Feature and space-based interference with functionally active and passive items in working memory

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

[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

Running head: Feature and space-based interference in working memory

20/02/2025

Word count: 7069

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

Abstract

Functionally active and passive states in working memory have been 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, in particular that passively maintained items 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 two experiments, on half the trials, we presented an interfering task during memory maintenance, in which the stimuli matched either the feature dimension of the memory items (colour or orientation), or their spatial location. Whether the interfering task appeared on a given trial was unpredictable. Finally, we conducted a third experiment where participants were given prior knowledge of the interference condition. We found that both active and passive memory items were affected by interference to a similar extent, with overall performance being closely matched in all experiments. We further investigated precision and probability of target response parameters from the standard mixture model, which also showed no differences between states. We conclude that active and passive items, although potentially stored in different neuronal states, do not show differential susceptibility to interference.

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

Public significance statement

The ability to briefly remember information is critical to human cognition. Our so-called working memory is nevertheless rather limited, able to hold only a few items at any one time, and prone to forgetting when we are briefly distracted. Yet, there is reason to believe that not all information in working memory is equally vulnerable. Items that are more passively stored, because they will only be required after some time, might be more resilient to interference. Items that are stored actively, for more immediate recall, might be more easily disrupted. Here, we investigated the effect of an interference task on the retention of both active and passive items in working memory. Our results showed that active and passive items are equally affected by interference, suggesting that resilience in working memory does not depend on the functional state of the items therein.

Feature and space-based interference with functionally active and passive items in working memory

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 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 for maintenance mechanisms of working memory that do not exclusively rely on persistent firing. Rather, there also seems to be an activity-silent, or at least activity-quiescent, state that can support the maintenance of working memory content (Stokes et al., 2015).

Theoretical models have proposed short-term synaptic plasticity (STSP) as a potential mechanism underlying the maintenance of quiescent WM traces (Kozachkov et al., 2022; Miller et al., 2018; Stokes, 2015). According to this view, after initial encoding items could be stored in rapid changes in synaptic weights within an assembly of neurons. Computational evidence supports this hypothesis: Mongillo and colleagues (2008) demonstrated that residual calcium from the initial encoding phase of WM could produce synaptic traces via STSP that could later be read out. Similarly, the inclusion of such STSP mechanisms in an artificial recurrent neural network was shown to be necessary to produce firing patterns like 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). Collectively, these findings suggest that STSP provides a plausible mechanism for activity-quiescent WM states.

It seems reasonable to assume that both persistent-activity and activity-quiescent states support working memory, because they serve different functional roles. Persistent-activity states have been associated with functionally active items, that is, items maintained in the focus of attention and/or prioritised items needed for immediate recall. Activity-quiescent states have in turn been associated with functionally passive items; items which are currently deprioritised, outside the focus of attention and only needed for later recall. For example, in a classic two-item retro-cueing paradigm, an item that is known to be relevant only later in the trial (i.e., deprioritised) first enters an activity-quiescent state, but is transformed back into a persistent-activity state once a response to the prioritised item has been made, and a response to the formerly deprioritised item will soon be needed (LaRocque et al., 2012; Wolff et al., 2017). Furthermore, it has also been shown that two items can adaptively be moved back between persistent-activity and activity-quiescent states on a single-trial basis, based on which item was cued to be relevant on the current trial (Muhle-Karbe et al., 2021).

The existence of these two functional states prompts the question of what the advantages are of having different ways to store items in WM in the brain, and in which situation these might come into play. While it has previously been argued that a quiescent state of WM maintenance would be more energy-efficient (Stokes, 2015), one could also wonder whether there could be functional benefits of these distinct states. One hypothesis is that quiescent states might offer greater protection against interference from incoming sensory information, as synaptic traces are less dynamic and therefore potentially less vulnerable than ongoing neuronal firing (Kozachkov et al., 2022; Lundqvist, Herman, Warden, et al., 2018; Miller et al., 2018). This is what we set out to test in the present study.

Prior research has yielded mixed findings regarding the susceptibility of prioritised versus deprioritised items to interference within WM. On the one hand, some studies suggest that prioritised items are protected from interference (Makovski et al., 2008; Makovski & Pertzov, 2015; Schneider et al., 2017; Souza & Oberauer, 2016). On the other hand, evidence also indicates that prioritised items may be more susceptible to interference (Hitch et al., 2018; Hu et al., 2014; Mallett & Lewis-Peacock, 2019). Moreover, some recent work has suggested that there is no difference in how prioritised and deprioritised items are affected by interference (Zhang & Lewis-Peacock, 2023).

Some of this ambiguity might have arisen because interference effects in WM are influenced by both the nature of the distractors and the mode of prioritisation of memory items. Factors such as featural similarity and spatial overlap between distractors and memory representations affect the extent of disruption (Hu et al., 2014; Pinto et al., 2013). For example, suffix tasks or visual distractions have been shown to disrupt WM representations only when distractors overlap with memory features (Allen et al., 2015; Allen & Ueno, 2018; Hu et al., 2014). Furthermore, while prioritisation through retro-cueing or reward-driven selection enhances the accessibility of prioritised items, it does not necessarily confer immunity against interference (Zhang & Lewis-Peacock, 2023).

In the current set of experiments, we employed a block design where memory items were explicitly prioritised (functionally active) or deprioritised (functionally passive). A distractor task, involving visual discrimination, was presented during the maintenance phase, and the interference was systematically manipulated to overlap with the active or passive items. Importantly, these interference manipulations included either featural (color or orientation) or spatial overlap, as low-level sensory representations in WM might encode spatial attributes, even when not explicitly relevant to the task. This approach allows us to make a comparison between feature and space-based interference, and their potentially differential effects on active versus passive states. We also included an experiment in which participants were cued about upcoming

interference, to assess the ability of our participants' to strategically shield active or passive memory representations. Previous research suggests that anticipation and strategic adjustments can impact WM robustness (Banks et al., 2023; Gresch et al., 2021; Klink et al., 2017), yet whether these effects differentially apply to active versus passive states remains unclear.

To preview the results, we consistently found that both active and passive items were susceptible to interference, with no evidence that a passive and potentially quiescent neuronal state provided a different (higher) level of protection against interference. Interference effects were driven by reductions in precision and probability of target responses for both states, and these effects were stronger when the interfering stimulus matched the memory item, regardless of priority order, and regardless of whether participants had explicit knowledge of the interference feature or location, or not. Interference effects thus appear to be robust across memory states and neither active nor passive memory states provide increased protection from interference.

Materials and Methods

Participants

In total, 145 participants (mean age = 21.3, SD = 3.7; 101 female) were either recruited as first year psychology students participating in the experiments for course credit, or through the paid-participant pool of the department where they received monetary reward in exchange for participation. We are interested in fundamental aspects of basic cognition and therefore believe that using young-healthy adults provides an appropriate sample. However, it should be noted that we are only sampling from first-year psychology students, which means our sample is also highly educated, and which might limit the translation to other groups within society. Our sample is also biased between sexes, as approximately 2/3 of our sample is female.

As we performed all analyses using Bayesian statistics, we opted for optional stopping to determine our sample size. We collected data until we found sufficient evidence for either the alternative or the null hypothesis (at least anecdotal evidence in either direction). 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. Orientation gratings were presented at a contrast of 80% relative to the background, with a spatial frequency of 0.05 cycles per degree. Colours were presented on the CIElab colour space (L*=70, a = 0, b = 0, radius = 38).

Figure 1

Overview of the trial structure per experiment

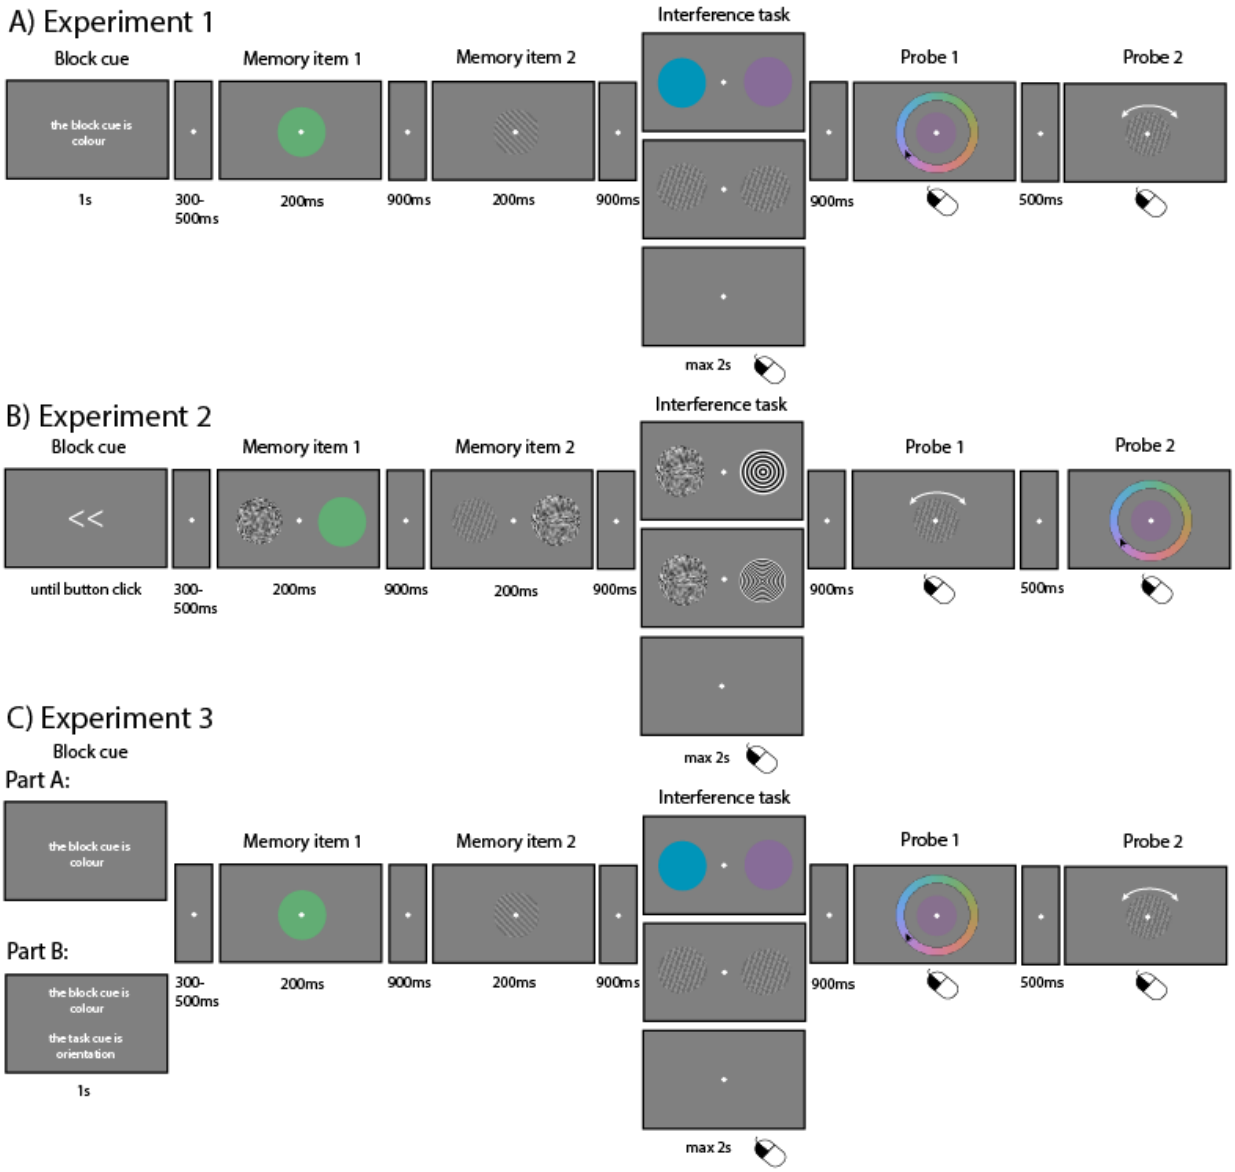

Note: Overview of the trial structure for each experiment. **A**: Experiment 1, **B**: Experiment 2, **C**: Experiment 3.

Experiment 1

In Experiment 1 (N= 51), we wanted to test if memory items in an active (prioritised) state were more prone to feature-based interference compared to items in a passive (deprioritised) state. At the start of the experiment, 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 questions before moving on to the experimental trials.

The experiment was divided into blocks. Each block started with 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). The presentation order of the memory items was counterbalanced but randomized and independent from the block cue; thus, the block cue did not indicate if the colour or orientation memory item was presented first but was purely informative for response order at the end of the trial. The block cue was only shown before the first trial of each block, participants then did 24 trials.

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. For the interference task participants were presented with two items 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. The two items were the same for 50% of the trials, and different for the other 50% of trials. We drew colours and orientations randomly from the full circle; thus between 1-180 degrees for orientations and 1-360° degrees for colours. Participants indicated their response by using the different mouse buttons to indicate if the items presented were the same or different from each other (left and right mouse buttons representing same or different conditions was counterbalanced between odd and even participant numbers). Participants had up to 2000ms to give a response, until the experiment would move on. We also included a baseline condition where instead of the interference display, the fixation dot stayed on the screen for 2000ms and participants were instructed to do nothing but to keep looking at the screen. After a response was given in the interference conditions, 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. Participants could give their response to 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 response screens for response 1 and response 2 were separated by a fixation dot, which was presented for 900ms. Participants' response time to respond to the memory items was not limited.

The logic behind the interference manipulation was as follows: On a given trial, if the feature dimension of the active item matched with that of the interference task, we call this a match trial for the active item. If the feature dimension of the active item did not match with the feature dimension of the interference task; we call this a no match trial for the active item. The same logic applies to the passive item. In both cases, a match between interference task and memory item should lead to higher interference than the absence thereof.

Participants completed 24 practice trials followed by 320 experimental trials divided into 16 blocks, which took approximately 75 minutes to complete. Practice trials were excluded from any analysis. All manipulations were randomised and counterbalanced.

Experiment 2

The research question of Experiment 2 (N=45) was identical to Experiment 1, but rather than interfering with memory items at the featural level, the manipulations were now space-specific. Unless stated otherwise, the procedure was exactly the same as in Experiment 1. 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 identical to Experiment 1 and the order in which they were presented was randomised and counterbalanced independent from the block cue. 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 cycles per degree (high spatial frequency). The ‘empty’ location would be filled with a noise patch again identical to the ones used during memory item presentation. Mouse buttons were used to respond to the interference task where one mouse button would correspond to high spatial frequency and the other would correspond to low spatial frequency (mouse buttons were counterbalanced between even and odd participants). At the end of the trial, the participant had to respond to both memory items in the order indicated by the block cue.

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

Participants completed 24 practice trials followed by 320 experimental trials divided into 16 blocks, which took approximately 75 minutes to complete. Practice trials were excluded from any analysis. All manipulations were randomised and counterbalanced.

Experiment 3

In experiment 3 (N=51) we wanted to test whether the predictability of the interference feature would influence the magnitude of the interference effect. This experiment had two phases. In the first part, we replicated the exact set-up of experiment 1. Participants did 160 trials in 10 blocks in this condition. Once they finished these, the participants moved onto the second part. In this part, we kept the previous block cue; however, we also blocked the interference condition and explicitly cued its feature dimension as well. Thus, participants had now prior, explicit knowledge of the interference condition. In this condition, participants performed 192 trials across 12 blocks. Participants performed 24 practice trials of part 1 which were not part of any analysis to familiarise themselves with the task. On average it took participants 90 minutes to complete the experiment. All manipulations were randomised and counterbalanced.

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 (Exp 1: N=4, Exp 2: N=2, Exp 3: N=7).

To quantify memory performance of participants we calculated the absolute error response for each trial for each memory item type (active versus passive), 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 were any differences in magnitude between the interference effect on active versus passive 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 passive 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, 2009). We fitted independent models for both active and passive 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). The raw datasets and datasets used to perform the statistical analysis in JASP will be made available on OSF upon publication.

Transparency and openness

In the above, we reported how we determine our sample size, criteria for data exclusions and all experimental manipulations. Experiments were performed in the same order as reported in the manuscript and all data was collected in 2024. Cleaned datasets, modelling outputs, and datasets used for statistical analysis in JASP are available through the link above. The experiments in this paper and their analyses were not preregistered.

Results

Experiment 1

In Experiment 1, we tested whether feature-based interference differed between active and passive items held in working memory. The results are shown in Figure 2. Across interference conditions, we found decisive evidence for better performance on active items (μ_{AS} = 18.6°, sd_{AS} = 6.6°), compared to passive items (μ_{AS} = 22.8°, sd_{AS} = 8.0°) (t_{paired} BF_{10} > 100). Furthermore, we observed decisive evidence that participants performed above chance level for the interference task (μ = 89.4%, sd = 5%) ($t_{one\text{-}sided}$ BF_{10} > 100), indicating that participants paid attention to the stimuli.

With respect to the interference, we performed separate repeated-measure Bayesian ANOVAs for the active and passive probes. For the probe of the active item, we found a main effect of interference (BF_{10} > 100, Figure 2A). We followed this up with one-sided Bayesian t-tests and found decisive evidence for an interference effect for both the match and no match condition

relative to the non-interference baseline (t_{paired} BF_{10} > 100 for both comparisons). We also found evidence that the interference effect was larger in the match compared to the no match condition (t_{paired} BF_{10} > 100). We performed the same analysis for the passive item where we found decisive evidence for a main effect of interference (BF_{10} > 100, Figure 2B). Using one-sided Bayesian t-tests we again found an interference effect for the match and no match condition, compared to baseline (t_{paired} BF_{10} > 100 for both comparisons). Furthermore, we found evidence that interference was stronger in the match compared to the no match condition (t_{paired} BF_{10} > 100). Contrary to our hypothesis, we thus found an interference effect for both the active and the passive item. We therefore wondered if the magnitude of interference differed between the two states. To control for the differences in average error between Probe 1 and Probe 2 we compared the z-scored absolute error between probes and found no difference in the magnitude of interference between states (t_{paired} BF_{10} = 0.158 (BF_{01} = 6.31), Figure 2C).

As a next step, we were interested in what drives the interference effect, and whether the underlying mechanisms are similar for active and passive items. We therefore fitted two separate standard-mixture models, one for the active item and one for the passive item (Zhang & Luck, 2008). As previously we fitted the model separately for the active and passive item and analysed results using repeated-measures Bayesian ANOVAs, followed by one-sided Bayesian t-tests for post-hoc comparisons.

For the active items we found decisive evidence for a main effect of interference on precision (k) (BF_{10} > 100, Figure 2D). Both match and no match conditions showed lower precision compared to the baseline condition (t_{paired} BF_{10} > 100 and t_{paired} BF_{10} = 64.95, respectively). Additionally, the precision was lower in the match compared to the no match condition (t_{paired} BF_{10} = 31.5). Furthermore, we found a decisive effect of interference on the probability of target response ($P_{(m)}$) (BF_{10} > 100, Figure 2F). Both match and no match conditions showed an effect on $P_{(m)}$ relative to the baseline condition (BF_{10} > 100 for both conditions). Additionally, $P_{(m)}$ was lower in the match condition compared to the no match condition (BF_{10} > 100).

For the passive item we found decisive evidence for a main effect of interference on precision (BF_{10} > 100, Figure 2E). Both match and no match condition showed evidence of lower precision compared to baseline (t_{paired} BF_{10} > 100 and t_{paired} BF_{10} = 1.45, respectively). Notably, the evidence for no match compared to baseline was only anecdotal. In line with this, there was decisive evidence for precision being lower in the match compared to the no match condition (t_{paired} BF_{10} > 100). We also found a main effect of interference on the probability of a target response (BF_{10} > 100, Figure 2G). The probability of a target response was lower for both the match and the no match condition compared to baseline (t_{paired} BF_{10} > 100 for both conditions). The probability of a target response was lower in the match compared to the no match condition (BF_{10} = 11.62).

Thus, we found that precision, that is, how well an item is represented in memory, and the probability of a target response were equally affected by the interference manipulation for items that are active or passive within working memory.

Figure 2

Results for experiment 1

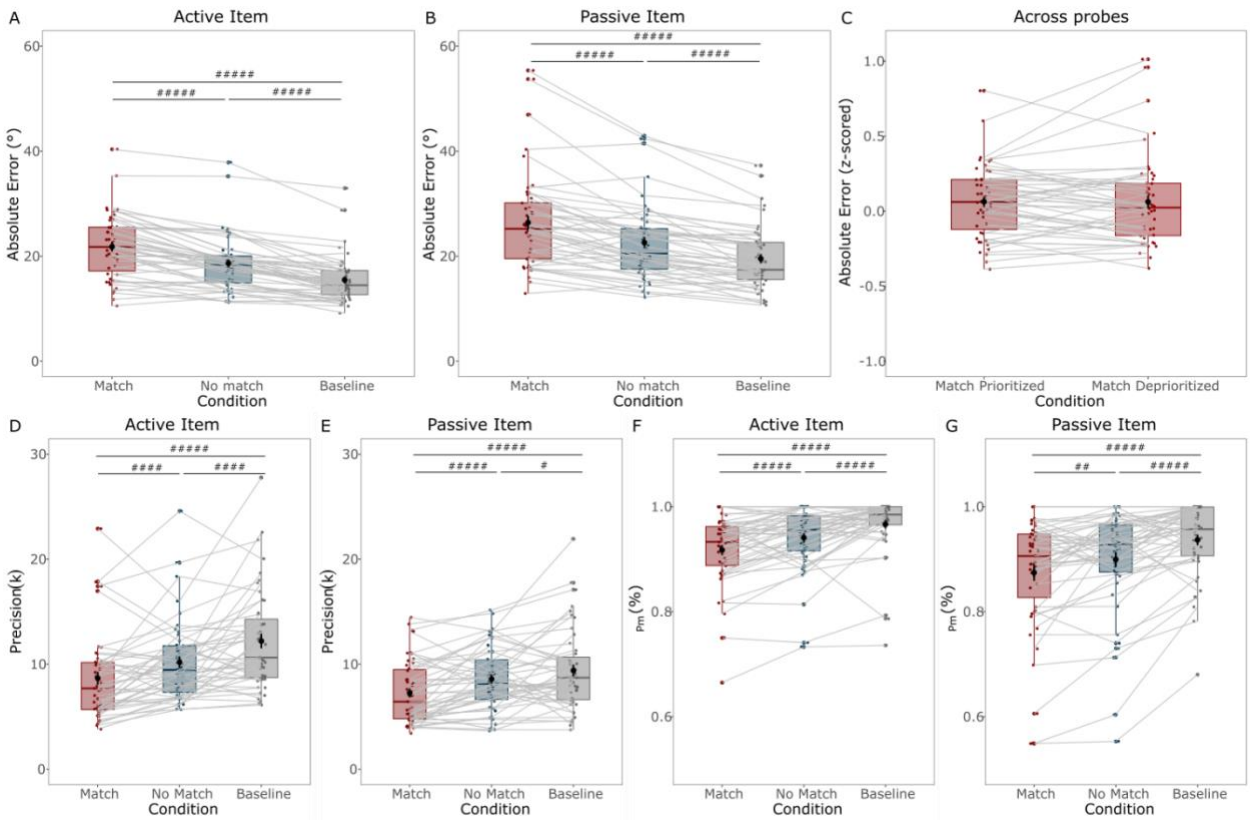

Note: A) Absolute error for the active item for the interference versus no interference condition. (B) Absolute error for the passive item for the match versus no match condition. (C) Z-scored absolute error for the match condition across both probes. (D) Precision (*k*) for the active item for match versus no match condition. (E) Precision (*k*) for the passive item for match versus no match condition. (F) Probability of a target response ($P_{(m)}$) for the active item for match versus no match condition. (G) Probability of a target response ($P_{(m)}$) for the passive 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 active and passive items in working memory; however, this time we interfered in a spatially-specific, rather than feature-specific manner. The results are shown in Figure 3. In accordance with the previous experiment, we found decisive evidence for better performance on the active item (μ_{AS} = 18.9°,

sd_{AS} = 5.4°) compared to the passive item (μ_{AS} = 20.9°, sd_{AS} = 4.7°) (t_{paired} BF_{10} > 100). We also observed decisive evidence that participants performed above chance in the interference task (μ = 95.9%, sd = 7.4%) ($t_{one\text{-}sided}$ BF_{10} > 100).

We repeated the same analyses as for experiment 1. We found a significant main effect of interference on the active item (BF_{10} > 100, Figure 3A). We found an effect of interference for match and no match condition relative to the baseline (BF_{10} > 100 for both comparisons). However, we found strong evidence against a differential effect between the match and the no match condition (BF_{10} = 0.086 (BF_{01} = 11.64)). Thus, while there was an interference effect relative to the baseline condition; whether the item was presented in the cued location or not did not modulate the interference effect. For the passive item, we also found a main effect for interference (BF_{10} > 100, Figure 3B). As for the active item, there was again an effect of interference for both the match and no match condition relative to the baseline condition (BF_{10} > 100 for both comparisons). However, we again did not find a modulation of the interference effect between the match compared to no match condition (BF_{10} = 0.189; BF_{01} = 5.3). Overall, these analyses suggest that while the interference manipulation worked, because absolute error was higher across both interference conditions when compared to baseline, the interference affected both items to similar extents, independent of spatial location. Comparing the active match condition with the passive match condition further showed that there was no different in the amount of interference between them (BF_{10} = 0.137 (BF_{01} = 7.3, Figure 3C).

As a next step, we looked at the model parameters of the standard-mixture model. We found no main effect on precision for the active item (BF_{10} = 0.6 (BF_{01} = 1.658), Figure 3D). We did find a main effect on the probability of a target response for the active item (BF_{10} = 79.42, Figure 3F). The probability of a target response was lower for the match and no match condition compared to baseline (BF_{10} = 13.97 and BF_{10} > 100, respectively). There was no effect on the probability of a target response between the match and no match condition (BF_{10} = 0.11 (BF_{01} = 9.273)). For the passive item there was no effect on precision either (BF_{10} = 0.21 (BF_{01} = 4.82),

Figure 3E). There was a main effect on the probability of a target response for the passive item (BF_{10} = 6.9, Figure 3G). We found a lower probability of a target response for the match and no match condition compared to the baseline (BF_{10} = 1.41 and BF_{10} > 100, respectively), and no evidence for a difference between match and no match condition BF_{10} = 0.079 (BF_{01} = 12.7)). Overall, we did not find any evidence in this study for location-specific modulation of interference, and neither did we find evidence for a modulation of interference between the functional states of the items in working memory. The general interference effect between both interference conditions compared to the baseline that we did find was driven by a decrease in the probability of a target response, which means there was increased guessing of the response by participants.

Figure 3

Results for experiment 2

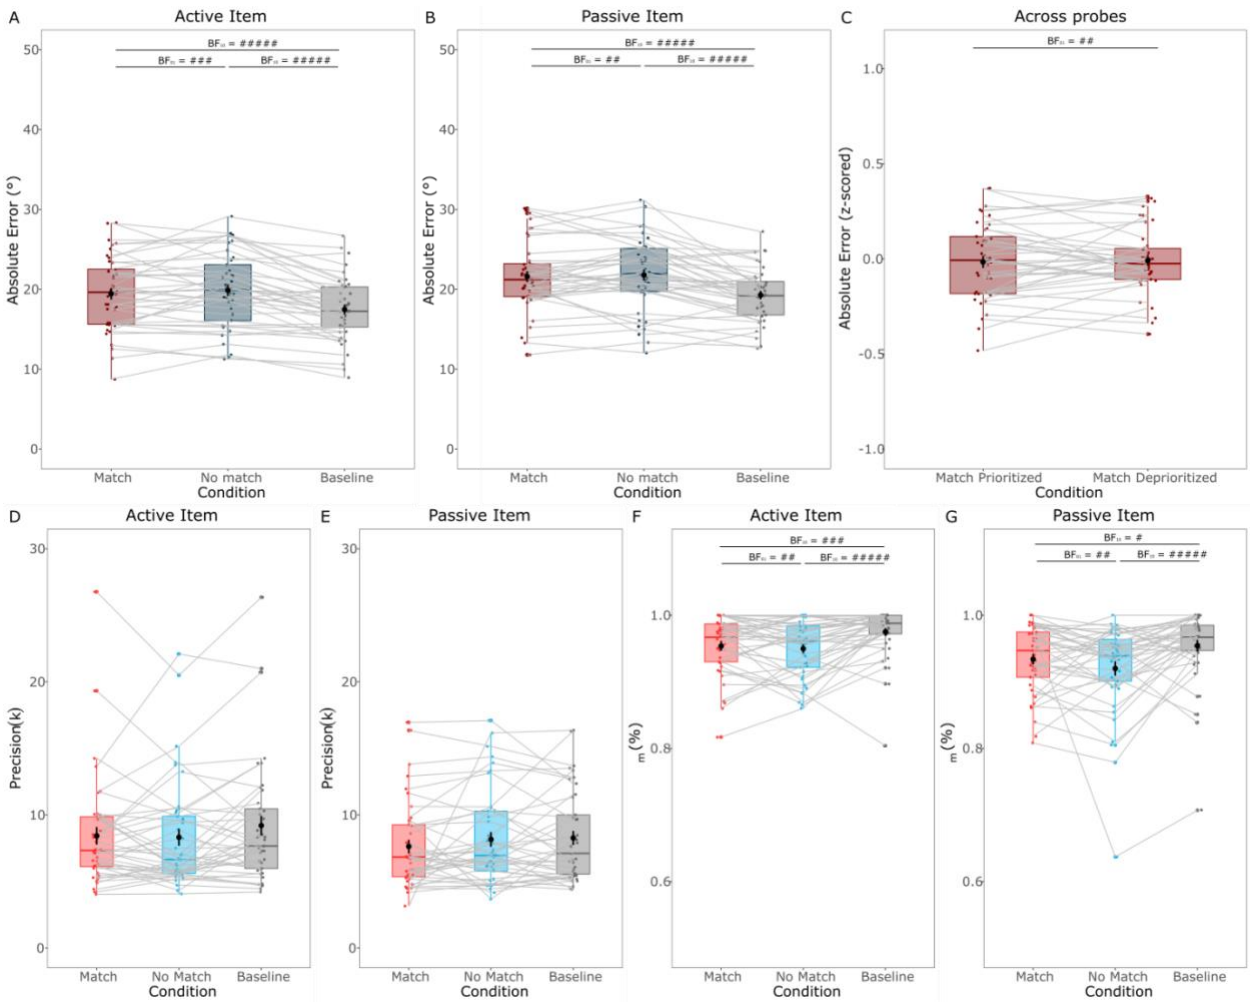

Note: (A) Absolute error for the active item for the interference versus no interference condition. (B) Absolute error for the passive item for the match versus no match condition. (C) Z-scored absolute error for the match condition across both probes. (D) Precision (*k*) for the active item for match versus no match condition. (E) Precision (*k*) for the passive item for match versus no match condition. (F) Probability of a target response ($P_{(m)}$) for the active item for match versus no match condition (G) Probability of a target response ($P_{(m)}$) for the passive 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 wanted to test if expectations about the upcoming interference could modulate the effect. Since we had a stronger interference effect in experiment 1 compared to experiment 2, we went back to feature-based interference. This time we did not only block the item priority, but we also included a blocked and explicit cue for the interference feature. Thus, participants could anticipate which item was going to be interfered with and we kept this constant across each block of trials. The idea was that this would provide an opportunity to the participants to exert strategic control over memory maintenance, if they could, possibly depending on the functional state of the items being maintained.

We found higher accuracy for the active item (μ_{AS} = 16.9°, sd_{AS} = 5.5°) compared to the passive item (μ_{AS} = 19.2°, sd_{AS} = 5.3°) (t_{paired} BF_{10} > 100). As in previous experiments, participants accurately responded in the interference condition, indicating that they paid attention to the task (μ = 89.9%, sd = 7%) ($t_{one\text{-}sided}$ BF_{10} > 100).

Again, we performed the same analysis as in the previous experiment, but we separated the blocks into random interference blocks (identical to experiment 1) and fixed interference blocks, where the interference condition stayed the same throughout the experiment and was cued to the participant.

Random condition. We found a main effect of interference on the active item (BF_{10} > 100). We found an effect of interference on both the match and no match condition compared to baseline (BF_{10} > 100 for both comparisons). Furthermore, we found an effect of interference between the match and no match condition (BF_{10} > 100). We found the same main effect of interference for the passive item (BF_{10} > 100). Interference was larger for both the match and no match condition relative to the baseline condition (BF_{10} > 100 and BF_{10} = 66.75, respectively). Additionally, interference was larger for the match compared to the no match condition (BF_{10} = 16.3). As before, the interference effect was not modulated by the functional state of the items in working memory (BF_{10} = 0.22 (BF_{01} = 4.56)).

Fitting the standard-mixture model to this data, we found a main effect of interference on precision for the active item ($BF_{10} > 100$). Precision was lower in the match and no match condition relative to baseline ($BF_{10} > 100$ and $BF_{10} = 2.67$, respectively). Furthermore, precision was lower in the match compared to the no match condition ($BF_{10} = 2.91$). There was also a main effect of interference on the probability of a target response for the active item ($BF_{10} = 10.12$). We found evidence that both the match and no match condition had a lower probability of a target response compared to the baseline condition ($BF_{10} = 23.31$ and $BF_{10} = 3.26$, respectively). Unlike in experiment 1, we did not find a difference between the match and no match condition ($BF_{10} = 0.757$ ($BF_{01} = 1.32$)). For the passive item, we found a main effect of interference on precision ($BF_{10} = 2.84$). We found evidence that the match condition had lower precision compared to the baseline ($BF_{10} = 8.28$); however, we found anecdotal evidence that there was no difference between the no match condition and the baseline ($BF_{10} = 0.75$ ($BF_{01} = 1.33$)). We did find evidence for a difference between the match and no match condition on precision ($BF_{10} = 1.42$). There was also evidence for a main effect on the probability of a target response ($BF_{10} = 5.26$). We found evidence for reduced probability of a target response in the match and no match condition relative to baseline ($BF_{10} = 23.32$ and $BF_{10} = 7.05$, respectively). However, we found no evidence for a difference between the match and no match condition on the probability of a target response ($BF_{10} = 0.173$ ($BF_{01} = 5.79$)).

Overall, we observed some discrepancies in the model parameters when compared to experiment 1. We found anecdotal evidence for the null hypothesis in two cases where we previously found strong evidence for the alternative hypothesis in experiment 1. Overall, these findings could be caused by the reduced trial numbers of this part of the experiment compared to experiment 1, which was due to the second condition that we included here – especially since the evidence for the null hypothesis was only anecdotal. Notably we did find substantial evidence for no difference between the match and no match condition on the probability of a target response for the passive item, which was different from the results found in experiment 1.

Fixed condition. There was a main effect of interference for the active item (BF_{10} > 100). We found a larger error in the match and no match interference condition compared to baseline (BF_{10} > 100 for both comparisons), and a larger error in the match compared to the no match condition (BF_{10} > 100). For the passive item we found the same main effect of interference (BF_{10} > 100). Again, the error was larger in the match and no match condition compared to baseline (BF_{10} > 100 for both comparisons), and larger in the match compared to the no match condition (BF_{10} = 16.68). The interference effect was not modulated by the priority state in working memory (BF_{10} = 0.93 (BF_{01} = 1.2))

We found a main effect of interference on precision for the active item (BF_{10} > 100). Precision was lower for both match and no match interference conditions compared to baseline (BF_{10} > 100 for both comparisons). Precision was lower for the match compared to the no match condition (BF_{10} > 100). Interference also showed a main effect on the probability of a target response (BF_{10} > 100). We found lower probability of a target response in both the match and no match condition compared to the baseline condition (BF_{10} > 100 and BF_{10} = 21.13, respectively). We also found evidence for lower probability of a target response in the match compared to the no match condition (BF_{10} = 3.89). For the passive item, we found a main effect of interference on precision (BF_{10} > 100). We found evidence of lower precision for the match and no match condition relative to baseline (BF_{10} > 100 and BF_{10} = 11.75, respectively). We did not find evidence for a difference between the match and no match condition on precision (BF_{10} = 0.514 (BF_{01} = 1.95). We also found evidence of a main effect on the probability of a target response (BF_{10} = 8.6). We found evidence that the probability of a target response was lower in the match compared to the baseline condition (BF_{10} = 28.77), as well as anecdotal evidence for a difference between the match and no match condition (BF_{10} = 1.55). We found no evidence for a difference between the no match and baseline condition (BF_{10} = 0.92 (BF_{01} = 1.09)).

Overall, the patterns we observed were very similar between the fixed and random condition trials, indicating that participants did not benefit from having explicit knowledge of the

interference condition, and that active and passive items were still equally affected by interference feature-based interference. Following this observation, we finally tested if participants had any benefit from knowledge of the interference feature. To this end, we calculated the difference between interference (match and no match condition combined) minus the baseline condition to test if the magnitude of interference was modulated by the fixed versus random condition. We then performed a paired-sample t-test between these difference scores for the active and passive item separately. We found no evidence that the magnitude of interference differed between random and fixed blocks for the active item (BF_{10} = 0.36 (BF_{01} = 2.75)) or passive item (BF_{10} = 0.17 (BF_{01} = 5.87)).

Figure 4

Results of experiment 3

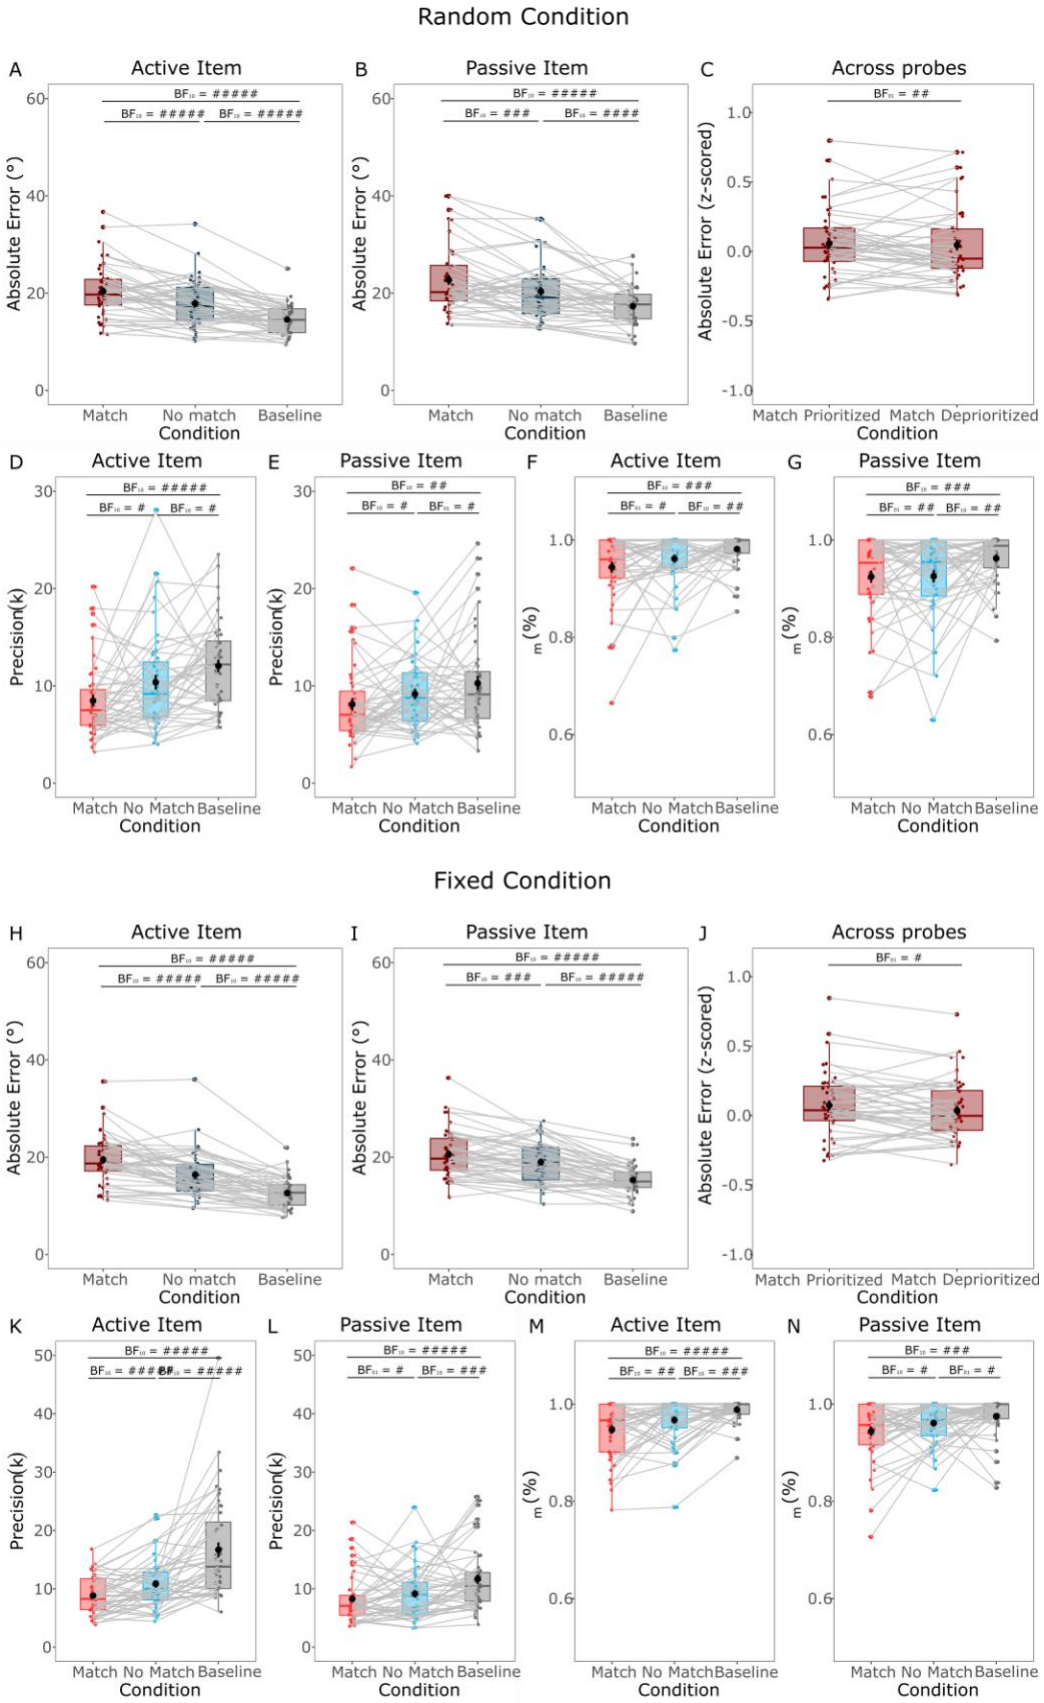
Random Condition
A
Active Item
B
Passive Item
C
Across probes
Absolute Error (°)
Absolute Error (z-scored)
Match
No match
Baseline
Match Prioritized
Match Deprioritized
Condition
D
Active Item
E
Passive Item
F
Active Item
G
Passive Item
Precision(k)
Match No Match Baseline
Fixed Condition
H
Active Item
I
Passive Item
J
Across probes
K
Active Item
L
Passive Item
M
Active Item
N
Passive Item

Note: Plots **A-G** are for the random condition. (A) Absolute error for the active item for the interference versus no interference condition. (B) Absolute error for the passive 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 active item for match versus no match condition. (E) Probability of a target response ($P_{(m)}$) for the active item for match versus no match condition. (F) Precision (*k*) for the passive item for match versus no match condition. (G) Probability of a target response ($P_{(m)}$) for the passive item for match versus no match condition. Plots **H-N** are for the fixed condition. (H) Absolute error for the active item for the interference versus no interference condition. (I) Absolute error for the passive item for the match versus no match condition. (J) Z-scored absolute error for the match condition across both probes. (K) Precision (*k*) for the active item for match versus no match condition. (L) Precision (*k*) for the passive item for match versus no match condition. (M) Probability of a target response ($P_{(m)}$) for the active item for match versus no match condition. (N) Probability of a target response ($P_{(m)}$) for the passive 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 persistent-activity or activity-quiescent states, and an accumulating amount of evidence has shown that these two mechanisms of WM storage are not mutually exclusive (Lundqvist, Herman, & Miller, 2018; Muhle-Karbe et al., 2021; Oberauer & Awh, 2022). However, little research to date has focused on the potential functional roles of these differing states in WM. In the current study, we set out to test one potential functional role of these varying states – namely, the hypothesis that functionally passive (and neuronally activity-quiescent) states could be protecting items from interference, compared to functionally active, (and neuronally persistent-activity) states.

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 the maintenance of active (item cued to be probed first and therefore thought to be prioritised) and passive (item cued to be probed second and therefore thought to be deprioritised) items, as they might be

supported by different neural mechanisms. 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 hypothesised that activity-quiescent maintenance through functional connectivity might afford better protection from interference (Kozachkov et al., 2022; Lundqvist, Herman, Warden, et al., 2018; Miller et al., 2018).

Contrary to our hypothesis, we found no evidence that passive items (assumed to be maintained in a quiescent state) were more protected from interference compared to active items (assumed to be in a persistent-activity state). Instead, both states showed similar susceptibility to interference across all three experiments, regardless of whether the distractor overlapped in feature (Experiment 1) or spatial dimension (Experiment 2), a pattern that persisted even when the participants had explicit knowledge about the nature of the upcoming interference (Experiment 3).

These findings are in line with recent results from Zhang and Lewis-Peacock (2023), who similarly found no difference in interference susceptibility between active and passive items – using face stimuli varying along sex and age dimensions. Earlier studies have suggested that prioritised items might *heighten* susceptibility to certain forms of disruption (Hitch et al., 2018; Hu et al., 2014; Mallett & Lewis-Peacock, 2019). Our data do not support this idea. Likewise, the present evidence also challenges earlier proposals that prioritised items always enjoy protection from interference (Makovski et al., 2008; Makovski & Pertzov, 2015; Schneider et al., 2017; Souza & Oberauer, 2016).

What could be causing such seemingly contradictory outcomes regarding the nature of interference effects in previous studies? One possibility involves the specifics of the task – for instance, the exact nature of the distractor (Hu et al., 2014; Pinto et al., 2013), the number of to-be-remembered items (Allen & Ueno, 2018), or the time points of distractor presentation. All of

these could determine whether priority confers resilience or risk. Additionally, how “priority” is operationalised – through retro-cued, reward or attentional modulations – may lead to divergent results. Here, we used explicit blocks to designate which item was relevant first versus second, offering a relatively direct manipulation of priority. This approach aligns with Muhle-Karbe and collagues (2021), who showed that prioritised items display heightened neuronal representations while deprioritised ones can exist in a silent trace. Nevertheless, in our case, the state the items were in did lead to a differential effect of interference.

In addition to the question of whether deprioritised items are more protected from interference, our data also allow for a test of whether interference might be different between spatial and featural overlap. Experiment 1 (feature-based interference) showed a robust match effect on accuracy and precision: Both active and passive items showed higher disruption when the distractor shared the same colour or orientation. Interestingly, in Experiment 2 (space-based interference), the match versus no-match distinction was less critical. The presence of a spatial distractor affected WM performance overall, when compared to the no distractor baseline condition, but matching locations did not systematically enhance interference. This outcome aligns with domain-specific interference theories, whereby spatial overlap may be less disruptive if the stored features (colour/orientation) do not critically rely on spatial codes (Awh & Jonides, 2001; Jerde et al., 2012).

Regardless of the specific dimension, the critical finding remains that the magnitude of interference did not systematically differ between active and passive states. This was further extended in Experiment 3, where participants had explicit knowledge about which feature would be targeted by the interference task. Even when provided with this opportunity to anticipate the nature of the interference task, our data suggests that participants were not able to make use of this strategic cue to protect the working memory content in either state from interference. Thus, while interference can take place in multiple forms – feature or space-based – and can be

anticipated, none of these manipulations selectively provided increased resilience of passive items.

If these presumably quiescent states of passive items in WM do not afford robust protection from external distractors, why might the brain maintain items in silent traces? One possibility could be that quiescent states help minimise internal interference among multiple items stored in WM (Chatham & Badre, 2015; Rose et al., 2016). By storing passive items – which are only relevant later in time – in a synaptic format, the system might free up active firing for the currently most relevant item, thereby reducing crosstalk (Muhle-Karbe et al., 2021). Indeed, it is a well-accepted fact that WM capacity is strongly limited, which would suggest that it is important to dynamically balance activity across items, and having only one, or few items fully “activated” at a time. This mechanism may be especially crucial under higher loads than tested here (only two items). Future work could explore whether the functional benefit of quiescent states is relevant for preventing overlap within WM itself, rather than insulating those items against external interference.

The current study comes with some 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 based on previous papers (e.g., LaRocque et al., 2012; Muhle-Karbe et al., 2021; Wolff et al., 2017). Second, we only tested memory for a pair of items, which is a relatively modest load on working memory, and which might not have required offloading of information in a quiescent state. A higher WM load might reveal different patterns: If load increases, there might be more need for items to be offloaded in a quiescent state to avoid cross-item interference. Finally, although we manipulated whether participants were cued to the nature of the interference task (Experiment 3), the precise temporal dynamics of the interference may play a major role in whether silent states become reactivated or remain quiescent, which we did not modulate. Future paradigms might systematically vary timing, load, and distractor complexity to better clarify when—if ever—silent states reduce interference from external inputs.

To conclude, across three experiments involving feature-based and space-based interference, we did not see protection from interference on passive items in WM. Even explicit knowledge of which feature dimension would be targeted by the interference task did not protect passive (nor active) items from interference. Rather than serving as an external interference shield, quiescent states may be functionally critical for mitigating internal interference among concurrently stored items. Future work that combines higher WM loads, neural measures, and different interference paradigms could further elucidate the possible functional advantages of short-term synaptic plasticity in WM, and clarify how the brain flexibly transitions items between active and quiescent states to meet external demands.

Author contributions

Sophia Wilhelm contributed to conceptualisation of the experiments, and served as lead for data curation, formal analysis, methodology, visualisation and writing – original draft. Yuanyuan Weng contributed to data curation and reviewed the final manuscript. Jelmer Borst and Robbert Havekes reviewed the final manuscript. Elkan Akyürek served as lead to conceptualisation of the experiments, project administration and supervision and contributed to writing – original draft.

References

Allen, R. J., Castellà, J., Ueno, T., Hitch, G. J., & Baddeley, A. D. (2015). What does visual suffix interference tell us about spatial location in working memory? *Memory & Cognition*, *43*(1), 133–142. https://doi.org/10.3758/s13421-014-0448-4

Allen, R. J., & Ueno, T. (2018). Multiple high-reward items can be prioritized in working memory but with greater vulnerability to interference. *Attention, Perception, & Psychophysics*, *80*(7), 1731–1743. https://doi.org/10.3758/s13414-018-1543-6

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*(3), 119–126. https://doi.org/10.1016/S1364-6613(00)01593-X

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

Banks, J. B., Mallick, A., Nieto, A. C., Zanesco, A. P., & Jha, A. P. (2023). The role of affective interference and mnemonic load in the dynamic adjustment in working memory. *Memory & Cognition*, *51*(1), 203–220. https://doi.org/10.3758/s13421-022-01303-4

Chatham, C. H., & Badre, D. (2015). Multiple gates on working memory. *Current Opinion in Behavioral Sciences*, *1*, 23–31. https://doi.org/10.1016/j.cobeha.2014.08.001

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

Gresch, D., Boettcher, S. E. P., van Ede, F., & Nobre, A. C. (2021). Shielding working-memory representations from temporally predictable external interference. *Cognition*, *217*, 104915. https://doi.org/10.1016/j.cognition.2021.104915

Hitch, G. J., Hu, Y., Allen, R. J., & Baddeley, A. D. (2018). Competition for the focus of attention in visual working memory: Perceptual recency versus executive control. *Annals of the New York Academy of Sciences*, *1424*(1), 64–75. https://doi.org/10.1111/nyas.13631

Hu, Y., Hitch, G. J., Baddeley, A. D., Zhang, M., & Allen, R. J. (2014). Executive and perceptual attention play different roles in visual working memory: Evidence from suffix and strategy effects. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1665–1678. https://doi.org/10.1037/a0037163

Jerde, T. A., Merriam, E. P., Riggall, A. C., Hedges, J. H., & Curtis, C. E. (2012). Prioritized Maps of Space in Human Frontoparietal Cortex. *The Journal of Neuroscience*, *32*(48), 17382–17390. https://doi.org/10.1523/JNEUROSCI.3810-12.2012

Klink, P. C., Jeurissen, D., Theeuwes, J., Denys, D., & Roelfsema, P. R. (2017). Working memory accuracy for multiple targets is driven by reward expectation and stimulus contrast with different time-courses. *Scientific Reports*, *7*(1), 9082. https://doi.org/10.1038/s41598-017-08608-4

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., & Miller, E. K. (2018). Working Memory: Delay Activity, Yes! Persistent Activity? Maybe Not. *Journal of Neuroscience*, *38*(32), 7013–7019. https://doi.org/10.1523/JNEUROSCI.2485-17.2018

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

Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *Quarterly Journal of Experimental Psychology*, *68*(9), 1735–1743. https://doi.org/10.1080/17470218.2015.1049623

Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 369–380. https://doi.org/10.1037/0278-7393.34.2.369

Mallett, R., & Lewis-Peacock, J. A. (2019). Working memory prioritization impacts neural recovery from distraction. *Cortex*, *121*, 225–238.https://doi.org/10.1016/j.cortex.2019.08.019

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. doi: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

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

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

Pinto, Y., Sligte, I. G., Shapiro, K. L., & Lamme, V. A. F. (2013). Fragile visual short-term memory is an object-based and location-specific store. *Psychonomic Bulletin & Review*, *20*(4), 732–739. https://doi.org/10.3758/s13423-013-0393-4

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

Schneider, D., Barth, A., & Wascher, E. (2017). On the contribution of motor planning to the retroactive cuing benefit in working memory: Evidence by mu and beta oscillatory activity in the EEG. *NeuroImage*, *162*, 73–85. https://doi.org/10.1016/j.neuroimage.2017.08.057

Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, *78*(7), 1839–1860. https://doi.org/10.3758/s13414-016-1108-5

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. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*(4), 423–428.

Zhang, Z., & Lewis-Peacock, J. A. (2023). Prioritization sharpens working memories but does not protect them from distraction. *Journal of Experimental Psychology: General*, *152*(4), 1158–1174. https://doi.org/10.1037/xge0001309