item and the interference task but rather by a match or non-match of spatial location with the interference item to make it an interference or no interference trial, respectively.

Participants completed 320 experimental trials divided into 20 blocks, and it took approximately 75 minutes to complete the experiment. All manipulations were randomised and counterbalanced.

Experiment 3

The aim of experiment 3 (N=34) was to combine the manipulations of experiment 1 and 2. Thus, here the memory items again were one colour and one orientation item (as in experiment

1), but now they were presented in distinct locations (as in experiment 2). This allowed for double overlap in the interference task, as we were now able to match an item on both the featural and spatial level. We used a block design again, where we cued the feature dimension that told participants which memory item would be probed first at the end of the trial. Memory items were presented laterally at the same time (Figure 1C).

For the interference task, participants were now shown either one coloured circle or one Gabor patch. This item was presented in one of the two locations of the memory items. The other location contained only a noise patch. The task of participants was to indicate on which side of the screen (thus, left or right) this item appeared. At the end of the trial, the responses to the memory items were given as in Experiment 1.

The interference manipulation in this experiment combined feature-based and location-based interference. Thus, in the interference condition the interference stimulus overlapped with both the spatial location and feature-dimension of a given item. In the no interference condition, both the spatial location and the feature-dimension differed.

Participants completed 320 experimental trials, divided into 16 blocks. On average it took participants 75 minutes to complete the experiment. All manipulations were randomised and counterbalanced.

Procedure

Participants were given written and verbal instructions and signed written informed consent. Participants started with several practice trials, after which they were given the opportunity to ask additional questions before moving on to the experimental trials. The experiment was divided into blocks, and at the beginning of each block the participants received a block cue for 1s which was predictive of the response order for the memory items at the end of the trial, as shown in Figure 1. Thus, for example for experiment 1, if the block cue was colour, participants first had to respond to the memorised colour at the end of the trial and to the memorised orientation second, or vice versa if the block cue was orientation. Each trial consisted of an initial fixation dot between 300-500ms (uniformly sampled), followed by the presentation of the first memory item for 200ms. After a fixation dot of 900ms, the second memory item was presented for 200ms. Another fixation dot was presented for 900ms, after which the interference task was presented. The interference task differed between the three experiments, but essentially here participants had to perform a simple visual discrimination task (details for each experiment below). Participants had a maximum of 2s to respond to this task. After a response was given, a fixation dot was presented for 900ms, after which the response screen for memory item 1 appeared, followed by the response screen for memory item 2. The response screens were separated by a fixation dot, which was presented for 900ms. Participants' response time to respond to the memory items was not limited.

Analysis

The data of all three experiments were analysed using the same approach. Participants whose response error was +/-2SD larger than the mean, or participants who performed (below) chance level in the interference task were excluded from the analysis (Exp1: N=4, Exp2: N=3, Exp3: N=1).

To quantify memory performance of participants we calculated the absolute error response for each trial for each memory item type (prioritised versus deprioritised), and each condition (match versus no match). To test for differences between conditions on each probe we used paired-samples Bayesian t-tests, which allowed us to quantify evidence for an interference effect (H_A), as well as to quantify evidence for no effect of interference (H_0).

To quantify if there are any differences in magnitude between the interference effect on cued versus uncued items we calculated z-scores and statistically compared those - we chose to compare z-scored responses rather than absolute error since the absolute error for the uncued item was on average higher. Thus, we z-scored all individual absolute error responses and calculated an average z-score for the interference and no interference condition per participant. Again, we used Bayesian t-tests to test for statistical differences in the magnitude of interference.

We furthermore wanted to quantify the nature of interference by fitting the two-component standard mixture model (Zhang & Luck, 2008). We fitted independent models for both prioritised and deprioritised items for each condition using the mixture package (Grange & Moore, 2022; version: 1.2.1). The standard mixture model provides two parameters based on the error distribution. It provides the probability of an item being stored in memory (P_m), with smaller values of P_m indicating an increase in guessing responses. As a second parameter it provides the precision (k) of the memory representation. Relating this to our experiments, this provided us the opportunity to test if interference leads to a gradual imprecision of the memory representation, which would be reflected in an effect on k, or if interference leads to a complete loss of the memory representation, which would be reflected in an effect on P_m. To test statistical differences between the parameters we used paired-sample Bayesian t-tests.

Pre-processing, mixture modelling, and plotting of the data was performed using R (R Core Team, 2018; version 4.3.2). Bayesian statistics were calculated using JASP (JASP Team, 2023; version: 0.18.1), and Bayesian values were interpreted based on Wetzels et al. (2011).

Results

Experiment 1

In experiment 1, we tested whether feature-based interference differed between prioritised and deprioritised items held in working memory. On average, we found strong evidence for better performance on prioritised items (μ_{AS} = 18.1°, sd_{AS} = 4.3°), compared to deprioritised items (μ_{AS} = 20.5°, sd_{AS} = 4.2°) (t_{paired} BF_{10} = 25.87). Furthermore, we observed definite evidence that participants performed above chance level for the interference task (μ = 89.03%, sd = 5.4%; $t_{one\text{-}sided}$ BF_{10} > 150), indicating that participants paid attention to the stimuli.

With respect to the interference, the absolute error for the prioritised item was larger when it matched the feature dimension of the interference stimulus, than when it did not (t_{paired} BF_{10} = 137.28, Figure 2A). For the deprioritised item, we found the same result (t_{paired} BF_{10} > 150, Figure 2B). Thus, we found an effect of interference both for prioritised as well as deprioritised items. To test if the magnitude of interference differed between the two states while controlling for differences in average error between Probe 1 and Probe 2, we compared z-scored values of the average error (t_{paired} BF_{10} = 0.224 (BF_{01} = 4.461), Figure 2C). Intruigingly, contrary to our

hypothesis, there seemed to be no difference in the magnitude of interference between prioritised and deprioritised items.

As a next step, we were interested in what drives the interference effect, and whether the underlying mechanisms are similar for prioritised and deprioritised items. We therefore fitted two separate standard-mixture models, one for the prioritised item and one for the deprioritised item (Zhang & Luck, 2008). For prioritised items we found very strong evidence that *k* differed depending on the interference condition (t_{paired} BF_{10} = 58.291, Figure 2D), while $P_{(m)}$ remained unaffected (t_{paired} BF_{10} = 0.47 (BF_{01} = 2.130), Figure 2E). However, for the deprioritised item we did not find evidence for an effect of interference on precision (t_{paired} BF_{10} = 0.879 (BF_{01} = 1.138), Figure 2F), but found anecdotal evidence for an effect on $P_{(m)}$ (t_{paired} BF_{10} = 1.159, Figure 2G). This led us to hypothesise that while both prioritised and deprioritised items seem susceptible to interference, the underlying mechanisms causing this behavioural impairment might be different between priority states.

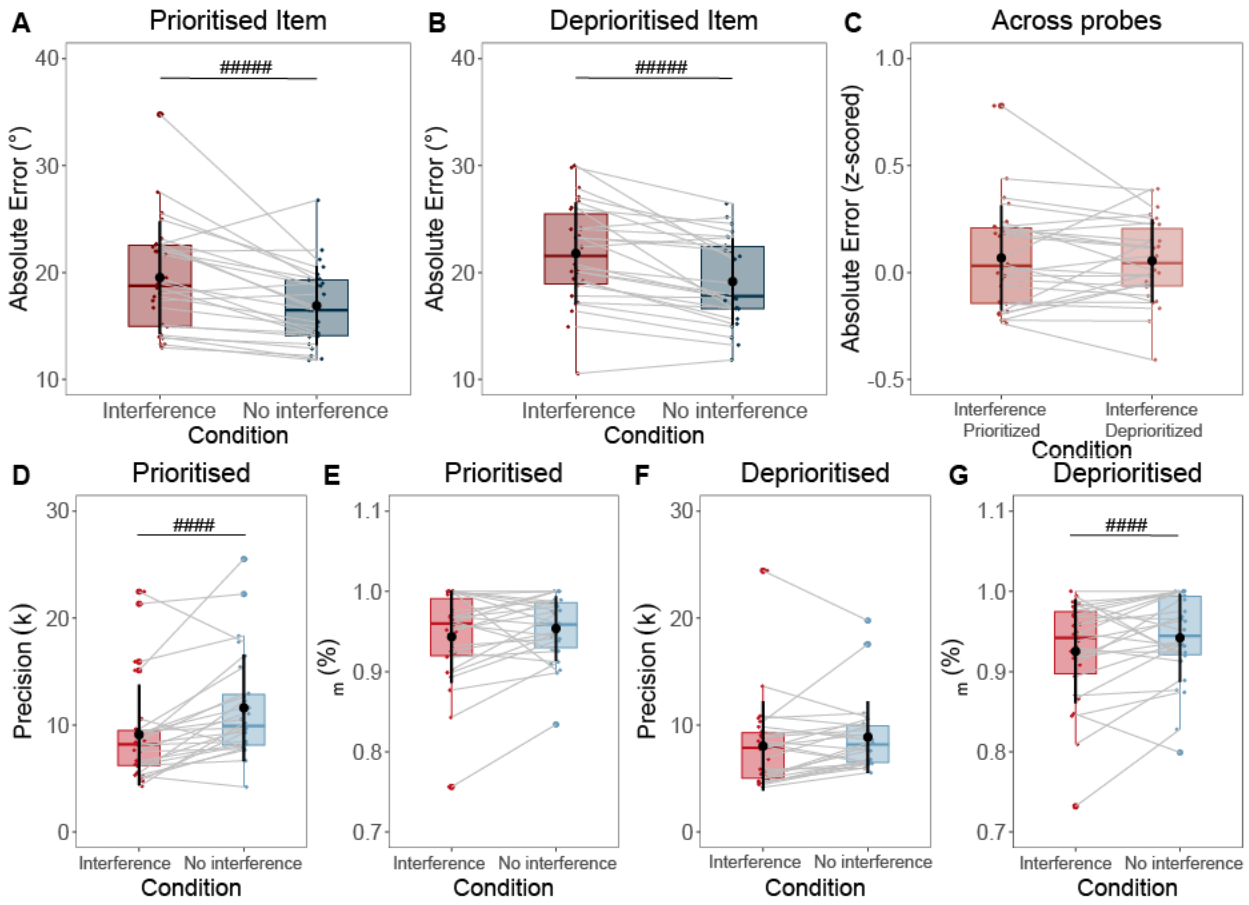

Figure 2: Results for experiment 1. (A) Absolute error for the prioritised item for the interference versus no interference condition. (B) Absolute error for the deprioritised item for the match versus no match condition. (C) Z-scored absolute error for the match condition across both probes. (D) Precision (*k*) for the prioritised item for match versus no match condition. (E) Guess rate ($P_{(m)}$) for the prioritised item for match versus no match condition. (F) Precision (*k*) for the deprioritised item for match versus no match condition. (G) Guess rate ($P_{(m)}$) for the deprioritised item for match versus no match condition. Coloured dots represent individual data points, black dots indicate averages and error bars are +/- SD. Red colour represents the match condition, blue colour represents no match condition. Grey lines connect each subject in the match versus no match condition. Boxplots show median and quartile values. ##### indicates decisive

evidence, #### indicates very strong evidence, ### indicates strong evidence, ## indicates substantial evidence and # indicates anecdotal evidence.

Experiment 2

In experiment 2, we again investigated the effects of interference between prioritised and deprioritised items in working memory; however, this time we interfered in a spatial-specific, rather than feature-specific manner. In accordance with the previous experiment, we found very strong evidence for better performance on the prioritised item (μ_{AS} = 27.71°, sd_{AS} = 4.9°) compared to the deprioritised item (μ_{AS} = 37.8°, sd_{AS} = 7.8°) (t_{paired} BF_{10} > 150). We also observed very strong evidence that participants performed above chance in the interference task (μ = 97.3%, sd = 2.5%; $t_{one\text{-}sided}$ BF_{10} > 150). In this paradigm, we found only anecdotal evidence that interfering with the prioritised item led to a larger response error (t_{paired} BF_{10} = 2.671, Figure 3A). However, we observed no effect of interference on the deprioritised item (t_{paired} BF_{10} = 0.393 (BF_{01} = 2.544), Figure 3B). Overall, there was no evidence for differential interference in working memory between the prioritised and deprioritised items (t_{paired} BF_{10} = 0.370 (BF_{01} = 2.703), Figure 3C).

Considering the mixture-model of visual working memory, we found no evidence for an effect of interference on *k* (t_{paired} BF_{10} = 0.488 (BF_{01} = 2.050), Figure 3C), nor on $P_{(m)}$ (t_{paired} BF_{10} = 0.237 (BF_{01} = 4.214), Figure 3D) for the prioritised item. This was somewhat unexpected, as we did find an effect of interference on *k* in experiment 1. We found no evidence for an effect of interference on *k* (t_{paired} BF_{10} = 0.211 (BF_{01} = 4.738), Figure 3E), nor on $P_{(m)}$ (t_{paired} BF_{10} = 0.235 (BF_{01} = 4.246), Figure 3F) for the deprioritised item, which was in line with our expectations, as we did not see an interference effect for the deprioritised item in this paradigm, and thus would not expect either parameter to differ between interference conditions.

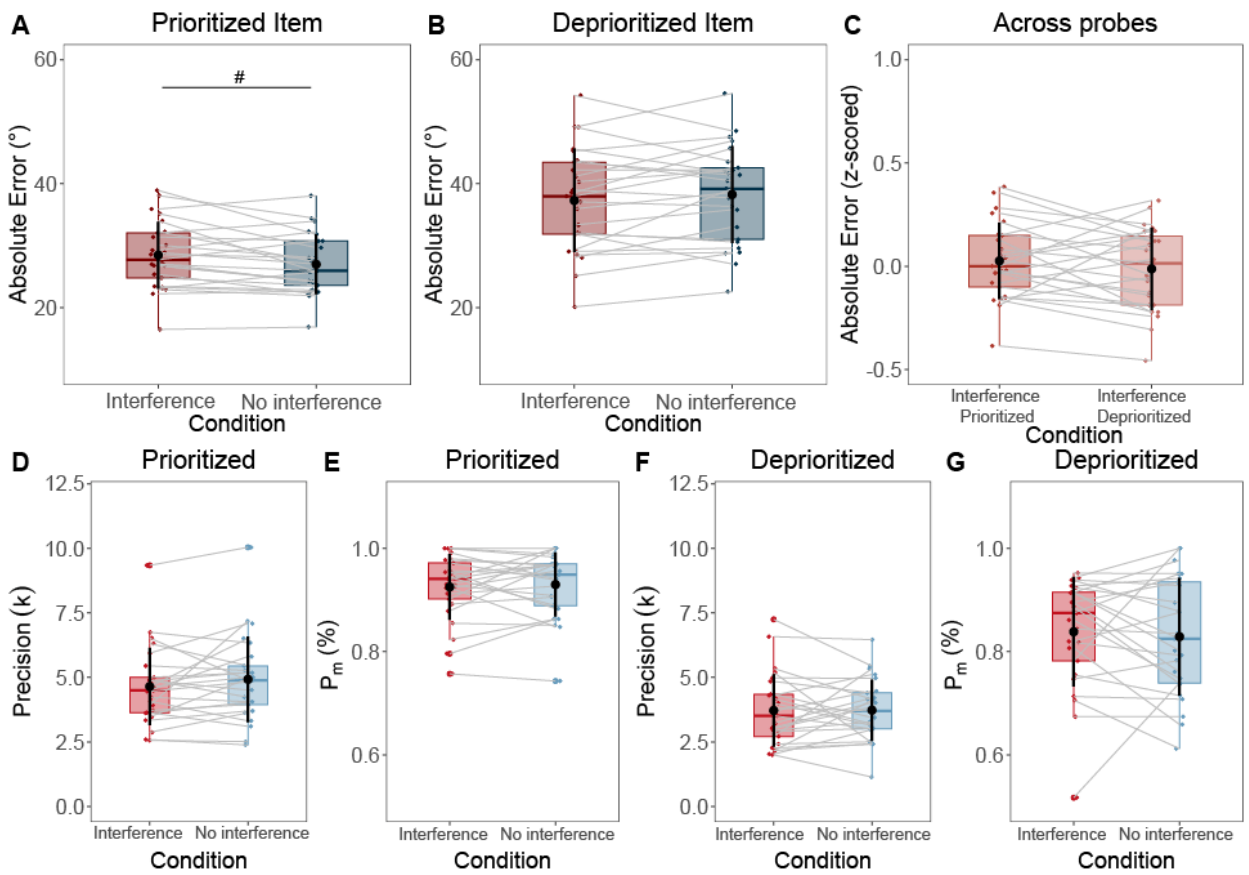

Figure 3: Results for experiment 2. (A) Absolute error for the prioritised item for the interference versus no interference condition. (B) Absolute error for the deprioritised item for the match versus no match condition. (C) Z-scored absolute error for the match condition across both probes. (D) Precision (*k*) for the prioritised item for match versus no match condition. (E) Guess rate ($P_{(m)}$) for the prioritised item for match versus no match condition. (F) Precision (*k*) for the deprioritised item for match versus no match condition. (G) Guess rate ($P_{(m)}$) for the deprioritised item for match versus no match condition. Coloured dots represent individual data points, black dots indicate averages and error bars are +/- SD. Red colour represents the match condition, blue colour represents no match condition. Grey lines connect each subject in the match versus no match condition. Boxplots show median and quartile values. ##### indicates decisive evidence, #### indicates very strong evidence, ### indicates strong evidence, ## indicates substantial evidence and # indicates anecdotal evidence.

Experiment 3

In experiment 3, we combined the manipulations of the previous two experiments. Thus, as in experiment one we matched the memory and interference stimuli by their feature dimension, and as in experiment 2, we also matched them by their spatial location. Therefore, rather than presenting both memory items at the centre of the screen serially, we presented them laterally at the same time, to add a unique spatial location to each item. Overall, participants performed better on the prioritised item (μ_{AS} = 21.9°, sd_{AS} = 6.4°), compared to the deprioritised item (μ_{AS} = 27.3°, sd_{AS} = 7.7°) (t_{paired} BF_{10} > 150). As before, we also found high accuracy for the behavioural response to the interference task (μ = 97.6%, sd = 2.2%; $t_{one\text{-}sided}$ BF_{10} > 150), indicating that participants were paying attention to the items presented.

Looking at the effect of the interference manipulation, for the prioritised item we found anecdotal evidence for interference (t_{paired} BF_{10} = 1.16, Figure 4A), while in this paradigm we also

found decisive evidence of interference for the deprioritised item (t_{paired} BF_{10} = 118.79, Figure 4B). To check if the magnitude of interference differed between the prioritised and deprioritised items we calculated the z-score again, but as in the previous experiments, we found no difference in the magnitude of interference between prioritised and deprioritised items (t_{paired} BF_{10} = 0.230, Figure 4C), indicating that the degree of interference was comparable between the two states.

Interestingly, while we did find (anecdotal) effect of interference for the prioritised item, we did not find changes for *k* (t_{paired} BF_{10} = 0.488 (BF_{01} = 2.051), Figure 4D), nor for $P_{(m)}$ (t_{paired} BF_{10} = 0.225 (BF_{01} = 4.439), Figure 4E) between interference conditions, just as in experiment 2. For the deprioritised item we found that interference differentially affected $P_{(m)}$ (t_{paired} BF_{10} = 1.921, Figure 4F). There was no effect on *k* (t_{paired} BF_{10} = 0.953 (BF_{01} = 1.049), Figure 4G). This is in accordance with experiment 1, suggesting that if there is an interference effect on the deprioritised item this mainly leads to an increase in guessing responses.

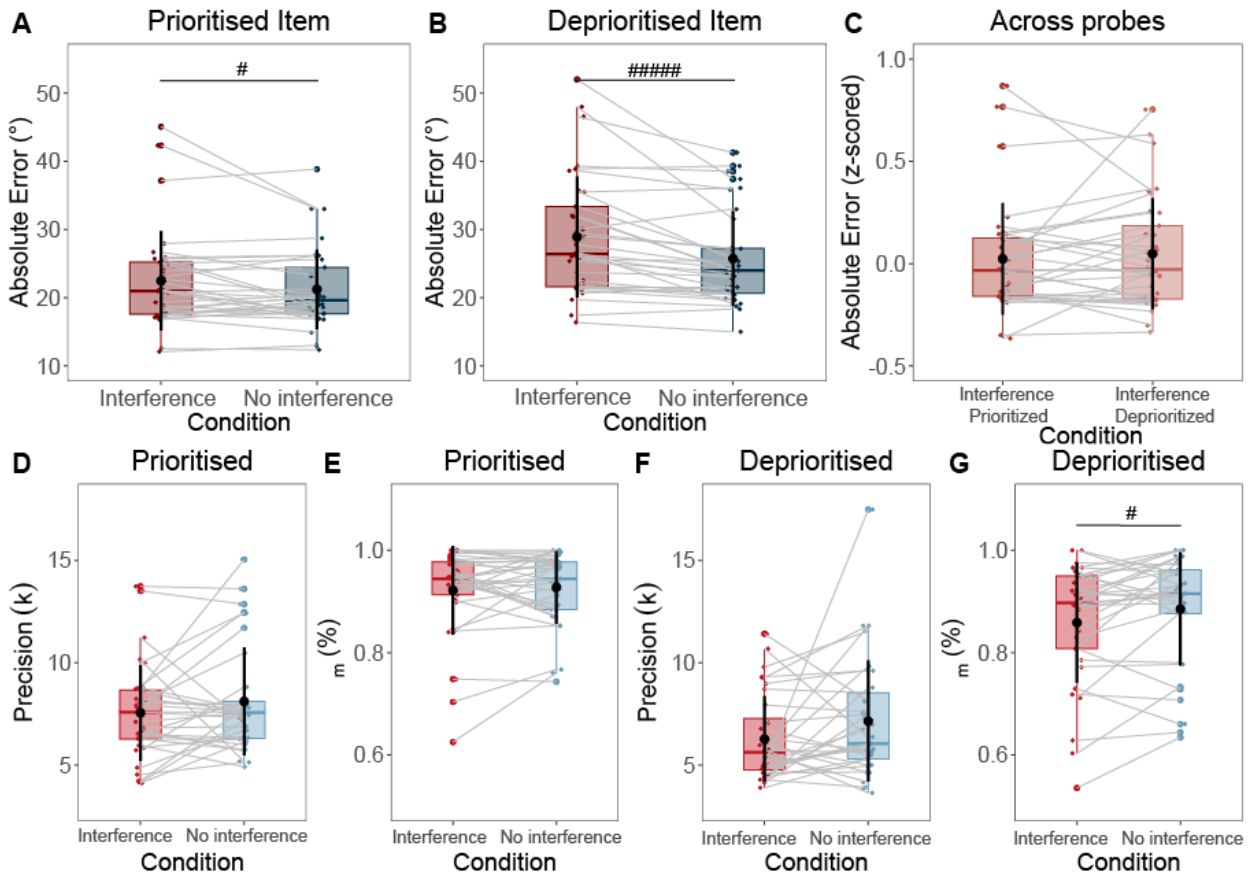

Figure 4: Results for experiment 3. (A) Absolute error for the prioritised item for the interference versus no interference condition. (B) Absolute error for the deprioritised item for the match versus no match condition. (C) Z-scored absolute error for the match condition across both probes. (D) Precision (*k*) for the prioritised item for match versus no match condition. (E) Guess rate ($P_{(m)}$) for the prioritised item for match versus no match condition. (F) Precision (*k*) for the deprioritised item for match versus no match condition. (G) Guess rate ($P_{(m)}$) for the deprioritised item for match versus no match condition. Coloured dots represent individual data points, black dots indicate averages and error bars are +/- SD. Red colour represents the match condition, blue colour represents no match condition. Grey lines connect each subject in the match versus no match condition. Boxplots show median and quartile values. ##### indicates decisive

evidence, #### indicates very strong evidence, ### indicates strong evidence, ## indicates substantial evidence and # indicates anecdotal evidence.

Discussion

Working memory can be stored in activity-dependent or activity-quiescent states; with an accumulating amount of evidence showing that these two mechanisms of WM storage are not mutually exclusive. However, little research to date has focused on the potential functional roles of these differing states in WM. In the current manuscript, we set out to test one potential functional role of these varying states – namely, the hypothesis that activity-quiescent states could be protecting items from interference.

For this, we tested the effect of an interference task, presented during WM maintenance in three experiments. We were primarily interested in potential differences between prioritised (active) and deprioritised (passive) items and hypothesised that different neural mechanisms might underlie the maintenance of active and passive items. Active items might be held by means of persistent neural firing, which could be mediated by internal attention (Bae & Luck, 2018; Lewis-Peacock & Postle, 2012). Conversely, passive items might be retained in an activity-quiescent manner, through functional connectivity that results from short-term synaptic plasticity (Mongillo et al., 2008; Stokes, 2015). We furthermore hypothesised that activity-quiescent maintenance through functional connectivity might afford better protection from interference (Kozachkov et al., 2022; Lundqvist et al., 2018; Miller et al., 2018). Our experiments provided a test of these two linked ideas.

The results of our experiments provided partial support for the hypotheses. On the one hand, the finding that there were qualitative differences in the effect of featural interference on working memory, dependent on whether the items were prioritised or deprioritised is at least compatible with the idea that these items were also held in different neural states. On the other hand, the prediction that deprioritised items would be less susceptible to interference was not borne out. Specifically, although the precision of the representation seemed to decline for prioritised items, the loss of the complete representation was at stake for deprioritised items. Overall, the net effect of these two kinds of impairment was approximately the same.

The effect of interference on the precision of prioritised items may point to the involvement of attention in storing items in a prioritized state. In a recent study on awareness during the attentional blink, we observed that having a broad attentional scale reduced the precision of target item reports (Karabay et al., 2022). Having a narrow attentional scale affected the guess rate exclusively. Importantly, the overall level of target reports was not affected, which is reminiscent of the current pattern of results. We might thus speculate that when an item is actively held in working memory, in the focus of attention, and attention is subsequently drawn to the stimuli of the interference task, which appeared at different spatial locations, a similar attentional spreading effect might have occurred. In this case, a precision effect would indeed be likely.

Nevertheless, the question remains why the interference effect on deprioritised items was expressed in guess rates, suggesting that the representations were fully lost on some of the trials

(but perfectly preserved on others). It has been argued that activity-quiescent working memory mechanisms should be more robust to interference, and if anything, show graceful degradation (Kozachkov et al., 2022; Lundqvist et al., 2018; Miller et al., 2018). One possible explanation might be that (episodic) long-term memory (LTM) was involved when items were deprioritised, even though it should be noted that the present load on working memory was not particularly high (Buschman & Miller, 2022; Oberauer & Awh, 2022). Episodic LTM would not necessarily have to respond differently to interference, but it is a system that might be more prone to interference from past items (i.e., from previous trials). Previous work has shown that regions of activated LTMs and WM content overlap and that LTM can influence WM, but this might only happen when WM representations are already so weak that reliance on LTM is necessary (Oberauer et al., 2017; Vo et al., 2021). Again speculatively, we may thus suppose that the specific combination of the present interference task and the trial history together might have led to an influence of LTM on a subset of trials driving the complete loss of deprioritised representations in certain cases, explaining the observed guess rate effect.

The current study has several limitations that should be noted. First, we did not take any direct neural measures, which means that our interpretations as to potential underlying mechanisms remain indirect inferences. Second, we only tested memory for a pair of items, which is a relatively modest load on working memory. It is conceivable that a higher load may produce further evidence with regard to the involvement of LTM, as higher loads would necessitate working memory offloading. Third, we only manipulated two feature dimensions (colour and orientation), next to spatial location. A more fine-grained and comprehensive test of the properties of items in working memory is needed to confirm and extend the present outcomes.

Based on the evidence presented in this manuscript, protecting of interference does not seem to be a function of activity-quiescent states. However, this does not mean that there could not be other functional benefits of contrasting neuronal mechanisms within WM. One hypothesis could be that rather than protection from incoming sensory information, the two states of WM could function to protect various items in WM to interfere with each other – this possibility is something that should be explored in future work. Additionally, while we did not find a difference in magnitude of interference between the two neuronal states, we did observe qualitative differences in the nature of interference that future work should link to neuronal signatures. The existence of these differences also underscore that it might be important to investigate these mechanisms at higher loads of WM, where difficulty to maintain accurate representations decreases and the need for protection is higher. Doing so might further our understanding of the functional relevance of activity-dependent versus activity-quiescent WM states.

Conclusion

Active (prioritised) and passive (deprioritised) items in working memory are similarly affected by interference from a distracting task during maintenance. However, featural interference is expressed as a precision effect on prioritised items, and as a guess rate effect on deprioritised items, suggesting that these functional states might indeed be supported by different neural mechanisms, such as persistent activity and short-term synaptic plasticity, respectively.

Open Practices Statement:
The cleaned datasets analysed in the current manuscript, plus all prepared datasets (restructured raw data and the output files of the behavioural modelling) that were used for statistical testing in JASP are available on Open Science Framework (https://osf.io/qpt5f/?view_only=644630a919ff4394905a20e7a296fb89). None of the experiments were pre-registered.

Declarations*:*

Funding:
Yuanyuan Weng was funded by the Chinese Scholarship Council.

Conflict of interest:
We have no conflict of interest to disclose.

Ethics approval:
These experiments were approved by the Ethics Committee at the Faculty of Behavioural of Social Sciences of the University of Groningen. The study was conducted in accordance with the 2013 Declaration of Helsinki.

Consent to participate:
All participants provided active, written informed consent before participating any of the experiments.

References

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 8, pp. 47–89). Academic Press. https://doi.org/10.1016/S0079-7421(08)60452-1

Bae, G.-Y., & Luck, S. J. (2018). Dissociable Decoding of Spatial Attention and Working Memory from EEG Oscillations and Sustained Potentials. *The Journal of Neuroscience*, *38*(2), Article 2. https://doi.org/10.1523/JNEUROSCI.2860-17.2017

Buschman, T. J., & Miller, E. K. (2022). Working Memory Is Complex and Dynamic, Like Your Thoughts. *Journal of Cognitive Neuroscience*, 1–7. https://doi.org/10.1162/jocn_a_01940

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron*, *14*(3), 477–485. https://doi.org/10.1016/0896-6273(95)90304-6

Grange, J. A., & Moore, S. B. (2022). mixtur: An R package for designing, analysing, and modelling continuous report visual short-term memory studies. *Behavior Research Methods*, *54*(5), 2071–2100. https://doi.org/10.3758/s13428-021-01688-1

Karabay, A., Wilhelm, S. A., de Jong, J., Wang, J., Martens, S., & Akyürek, E. G. (2022). Two faces of perceptual awareness during the attentional blink: Gradual and discrete. *Journal of Experimental Psychology. General*, *151*(7), 1520–1541. https://doi.org/10.1037/xge0001156

Kozachkov, L., Tauber, J., Lundqvist, M., Brincat, S. L., Slotine, J.-J., & Miller, E. K. (2022). Robust and brain-like working memory through short-term synaptic plasticity. *PLOS Computational Biology*, *18*(12), e1010776. https://doi.org/10.1371/journal.pcbi.1010776

LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Decoding Attended Information in Short-term Memory: An EEG Study. *Journal of Cognitive Neuroscience*, *25*(1), Article 1. https://doi.org/10.1162/jocn_a_00305

Lewis-Peacock, J. A., & Postle, B. R. (2012). Decoding the internal focus of attention. *Neuropsychologia*, *50*(4), 470–478. https://doi.org/10.1016/j.neuropsychologia.2011.11.006

Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L., & Miller, E. K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature Communications*, *9*(1), 394. https://doi.org/10.1038/s41467-017-02791-8

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, *44*(2), 314–324. https://doi.org/10.3758/s13428-011-0168-7

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*(2), 463–475. https://doi.org/10.1016/j.neuron.2018.09.023

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic Theory of Working Memory. *Science*, *319*(5869), Article 5869. https://doi.org/10.1126/science.1150769

Muhle-Karbe, P. S., Myers, N. E., & Stokes, M. G. (2021). A Hierarchy of Functional States in Working Memory. *The Journal of Neuroscience*, *41*(20), 4461–4475. https://doi.org/10.1523/JNEUROSCI.3104-20.2021

Ngiam, W. X. Q., Brissenden, J. A., & Awh, E. (2019). “Memory compression” effects in visual working memory are contingent on explicit long-term memory. *Journal of Experimental Psychology: General*, *148*(8), 1373–1385. https://doi.org/10.1037/xge0000649

Oberauer, K., & Awh, E. (2022). Is There an Activity-silent Working Memory? *Journal of Cognitive Neuroscience*, *34*(12), 2360–2374. https://doi.org/10.1162/jocn_a_01917

Oberauer, K., Awh, E., & Sutterer, D. W. (2017). The role of long-term memory in a test of visual working memory: Proactive facilitation but no proactive interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(1), 1–22. https://doi.org/10.1037/xlm0000302

Pals, M., Stewart, T. C., Akyürek, E. G., & Borst, J. P. (2020). A functional spiking-neuron model of activity-silent working memory in humans based on calcium-mediated short-term synaptic plasticity. *PLOS Computational Biology*, *16*(6), Article 6. https://doi.org/10.1371/journal.pcbi.1007936

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*(6316), 1136–1139. https://doi.org/10.1126/science.aah7011

Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), Article 7. https://doi.org/10.1016/j.tics.2015.05.004

Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E.-J. (2011). Statistical Evidence in Experimental Psychology: An Empirical Comparison Using 855 t Tests. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *6*(3), 291–298. https://doi.org/10.1177/1745691611406923

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*(6), Article 6. https://doi.org/10.1038/nn.4546

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), Article 7192. https://doi.org/10.1038/nature06860