

Jointly looking to the past and the future in visual working memory

Baiwei Liu , Zampeta-Sofia Alexopoulou, Freek van Ede 

Institute for Brain and Behavior Amsterdam, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, The Netherlands

Reviewed Preprint

Revised by authors after peer review.

[About eLife's process](#)

Reviewed preprint version 2

April 16, 2024 (this version)

Reviewed preprint version 1

November 28, 2023

Posted to preprint server

August 2, 2023

Sent for peer review

July 29, 2023

 https://en.wikipedia.org/wiki/Open_access

 Copyright information

Abstract

Working memory enables us to bridge past sensory information to upcoming future behaviour. Accordingly, by its very nature, working memory is concerned with two components: the past and the future. Yet, in conventional laboratory tasks, these two components are often conflated, such as when sensory information in working memory is encoded and tested at the same location. We developed a task in which we dissociated the past (encoded location) and future (to-be-tested location) attributes of visual contents in working memory. This enabled us to independently track the utilisation of past and future memory attributes through gaze, as observed during mnemonic selection. Our results reveal the joint consideration of past and future locations. This was prevalent even at the single-trial level of individual saccades that were jointly biased to the past and future. This uncovers the rich nature of working memory representations, whereby both past and future memory attributes are retained and can be accessed together when memory contents become relevant for behaviour.

eLife assessment

This **important** study advances our understanding of how past and future information is jointly considered in visual working memory by studying gaze biases in a memory task that dissociates the locations during encoding and memory tests. The evidence supporting the conclusions is **convincing**, with state-of-the-art gaze analyses that build on a recent series of experiments introduced by the authors. This work will be of broad interest to vision scientists interested in the interplay of vision, eye movements, and memory.

Introduction

Working memory is a fundamental cognitive function that enables us to hold onto past sensory information in service of upcoming future behaviour ¹⁻³. Accordingly, by its very nature, working memory is concerned with two components: the past and the future.

In conventional laboratory tasks, past and future components are often conflated such as when sensory information in working memory is tested at the same location as where it was encoded. By contrast, in the dynamic situations we face every day, sensory information often disappears at one specific location (where it enters visual working memory) but becomes relevant at another location (as also in [4–8](#)). Imagine trying to capture a photograph of a precious bird species that you just saw disappear behind a building. Your working memory of the bird is likely to consider not only where you last saw the bird, but also where you expect it to re-appear to capture it on camera. Such situations raise an interesting, underexplored question: when past and future locations of memory contents are not the same, does the brain code internal representations with regard to past, future, or both?

To address this question, we developed a task in which we dissociated the encoding (past) and to-be-tested (future) locations associated with visual representations in working memory. This enabled us to experimentally isolate past and future memory attributes and to track their respective utilisation through spatial biases in gaze behaviour in healthy human volunteers.

Results

Twenty-five human volunteers performed a working-memory task in which visual memory items were encoded and tested at different locations (**Fig. 1a**). Participants memorised two coloured gratings with different orientations presented either vertically or horizontally. The crucial manipulation was that we always tested memory content in the orthogonal axis and at a predictable location (depending on the future rule that we varied across sessions; see **Supplementary Fig. 1** for the four possible rules).

After a delay period, we cued the relevant memory item by changing the colour of the central fixation dot. At this stage, participants were required to select the colour-matching grating from working memory in order to compare it to the upcoming test stimulus (clockwise/counter-clockwise judgement). Crucially, we always presented two stimuli at the test-phase of which only one was relevant, as determined by the future rule. For example, in **Figure 1a**, after the green memory item is cued, the relevant test stimulus will be the right stimulus, given the future rule (top item tested on right) in this session. Note how the colour cue only ever informed the to-be-tested memory content but never directly informed the to-be-tested future location. The to-be-tested location was an attribute of the cued memory content but not of the cue itself.

Participants were able to perform this dynamic visual working-memory task, with an average accuracy of 70 ± 2 percent correct (mean $\pm$ SEM) and an average reaction time of 1218 ± 125 ms.

Gaze reveals the use of both past and future memory attributes that are considered at overlapping time windows

To track the utilisation of past and/or future locations associated with working memory contents, we tracked spatial biases in gaze following the cue to select either memory item. Specifically, we focused on directional biases in saccades, that have previously been shown to be sensitive to selective spatial attention [9–16](#), even when directed internally [17–20](#).

As shown in **Figure 1b**, after cue onset, saccades became biased in the direction of the encoded (past) location of the selected memorized target, as demonstrated by significantly more gaze shifts toward vs. away from the encoded location of the target (**Fig. 1b**; cluster $P < 0.001$), starting from around 200 ms after the cue. This is consistent with our prior demonstrations of directional eye-movement biases within the spatial lay-out of working memory [17–20](#). In our current task, this uniquely reveals how the retention of items in working memory continues to rely on their past encoded locations, even if items are known to become relevant (tested) at another location.




Figure 1.

Directional gaze biases by past and future locations during mnemonic selection overlap in time.

a) Task schematic. Participants memorised two oriented gratings with different colours presented either vertically or horizontally. Following a delay, a colour change of the central fixation dot prompted participants to select the colour-matching item from working memory to report its orientation later. After another delay, two test gratings appeared transiently and participants compared the cued memory grating to the relevant test grating (clockwise/counter-clockwise judgment) that was determined by the 'future rule' that was stable within each block. After a response, feedback (0: wrong, 1: correct) was presented at the side of the relevant test grating. Dash lines serve to explain the association between the encoding and test locations and were never presented in the actual experiment. **b-c)** Time courses of gaze shift rates (number of saccades per second) for shifts toward and away from the encoded (panel b) and to-be-tested (panel c) locations. **d)** Overlays and comparisons of the difference in gaze-shift rates (toward minus away) for the past (encoded) location and the future (to-be-tested) location. Horizontal lines indicate significant temporal clusters (cluster-based permutation test, $P < 0.001$). Data are presented as mean values with shading reflecting 95% confidence intervals, calculated across participants ($n = 25$).

Having established that past (encoded) memory locations were still utilised by participants in our task (despite us never asking about memory-item locations), an interesting question becomes when the future memory attribute (to-be-tested location) would be considered after the selection cue. Intuitively, participants may select the relevant item at its past location, before considering the relevant future test location – which would yield a serial pattern of past-before-future. In contrast, as shown in **Figure 1c** and **d**, we found a similar saccade bias to the relevant future location (**Fig. 1c**; cluster $P < 0.001$). Strikingly, this future bias also emerged early after the cue. An overlay of the spatial biases (toward vs. away) in the orthogonally manipulated past and future axes (**Fig. 1d**), revealed consideration of past and future locations at overlapping time windows. Gaze biases in both axes were driven predominantly by microsaccades (**Supplementary Fig. 2**) and occurred similarly in horizontal-to-vertical and vertical-to-horizontal trials (**Supplementary Fig. 3**). Moreover, while the past bias was relatively transient, the future bias continued to increase in anticipation of the test stimulus and increasingly incorporated eye-movements beyond the microsaccade range (see **Supplementary Fig. 4** for a more extended time range).

These data thus suggest the joint consideration – or “activation” – of past and future memory attributes, at least when analysing past and future memory attributes separately. Below we provide additional single-trial (single-saccade) evidence for this interpretation.

Individual saccades reveal truly joint consideration of past and future memory attributes

In principle, the observed joint activation of the past and future locations associated with the cued memory content in the trial-averaged and participant-averaged data could result from two alternative scenarios with different interpretations. First, either the past or the future alone may be considered in different trials and/or participants, without past and future memory attributes ever being considered together. Alternatively, participants may truly consider both past and future memory attributes jointly at the single-trial, single-saccade level.

While it is notoriously hard to disentangle the single-trial interpretation of averaged data²¹, we were here able to do so by interrogating the individual saccade characteristics for which the two alternative scenarios make different predictions. In the first scenario, saccades should be biased to either the past or the future location, but there should be no dependency between them (i.e. a saccade may be biased to the past location regardless of the future location, and vice versa). In contrast, in the second scenario, with truly joint consideration of the past and future, there should be a clear dependency: it should be those saccades that are biased to the past that are also biased to the future. In other words, the future-biased saccades should predominantly be driven by the past-biased saccades, and vice versa.

To disentangle these alternatives at the single-trial level of individual saccades, we focused on the first saccades after the cue, in the 200-600 ms window. We previously identified this window as the relevant window for microsaccade biases by internal selective attention¹⁷, and this is also where we observed the overlapping past and future biases in the average here (see **Figure 1d**). To facilitate visualisation and quantification, we rotated all detected saccades to match a consistent coordinate frame, in which the past location is represented horizontally (right=towards, left=away), and the future location vertically (top=toward, bottom=away; see **Fig. 2a**).

As shown in **Figure 2a**, the majority of first-detected microsaccades in our window of interest were made toward the past (right > left) and toward the future (top > bottom), replicating our prior analyses. Critically, this visualisation and quantification enabled us to disentangle the two alternative sketched above in which past and future saccades at the single-trial (individual-saccade) level were either independent (past bias regardless of future and future bias regardless of past) or dependent (joint past and future bias). Our data supported the latter.




Figure 2.

Individual saccades are jointly biased to past and future memory attributes.

a) The distribution of the direction of the first saccades we detected after cue onset relative to past (horizontal) and future (vertical) locations. Data from the different sessions were rotated to match a common coordinate frame. **b)** The bias toward the past (x-axis) as a function of saccade direction with regard to the future (y-axis). **c)** The bias toward the future (y-axis) as a function of saccade direction with regard to the past (x-axis). In **b-c**, The bold black line indicates the significant temporal cluster (cluster-based permutation test, $P < 0.001$). Data are presented as mean values with shading indicating 95% confidence intervals, calculated across participants ($n = 25$). **d)** The percentage of identified first saccades toward or away from the future location as a function of whether the same saccades were also biased toward or away from the past location. Error bars in panel d indicate ± 1 SEM calculated across participants ($n = 25$).

Figure 2b shows the bias toward the past as a function of saccade direction with regard to the future. Likewise, **Figure 2c**, shows the future bias as a function of whether saccades were also biased to the past. In both cases, we see a clear dependency: the past bias is particularly pronounced for saccades that also have a future bias (**Fig. 2b**) and, vice versa, the future bias is most pronounced for saccades that also have a past bias (**Fig 2c**). This is perhaps best appreciated by the binarised quantification of these same data in **Figure 2d**: showing the percentage of identified first saccades toward or away from the future location, as a function of whether the same saccades were also biased toward or away from the past location. We found a clear interaction ($F(1, 24) = 18.1$, $P < 0.001$, partial $\eta^2 = 0.43$), whereby the bias toward the future location was exclusively observed for those saccades that were also biased toward the past location. Indeed, follow-up t-tests revealed no difference future bias for saccades that were away from the past ($t(24) = 1.13$, $P_{\text{Bonferroni}} = 1$, $d = 0.23$), but a clear future effect for saccades that were toward the past ($t(24) = 4.65$, $P_{\text{Bonferroni}} < 0.001$, $d = 0.93$). This would not be expected if single saccades cared exclusively about either the past or the future location. Instead, this provides single-trial-level support with the truly joint consideration of past and future memory attributes.

Discussion

Here, we brought the study of visual working memory into a dynamic context by experimentally dissociating (orthogonalising) the past (last-seen location) and future (to-be-tested location) attributes of visual memory contents, and independently tracking the utilisation of these two attributes through gaze. Doing so, we unveil a novel, fundamental property of working memory – the joint availability and utilisation (i.e. selection) of past and future memory attributes. As such, our data provide key support for the proposal that memory is fundamentally future-oriented [2,3,22,23](#), while also reminding us that past memory attributes are not forgotten, even when these do not become relevant again.

Our finding of joint utilisation of past and future memory attributes emerged from at least two alternative scenarios of how the brain may deal with dynamic everyday working memory demands in which memory content is encoded at one location but needed at another. First, memory contents could have directly been remapped (cf. [4,24–26](#)) to their future-relevant location. However, in this case, one may have expected to exclusively find a future-directed gaze bias, unlike what we observed. Moreover, using a complementary analysis of saccade directions along the axis of the future rule (cf. [24](#)), we found no direct evidence for remapping in the period between encoding and cue (**Supplementary Fig. 5**). Second, when dealing with multiple memory contents, contents could be stored at the past location at first and the future location could be considered only after relevant memory content has been selected (in which case the past bias should have *preceded* the future bias). In contrast, our data suggest that the brain simultaneously retains the copy of both past and future-relevant locations in working memory, and (re)activates each during mnemonic selection. Thus, while it is not surprising that the future location is considered [25–27](#), it is far less trivial that both past and future attributes would be retained and (re)activated together. This is our central contribution.

By capitalising on the discrete nature of saccades, we were able to demonstrate the truly joint consideration of past and future attributes, at the single-trial level. As such we were able to bypass a fundamental challenge of disentangling multiple potential single-trial interpretations when only having trial-average data available (for related discussions, see [21,28,29](#)). For example, when only considering the temporally overlapping past and future signals at the trial-average level, it was impossible to tell whether these joint effects resulted from a mix of trials and/or participants that relied on either the past *or* the future. By considering the discrete events – the individual saccades at the single-trial level – we could demonstrate how biases to the past and the future co-existed at the single-saccade level, supporting a truly joint consideration of past and future.

In the current study, we tracked spatial attention using microsaccades (cf. [9](#)–[20](#)). Compared to the commonly used online indicator of spatial attention, such as electrophysiological measures, microsaccades have important features that make them a promising complementary tool for uncovering the mechanisms of spatial attention [11](#),[30](#), including when directed internally as we have shown here. First, as we have discussed above, microsaccades are discrete events, allowing to track spatial attention at the single-trial level. Second, microsaccades are not limited to tracking spatial attention toward the left or right visual field, as electrophysiology indicators often are. These two aspects were paramount to our demonstration of joint single-trial consideration of past and future memory attributes, and are likely to open additional doors for future investigations.

While the past gaze bias that we report here replicates our own prior studies [17](#)–[20](#), here we for the first time demonstrate a similar bias to the future-relevant memory location that is similarly driven by microsaccades. This signal may reflect either of two situations: the selection of a future-copy of the cued memory content or anticipatory attention to its the anticipated location of its associated test-stimulus. Either way, by the nature of our experimental design, this future signal should be considered a content-specific memory attribute for two reasons. First, the two memory contents were always associated with opposite testing locations, hence the observed bias to the relevant future location must be attributed specifically to the cued memory content. Second, we cued which memory item would become tested based on its colour, but the to-be-tested location was dependent on the item's encoding location, regardless of its colour. Hence, consideration of the item's future-relevant location must have been mediated by selecting the memory item itself, as it could not have proceeded via cue colour directly. Accordingly, our data reveal how this future feature can be accessed from memory together with the specific memory content that it is associated with, implying joint storage and utilisation of past and future memory attributes.

Building on the above, at face value, our task may appear like a study that simply combines two established tasks: tasks using retro-cues to study attention in working memory (e.g., [2](#),[31](#)–[33](#)) and tasks using pre-cues to study orienting of spatial attention to an upcoming external stimulus (e.g., [31](#),[32](#),[34](#)–[36](#)). A critical difference with common pre-cue studies, however, is that the cue in our task never directly informed the relevant future location. Rather, as also stressed above, the future location was a feature of the cued memory item (according to the future rule), and not of the cue itself. Note how this type of scenario may not be uncommon in everyday life, such as in our opening example of a bird flying behind a building. Here too, the future relevant location is determined by the bird – i.e. the memory content – itself.

In our study, the past location of the memory items was technically irrelevant for the task and could thus, in principle, be dropped after encoding. One possibility is that participants remapped the two memory items to their future locations soon after encoding, and had started – but not finished – dropping the past location by the time the cue arrived. In such a scenario, the past signal is merely a residual trace of the memory items that serves no purpose but still pulls gaze. Alternatively, however, the past locations may be utilised by the brain to help individuate/separate the two memory items. Moreover, by storing items with regard to multiple spatial frames (cf. [37](#)) – here with regard to both past and future visual locations – it is conceivable that memories may become more robust to decay and/or interference. Also, while in our task past locations were never probed, in everyday life it may be useful to remember where you last saw something before it disappeared behind an occluder. In future work, it will prove interesting to systematically vary to the delay between encoding and cue to assess whether the reliance on the past location gradually dissipates with time (consistent with dropping an irrelevant feature), or whether the past trace remains preserved despite longer delays (consistent with preserving utility for working memory).

Our data complement other recent studies investigating visual working memory in more dynamic contexts [4](#), [18](#), [37](#)–[40](#), and showcase the rich nature of working memory representations. For example, akin to our finding of joint consideration of past and future locations in working memory, recent studies have uncovered the joint consideration of allocentric and egocentric spatial frames for working memory [37](#), [41](#). While storing memory content at a single location (or with reference to a single frame) would appear more intuitive and efficient, our data reveal an intriguing alternative. We speculate that the joint retention of multiple spatial attributes may make memories more robust, as well as more flexible for serving continuously evolving demands during everyday behaviour.

Methods

Ethics

Experimental procedures were reviewed and approved by the local Ethics Committee at the Vrije Universiteit Amsterdam. Each participant provided written consent before participation and was reimbursed €10/hour.

Participants

Twenty-five healthy human volunteers participated in the study (age range: 20–27; 10 male and 15 female; 23 right-handed; 10 corrected-to-normal vision: 5 glasses and 5 lenses). Sample size of 25 was determined a-priori based on previous publications from the lab with similar experimental designs and that relied on the same outcome measure [17](#), [18](#), [20](#). To achieve the intended sample size, three participants were replaced due to chance-level performance.

Stimuli and procedure

Participants performed a visual working memory task in which we orthogonalised the encoding and to-be-tested location in order to track the utilisation of past and future working-memory attributes.

Participants were required to encode and maintain two visual items (tilted coloured gratings) in working memory to later report one of their orientations (**Fig. 1a**). Each trial began with a brief (250 ms) encoding display in which two to-be-memorised gratings with different colours and orientations appeared vertically or horizontally on either side of the fixation dot, at 4 degrees visual angle. After a retention delay of 1250 ms, the fixation dot changed colour for 1000 ms. This colour change served as a 100% valid retro-cue, prompting participants to select the colour-matching target memory item. The retro-cue was followed by another retention delay of 500 ms, before the test display. The test display always contained two black gratings with different orientations that were presented vertically or horizontally on either side of the fixation dot (again at 4 visual degrees to each side). Based on a rule that described in the following paragraph, one of the black gratings was relevant to the task (the test grating), while the other merely served as a filler. After seeing the test display, participants were required to compare the cued memory grating to the relevant test grating and report whether the memory grating should be turned clockwise (using the keyboard button ‘j’) or counterclockwise (using the keyboard button ‘f’) rotated to match the relevant test grating in the test display. After responding, feedback (‘0’ for wrong, or ‘1’ for correct) would be presented for 250 ms at the side of the relevant test grating, indicating the correctness of the response and reinforcing the future rule.

In the encoding display, the gratings were randomly assigned two distinct colours: green (RGB: 133, 194, 18) and purple (RGB: 197, 21, 234) and two distinct orientations ranging from 0° to 180° with a minimum difference of 20° between each other. During the test display, the gratings are

always black (RGB: 64, 64, 64). The relevant grating was always rotated 20 degrees in either a clockwise or counter-clockwise direction compared to the to-be-tested memory grating. The orientation of the irrelevant grating in the test display was chosen randomly.

The unique element of our task was that we dissociated the encoding and testing location by presenting the gratings in the test display on the orthogonal axis as where the items were presented on the encoding display. Relevant examples can be found in **Figure 1a** [↗](#) and **Supplementary Figure 1** [↗](#). For example, if the two memory items appeared vertically at the top and bottom at encoding, the test gratings would appear horizontally to the left and right in the test display (see **Fig. 1a** [↗](#)).

To to-be-tested location was always linked to the encoded location by virtue of a future rule. For counterbalancing reasons, we used four unique rules that were presented across four sessions (**Supplementary Fig. 1** [↗](#)) and that always remained stable within a session. In rule 1: the two memory items were encoded vertically on the top and bottom. If the top memory item was cued, the relevant test grating would be on the right, while if the bottom item would be cued, the relevant test grating would be on the left in the test display. In rule 2, the mapping was reversed: the two memory items would again be encoded vertically, but this time if the top memory item would be cued the relevant test grating is on the left, while if the bottom item is cued the relevant test grating is on the right. Rules 3 and 4 follow the same logic, except that now the memory items are presented horizontally, and the test gratings vertically. Before every session, participants were notified of the future rule that applied to the upcoming session.

To ensure the use of the future rule, we made it task relevant by always presenting two test gratings in the test display of which only one was relevant: the one that matched the future rule (i.e., without applying the future rule, one would not know which test grating was to be used). It is for this reason that we consider the future location an attribute of the memory items, as each memory item was associated with its own unique test location. Note also how our colour cue never directly informed about the future rule or the to-be-tested location. The cue informed which memory item would become tested based on its colour, but the to-be-tested location was dependent on the cued item's encoding location, regardless of its colour (with item colour and location varying randomly across trials). Accordingly, consideration of the item's future-relevant location must have been mediated by selecting the memory item itself, as it could not proceed via the colour cue directly.

In total, the study consisted of 4 sessions, that each contained 5 blocks of 32 trials each. At the start of each session, participants were notified of the session-specific future rule and then practiced the task with the current rule for 16 trials before starting the formal session. We did not include practice trials in our analyses. The study lasted approximately 70 minutes per participant.

Eye-tracking acquisition and pre-processing

Gaze was tracked from a single eye (right eye in all participants except 1 for which the left eye provided a better signal) using an EyeLink 1000 (SR Research) at a sampling rate of 1000 Hz. The eye tracker camera was positioned on the table ~5 cm in front of the monitor and ~65 cm away from the eyes. Gaze position was tracked continuously along the horizontal and vertical axes. Before recording, the built-in calibration and validation protocols from the EyeLink software were used to calibrate the eye tracker.

After recording, the eye-tracking data were converted from the original .edf to the .asc format and analysed in Matlab using the Fieldtrip analysis toolbox [42](#) [↗](#) in combination with custom code. Blinks were marked by detecting NaN clusters in the eye-tracking data. All data from 100 ms before to 100 ms after the detected NaN clusters were also set to NaN to eliminate residual blink artefacts. Finally, data were epoched from -1000 to +2000 ms relative to after the onset of the retro-cue.

Gaze-shift detection

We focused our analysis on spatial biases in gaze shifts (saccades/microsaccades). To identify gaze shifts, we employed a velocity-based method that we established in our prior studies [17](#), [43](#), that builds on other velocity-based detection methods (e.g. [9](#)). First, gaze velocity was calculated by taking the Euclidean distance between temporally successive gaze-position values in the 2-dimensional plane (horizontal and vertical gaze position). Velocity was smoothed with a Gaussian-weighted moving average filter with a 7-ms sliding window (using the built-in function “smoothdata” in MATLAB). When the velocity exceeded a trial-based threshold of 5 times the median velocity, we marked the first sample after the threshold crossing as the onset of a saccade. To avoid counting the same saccade multiple times, a minimum delay of 100 ms between successive saccades was imposed. Saccade magnitude and direction were calculated by estimating the difference between pre-saccade gaze position (–50 to 0 ms before saccade onset) vs. the post-saccade gaze position (50 to 100 ms after saccade onset).

To focus our analysis on (micro)saccades that were driven by attention, we here focused on the “start microsaccade” defined as the saccade that moves gaze away from the fixation dot (as opposed to saccades that bring gaze back to fixation). We extracted start saccades using a method that we recently validated in another study [43](#). For each detected saccade, we estimated the pre- and post-saccade distance from the central fixation dot. If the post-saccade distance was larger than the pre-saccade distance, we defined the saccade as a start microsaccade. To minimise the contribution of gaze drift, we defined gaze positions associated with looking at the central fixation dot using the median gaze position in the fixation period from [–0.8 to –0.2 ms] relative to cue onset.

Time courses of gaze-shift rates (in Hz) were quantified using a sliding time window of 50 ms, advanced in steps of 1 ms. We additionally decomposed shift-rates into a time-magnitude representation (as in [17](#)), showing the time-resolved rate of attention-driven shifts (toward vs. away) per second, as a function of the saccade size (see **Supplementary Fig. 2**). For magnitude sorting, we used successive magnitude bins of 0.2 visual degree in steps of 0.04 visual degree.

Individual-shift level analysis

For our analysis at the single-trial, single-saccade level, we focused on the first start saccade observed during the 200–600 ms post-cue period that we have previously identified as the relevant time window for the effect of interest [17](#). For all ‘first saccades’, we then looked at the spatial distribution of saccade directions relative to the relevant past (encoding) and future (to-be-tested) locations associated with the cued memory item. To facilitate visualisation and quantification, we rotated all detected saccades to a common coordinate system, in which the past location was represented horizontally (left=away, right=toward), and the future location vertically (top=toward, bottom=away).

To obtain the proportional distribution of saccades in this coordinate system (see **Fig. 2A**), we used successive angular bins of 20 degrees in steps of 1 angle degree and simply calculated the proportion of saccades in each angular bin (with angles being defined with reference to past and future locations).

To quantify whether saccades were jointly biased to the past and future (as opposed to either alone), we decomposed these distributions into the bias toward the past (toward vs. away from past) as a function of saccade direction with regard to the future (**Fig. 2b**) and the bias toward the future (toward vs. away from future) as a function of saccade direction with regard to the past (**Fig. 2c**). Finally, we binarised saccades as either going toward or away from the past and

toward or away from the future (**Fig. 2d** [↗](#)). This enabled us to test whether saccades were jointly biased to the past and the future (which would predict an interaction whereby those same saccades that were biased to the past were also biased to the future).

Statistical analysis

To evaluate the reliability of patterns in our gaze data, we used a cluster-based permutation approach ⁴⁴ [↗](#). This method is ideal for evaluating patterns at multiple neighbouring points while circumventing the problem of multiple comparisons. We used this approach for all evaluations involving a series of data, such as along a time axis (**Fig. 1b-d** [↗](#); **Fig. 2b-c** [↗](#)).

To create a permutation distribution, we randomly permuted the trial-average data at the participant level 10000 times and identified the largest clusters found at each time. The p-values of the clusters in the original data were calculated as the proportion of permutations for which the size of the largest cluster after permutation was larger than the size of the observed cluster in the original, non-permuted data. We created the permutation distribution using Fieldtrip with default cluster settings (grouping adjacent same-signed data points that were significant in a mass univariate t-test at a two-sided alpha level of 0.05 and defining cluster size as the sum of all t-values in a cluster).

In addition, statistical evaluation for the single-trial, single-saccade analysis was performed on the binarised quantification of saccade proportions using a two-way repeated-measures ANOVA with the factors past (toward/away) and future (toward/away). ANOVA results were complemented with

Bonferroni-corrected post-hoc t-tests. For measures of effect size we used Partial eta squared (for our ANOVA) and Cohen's d (for follow-up t-tests). P-values of follow-up t-tests were Bonferroni corrected.

Data and code availability

The corresponding data and code can be found at: <https://osf.io/xukzs/> [↗](#) at OSF (doi: 10.17605/OSF.IO/XUKZS)

Acknowledgements

This research was supported by an ERC Starting Grant from the European Research Council (MENTICIPATION, 850636) and an NWO Vidi grant by the Dutch Research Council (grant number 14721) to F.v.E.


Author contributions

B.L. and F.v.E. designed and programmed the experiments, B.L. and S.A. acquired the data. B.L., S.A., and F.v.E. analysed and interpreted the data. B.L. and F.v.E. drafted and revised the manuscript.

Competing interests statement


The authors declare no competing interests.

Supplementary Information – Figures S1-S3



Supplementary Figure 1.


Four possible associations between encoding and test locations.



Supplementary Figure 2.

Early saccade biases by past and future memory attributes are predominantly driven by microsaccades.


Difference in gaze-shift rates toward minus away relative to past location (panel **a**) or future location (panel **b**), as a function of saccade size (y axes). For reference, dashed horizontal lines indicate 1° visual angle. Additionally, for each panel, we separately show the difference in gaze-shift rates (toward minus away) in time course at the top (collapsed over all depicted saccade sizes) and the difference in number of gaze-shift as a function of gaze-shift magnitude to the right (collapsed over all depicted times).



Supplementary Figure 3.

Trials with vertical or horizontal configurations show similar joint consideration of past and future memory attributes.


Conventions as in Main [Figure 1b-d](#), separately for trials that encoded items horizontally and tested items vertically (panel **a**) and trials that encoded items vertically and tested items horizontally (panel **b**).



Supplementary Figure 4.

Gaze biases in extended time window as a complement to [Figure 1](#) and [Supplementary Figure 2](#).

This extended analysis reveals that while the gaze bias towards the past location disappears around 600 ms after cue onset, the gaze bias towards the future location persists (panel **a**) and that while the early (joint) future bias occurs predominantly in the microsaccade range below 1 degree visual angle, the later bias to the future location incorporates larger eye movement that likely involve preparing for optimally perceiving the anticipated test stimulus (panel **b**).



Supplementary Figure 5.

Distribution of saccade directions relative to the future rule from encoding onset.

(**Top** panel) The spatial layouts in the four future rules. (**Middle** panel) Polar distributions of saccades during 0 to 1500 ms after encoding onset (i.e., the period between encoding onset and cue onset). The purple quadrants represent the axis of the future rule and the grey quadrants the orthogonal axis. (**Bottom** panel) Time courses of saccades along the above two axes. We did not observe any sign of a bias along the axis of the future rule itself.

References

1. D'Esposito M., Postle B. R. (2015) **The Cognitive Neuroscience of Working Memory** *Annu. Rev. Psychol* **66**:115–142
2. van Ede F., Nobre A. C. (2023) **Turning Attention Inside Out: How Working Memory Serves Behavior** *Annu. Rev. Psychol* <https://doi.org/10.1146/annurev-psych-021422-041757>
3. Rainer Gregor, Rao S. Chenchal, Miller Earl K. (1999) **Prospective Coding for Objects in Primate Prefrontal Cortex** *J. Neurosci* **19**
4. Brincat S. L., et al. (2021) **Interhemispheric transfer of working memories** *Neuron* **109**:1055–1066
5. Doherty J. R., Rao A., Mesulam M. M., Nobre A. C. (2005) **Synergistic Effect of Combined Temporal and Spatial Expectations on Visual Attention** *J. Neurosci* **25**
6. Woodman G. F., Vogel E. K., Luck S. J. (2012) **Flexibility in visual working memory: Accurate change detection in the face of irrelevant variations in position** *Vis. cogn* **20**:1–28
7. Kahneman D., Treisman A., Gibbs B. J. (1992) **The reviewing of object files: Object-specific integration of information** *Cogn. Psychol* **24**:175–219
8. Zaksas D., Bisley J. W., Pasternak T. (2001) **Motion Information Is Spatially Localized in a Visual Working-Memory Task** *J. Neurophysiol* **86**:912–921
9. Engbert R., Kliegl R. (2003) **Microsaccades uncover the orientation of covert attention** *Vision Res* **43**:1035–1045
10. Hafed Z. M., Clark J. J. (2002) **Microsaccades as an overt measure of covert attention shifts** *Vision Res* **42**:2533–2545
11. Rolfs M. (2009) **Microsaccades: Small steps on a long way** *Vision Res* **49**:2415–2441
12. Laubrock J., Kliegl R., Rolfs M., Engbert R. (2010) **When do microsaccades follow spatial attention?** *Atten. Percept. Psychophys* **72**:683–694
13. Yuval-Greenberg S., Merriam E. P., Heeger D. J. (2014) **Spontaneous Microsaccades Reflect Shifts in Covert Attention** *J. Neurosci* **34**
14. Natalya Shelchkova, Martina Poletti (2020) **Modulations of foveal vision associated with microsaccade preparation** *Proc. Natl. Acad. Sci. U.S.A* **117**:11178–11183
15. Fernández A., Hanning N. M., Carrasco M. (2023) **Transcranial magnetic stimulation to frontal but not occipital cortex disrupts endogenous attention** *Proc. Natl. Acad. Sci. U.S.A* **120**
16. Lowet E., et al. (2018) **Enhanced Neural Processing by Covert Attention only during Microsaccades Directed toward the Attended Stimulus** *Neuron* **99**:207–214

17. Liu B., Nobre A. C., van Ede F. (2022) **Functional but not obligatory link between microsaccades and neural modulation by covert spatial attention** *Nat. Commun* **13**
18. van Ede F., Deden J., Nobre A. C. (2021) **Looking ahead in working memory to guide sequential behaviour** *Curr. Biol* **31**:R779–R780
19. van Ede F., Chekroud S. R., Nobre A. C. (2019) **Human gaze tracks attentional focusing in memorized visual space** *Nat. Hum. Behav* **3**:462–470
20. van Ede F., Board A. G., Nobre A. C. (2020) **Goal-directed and stimulus-driven selection of internal representations** *Proc. Natl. Acad. Sci. U.S.A* **117**:24590–24598
21. Stokes M., Spaak E. (2016) **The Importance of Single-Trial Analyses in Cognitive Neuroscience** *Trends Cogn. Sci* **20**:483–486
22. Nobre A. C., Stokes M. G. (2019) **Premembering Experience: A Hierarchy of Time-Scales for Proactive Attention** *Neuron* **104**:132–146
23. Schacter D. L., Addis D. R., Buckner R. L. (2007) **Remembering the past to imagine the future: the prospective brain** *Nat. Rev. Neurosci* **8**:657–661
24. de Vries Eelke, van Ede Freek (2024) **Microsaccades Track Location-Based Object Rehearsal in Visual Working Memory** *eNeuro* **11**
25. Hayhoe M., Ballard D. (2005) **Eye movements in natural behavior** *Trends Cogn. Sci* **9**:188–194
26. Pelz J. B., Canosa R. (2001) **Oculomotor behavior and perceptual strategies in complex tasks** *Vision Res* **41**:3587–3596
27. Mennie N., Hayhoe M., Sullivan B. (2007) **Look-ahead fixations: anticipatory eye movements in natural tasks** *Exp. Brain Res* **179**:427–442
28. Jones S. R. (2016) **When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning** *Curr. Opin. Neurobiol* **40**:72–80
29. Nobre A. C., van Ede F. (2020) **Under the Mind's Hood: What We Have Learned by Watching the Brain at Work** *J. Neurosci* **40**:89–100
30. Engbert R., Martinez-Conde S., Macknik S. L., Martinez L. M., Alonso J.-M., Tse P. U (2006) **Microsaccades: a microcosm for research on oculomotor control, attention, and visual perception** *Progress in Brain Research* :177–192
31. Griffin I. C., Nobre A. C. (2003) **Orienting Attention to Locations in Internal Representations** *J. Cogn. Neurosci* **15**:1176–1194
32. Panichello M. F., Buschman T. J. (2021) **Shared mechanisms underlie the control of working memory and attention** *Nature* **592**:601–605
33. Souza A. S., Oberauer K. (2016) **In search of the focus of attention in working memory: 13 years of the retro-cue effect** *Atten. Percept. Psychophys* **78**:1839–1860
34. Posner M. I. (1980) **Orienting of Attention** *Q. J. Exp. Psychol* **32**:3–25

35. Worden M. S., Foxe J. J., Wang N., Simpson G. V. (2000) **Anticipatory Biasing of Visuospatial Attention Indexed by Retinotopically Specific α -Bank Electroencephalography Increases over Occipital Cortex** *J. Neurosci* **20**:RC63–RC63
36. Schmidt B. K., Vogel E. K., Woodman G. F., Luck S. J. (2002) **Voluntary and automatic attentional control of visual working memory** *Percept. psychophys* **64**:754–763
37. Draschkow D., Nobre A. C., van Ede F. (2022) **Multiple spatial frames for immersive working memory** *Nat. Hum. Behav* **6**:536–544
38. Chung Y. H., Schurgin M., Brady T. (2022) **The role of motion in visual working memory for dynamic stimuli: more lagged but more precise representations of moving objects** *PsyArXiv* <https://doi.org/10.31234/osf.io/cu3zg>
39. de Vries I. E. J., Slagter H. A., Olivers C. N. L. (2020) **Oscillatory Control over Representational States in Working Memory** *Trends Cogn. Sci* **24**:150–162
40. Xie Y., et al. (2022) **Geometry of sequence working memory in macaque prefrontal cortex** *Science* **375**:632–639
41. Fiehler K., Wolf C., Klinghammer M., Blohm G. (2014) **Integration of egocentric and allocentric information during memory-guided reaching to images of a natural environment** *Front. Hum. Neurosci* **8**
42. Oostenveld R., Fries P., Maris E., Schoffelen J.-M (2011) **FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data** *Comput. Intel. Neurosci* **2011**
43. Liu B., Nobre A. C., van Ede F. (2022) **Microsaccades transiently lateralise EEG alpha activity** *bioRxiv* <https://doi.org/10.1101/2022.09.02.506318>
44. Maris E., Oostenveld R. (2007) **Nonparametric statistical testing of EEG-and MEG-data** *J. Neurosci. Methods* **164**:177–190

Article and author information

Baiwei Liu

Institute for Brain and Behavior Amsterdam, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, The Netherlands

For correspondence: b.liu@vu.nl

ORCID iD: [0000-0002-7857-4456](https://orcid.org/0000-0002-7857-4456)

Zampeta-Sofia Alexopoulou

Institute for Brain and Behavior Amsterdam, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, The Netherlands

Freek van Ede

Institute for Brain and Behavior Amsterdam, Department of Experimental and Applied Psychology, Vrije Universiteit Amsterdam, The Netherlands

For correspondence: freek.van.ede@vu.nl

ORCID iD: [0000-0002-7434-1751](https://orcid.org/0000-0002-7434-1751)

Copyright

© 2023, Liu et al.

This article is distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use and redistribution provided that the original author and source are credited.

Editors

Reviewing Editor

Miriam Spering

The University of British Columbia, Vancouver, Canada

Senior Editor

Yanchao Bi

Beijing Normal University, Beijing, China

Reviewer #1 (Public Review):

In this study, the authors offer a fresh perspective on how visual working memory operates. They delve into the link between anticipating future events and retaining previous visual information in memory. To achieve this, the authors build upon their recent series of experiments that investigated the interplay between gaze biases and visual working memory. In this study, they introduce an innovative twist to their fundamental task. Specifically, they disentangle the location where information is initially stored from the location where it will be tested in the future. Participants are tasked with learning a novel rule that dictates how the initial storage location relates to the eventual test location. The authors leverage participants' gaze patterns as an indicator of memory selection. Intriguingly, they observe that microsaccades are directed towards both the past encoding location and the anticipated future test location. This observation is noteworthy for several reasons. Firstly, participants' gaze is biased towards the past encoding location, even though that location lacks relevance to the memory test. Secondly, there's a simultaneous occurrence of an increased gaze bias towards both the past and future locations. To explore this temporal aspect further, the authors conduct a compelling analysis that reveals the joint consideration of past and future locations during memory maintenance. Notably, microsaccades biased towards the future test location also exhibit a bias towards the past encoding location. In summary, the authors present an innovative perspective on the adaptable nature of visual working memory. They illustrate how information relevant to the future is integrated with past information to guide behavior.

<https://doi.org/10.7554/eLife.90874.2.sa1>

Reviewer #2 (Public Review):

Summary:

The manuscript by Liu et al. reports a task that is designed to examine the extent to which "past" and "future" information is encoded in working memory that combines a retrocue with rules that indicate the location of an upcoming test probe. An analysis of microsaccades on a fine temporal scale shows the extent to which shifts of attention track the location of the encoded item (past) and the location of the future item (test probe). The location of the

encoded grating and test probe were always on orthogonal axes (horizontal, vertical) so that biases in microsaccades could be used to track shifts of attention to one or the other axis (or mixtures of the two). The overall goal here was then to (1) create a methodology that could tease apart memory for the past and future, respectively, (2) to look at the time-course attention to past/future, and (3) to test the extent to which microsaccades might jointly encode past and future memoranda. Finally, some remarks are made about the plausibility of various accounts of working memory encoding/maintenance based on the examination of these time-courses.

Strengths:

This research has several notable strengths. It has a clear statement of its aims, is lucidly presented, and uses a clever experimental design that neatly orthogonalized "past" and "future" as operationalized by the authors. Figure 1b-d shows fairly clearly that saccade directions have an early peak (around 300ms) for the past and a "ramping" up of saccades moving in the forward direction. This seems to be a nice demonstration that the method can measure shifts of attention at a fine temporal resolution and differentiate past from future oriented saccades due to the orthogonal cue approach. The second analysis shown in Figure 2, reveals a dependency in saccade direction such that saccades toward the probe future were more likely also to be toward the encoded location than away from the encoded direction. This suggests saccades are jointly biased by both locations "in memory". The "central contribution" (as the authors characterize it) is that "the brain simultaneously retains the copy of both past and future-relevant locations in working memory, and (re)activates each during mnemonic selection", and that: "... while it is not surprising that the future location is considered, it is far less trivial that both past and future attributes would be retained and (re)activated together. This is our central contribution." The authors provide a nuanced analysis that offers persuasive evidence that past and future representations are jointly maintained in memory.

<https://doi.org/10.7554/eLife.90874.2.sa0>

Author response:

The following is the authors' response to the previous reviews.

eLife assessment

This important study advances our understanding of how past and future information is jointly considered in visual working memory by studying gaze biases in a memory task that dissociates the locations during encoding and memory tests. The evidence supporting the conclusions is convincing, with state-of-the-art gaze analyses that build on a recent series of experiments introduced by the authors. This work, with further improvements incorporating the existing literature, will be of broad interest to vision scientists interested in the interplay of vision, eye movements, and memory.

We thank the Editors and the Reviewers for their enthusiasm and appreciation of our task, our findings, and our article. We also wish to thank the Reviewers for their constructive comments that we have embraced to improve our article. Please find below our point-by-point responses to this valuable feedback, where we also state relevant revisions that we have made to our article.

In addition, please note that we have now also made our data and code publicly available.

Reviewer 1, Comments:

In this study, the authors offer a fresh perspective on how visual working memory operates. They delve into the link between anticipating future events and retaining previous visual information in memory. To achieve this, the authors build upon their recent series of experiments that investigated the interplay between gaze biases and visual working memory. In this study, they introduce an innovative twist to their fundamental task. Specifically, they disentangle the location where information is initially stored from the location where it will be tested in the future. Participants are tasked with learning a novel rule that dictates how the initial storage location relates to the eventual test location. The authors leverage participants' gaze patterns as an indicator of memory selection. Intriguingly, they observe that microsaccades are directed toward both the past encoding location and the anticipated future test location. This observation is noteworthy for several reasons. Firstly, participants' gaze is biased towards the past encoding location, even though that location lacks relevance to the memory test. Secondly, there's a simultaneous occurrence of an increased gaze bias towards both the past and future locations. To explore this temporal aspect further, the authors conduct a compelling analysis that reveals the joint consideration of past and future locations during memory maintenance. Notably, microsaccades biased towards the future test location also exhibit a bias towards the past encoding location. In summary, the authors present an innovative perspective on the adaptable nature of visual working memory. They illustrate how information relevant to the future is integrated with past information to guide behavior.

Thank you for your enthusiasm for our article and findings as well as for your constructive suggestions for additional analyses that we respond to in detail below.


This short manuscript presents one experiment with straightforward analyses, clear visualizations, and a convincing interpretation. For their analysis, the authors focus on a single time window in the experimental trial (i.e., 0-1000 ms after retro cue onset). While this time window is most straightforward for the purpose of their study, other time windows are similarly interesting for characterizing the joint consideration of past and future information in memory. First, assessing the gaze biases in the delay period following the cue offset would allow the authors to determine whether the gaze bias towards the future location is sustained throughout the entire interval before the memory test onset. Presumably, the gaze bias towards the past location may not resurface during this delay period, but it is unclear how the bias towards the future location develops in that time window. Also, the disappearance of the retro cue constitutes a visual transient that may leave traces on the gaze biases which speaks again for assessing gaze biases also in the delay period following the cue offset.

Thank you for raising this important point. We initially focused on the time window during the cue given that our central focus was on gaze-biases associated with mnemonic item selection. By zooming in on this window, we could best visualize our main effects of interest: the joint selection (in time) of past and future memory attributes.

At the same time, we fully agree that examining the gaze biases over a more extended time window yields a more comprehensive view of our data. To this end, we have now also extended our analysis to include a wider time range that includes the period between cue offset (1000 ms after cue onset) and test onset (1500 ms after cue onset). We present these data below. Because we believe our future readers are likely to be interested in this as well, we have now added this complementary visualization as Supplementary Figure 4 (while preserving the focus in our main figure on the critical mnemonic selection period of interest).

Author response image 1.

Supplementary Figure 4. Gaze biases in extended time window as a complement to Figure 1 and Supplementary Figure 2. This extended analysis reveals that while the gaze bias towards the past location disappears around 600 ms after cue onset, the gaze bias towards the future location persists (panel a) and that while the early (joint) future bias occurs predominantly in the microsaccade range below 1 degree visual angle, the later bias to the future location incorporates larger eye movement that likely involve preparing for optimally perceiving the anticipated test stimulus (panel b).



This extended analysis reveals that while the gaze bias towards the past location disappears around 600 ms after cue onset (consistent with our prior reports of this bias), the gaze bias towards the future location persists. Moreover, as revealed by the data in panel b above, while the early (joint) future bias occurs predominantly in the microsaccade range below 1 degree visual angle, the later bias to the future location incorporates larger eye movement that likely involve preparing for optimally perceiving the anticipated test stimulus.

We now also call out these additional findings and figure in our article:

Page 2 (Results): “Gaze biases in both axes were driven predominantly by microsaccades (Supplementary Fig. 2) and occurred similarly in horizontal-to-vertical and vertical-to-horizontal trials (Supplementary Fig. 3). Moreover, while the past bias was relatively transient, the future bias continued to increase in anticipation of the of the test stimulus and increasingly incorporated eye-movements beyond the microsaccade range (see Supplementary Fig. 4 for a more extended time range)”.

Moreover, assessing the gaze bias before retro-cue onset allows the authors to further characterize the observed gaze biases in their study. More specifically, the authors could determine whether the future location is considered already during memory encoding and the subsequent delay period (i.e., before the onset of the retro cue). In a trial, participants encode two oriented gratings presented at opposite locations. The future rule indicates the test locations relative to the encoding locations. In their example (Figure 1a), the test locations are shifted clockwise relative to the encoding location. Thus, there are two pairs of relevant locations (each pair consists of one stimulus location and one potential test location) facing each other at opposite locations and therefore forming an axis (in the illustration the axis would go from bottom left to top right). As the future rule is already known to the participants before trial onset it is possible that participants use that information already during encoding. This could be tested by assessing whether more microsaccades are directed along the relevant axis as compared to the orthogonal axis. The authors should assess whether such a gaze bias exists already before retro cue onset and discuss the theoretical consequences for their main conclusions (e.g., is the future location only jointly used if the test location is implicitly revealed by the retro cue).


Thank you – this is another interesting point. We fully agree that additional analysis looking at the period prior to retrocue onset may also prove informative. In accordance with the suggested analysis, we have therefore now also analysed the distribution of saccade directions (including in the period from encoding to retrocue) as a function of the future rule (presented below, and now also included as Supplementary Fig. 5). Complementary recent work from our lab has shown how microsaccade directions can align to the axis of memory contents during retention (see de Vries & van Ede, eNeuro, 2024). Based on this finding, one may predict that if participants retain the items in a remapped fashion, their microsaccades may align with the axis of the future rule, and this could potentially already happen prior to cue onset.

These complementary analyses show that saccade directions are predominantly influenced by the encoding locations rather than the test locations, as seen most clearly by the saccade distribution plots in the middle row of the figure below. To obtain time-courses, we categorized saccades as occurring along the axis of the future rule or along the orthogonal axis (bottom row of the figure below). Like the distribution plots, these time course plots also did not reveal any sign of a bias along the axis of the future rule itself.

Importantly, note how this does not argue against our main findings of joint selection of past and future memory attributes, as for that central analysis we focused on saccade biases that were specific to the selected memory item, whereas the analyses we present below focus on biases in the axes in which both memory items are defined; not only the cued/selected memory item.

Author response image 2.

Supplementary Figure 5. Distribution of saccade directions relative to the future rule from encoding onset. (Top panel) The spatial layouts in the four future rules. (Middle panel) Polar distributions of saccades during 0 to 1500 ms after encoding onset (i.e., the period between encoding onset and cue onset). The purple quadrants represent the axis of the future rule and the grey quadrants the orthogonal axis. (Bottom panel) Time courses of saccades along the above two axes. We did not observe any sign of a bias along the axis of the future rule itself.



We agree that these additional results are important to bring forward when we interpret our findings. Accordingly, we now mention these findings at the relevant section in our Discussion:

Page 5 (Discussion): “First, memory contents could have directly been remapped (cf. 4,24–26) to their future-relevant location. However, in this case, one may have expected to exclusively find a future-directed gaze bias, unlike what we observed. Moreover, using a complementary analysis of saccade directions along the axis of the future rule (cf. 24), we found no direct evidence for remapping in the period between encoding and cue (Supplementary Fig. 5)”.

Reviewer 2, Comments:

The manuscript by Liu et al. reports a task that is designed to examine the extent to which "past" and "future" information is encoded in working memory that combines a retro cue with rules that indicate the location of an upcoming test probe. An analysis of microsaccades on a fine temporal scale shows the extent to which shifts of attention track the location of the location of the encoded item (past) and the location of the future item (test probe). The location of the encoded grating of the test probe was always on orthogonal axes (horizontal, vertical) so that biases in microsaccades could be used to track shifts of attention to one or the other axis (or mixtures of the two). The overall goal here was then to (1) create a methodology that could tease apart memory for the past and future, respectively, (2) to look at the time-course attention to past/future, and (3) to test the extent to which microsaccades might jointly encode past and future memoranda. Finally, some remarks are made about the plausibility of various accounts of working memory encoding/maintenance based on the examination of these time courses.

Strengths:

This research has several notable strengths. It has a clear statement of its aims, is lucidly presented, and uses a clever experimental design that neatly orthogonalizes "past" and "future" as operationalized by the authors. Figure 1b-d shows fairly clearly that saccade

directions have an early peak (around 300ms) for the past and a "ramping" up of saccades moving in the forward direction. This seems to be a nice demonstration the method can measure shifts of attention at a fine temporal resolution and differentiate past from future-oriented saccades due to the orthogonal cue approach. The second analysis shown in Figure 2, reveals a dependency in saccade direction such that saccades toward the probe future were more likely also to be toward the encoded location than away from the encoded direction. This suggests saccades are jointly biased by both locations "in memory".

Thank you for your overall appreciation of our work and for highlighting the above strengths. We also thank you for your constructive comments and call for clarifications that we respond to below.

Weaknesses:

(1) The "central contribution" (as the authors characterize it) is that "the brain simultaneously retains the copy of both past and future-relevant locations in working memory, and (re)activates each during mnemonic selection", and that: "... while it is not surprising that the future location is considered, it is far less trivial that both past and future attributes would be retained and (re)activated together. This is our central contribution." However, to succeed at the task, participants must retain the content (grating orientation, past) and probe location (future) in working memory during the delay period. It is true that the location of the grating is functionally irrelevant once the cue is shown, but if we assume that features of a visual object are bound in memory, it is not surprising that location information of the encoded object would bias processing as indicated by microsaccades. Here the authors claim that joint representation of past and future is "far less trivial", this needs to be evaluated from the standpoint of prior empirical data on memory decay in such circumstances, or some reference to the time-course of the "unbinding" of features in an encoded object.

Thank you. We agree that our participants have to use the future rule – as otherwise they do not know to which test stimulus they should respond. This was a deliberate decision when designing the task. Critically, however, this does not require (nor imply) that participants have to incorporate and apply the rule to both memory items already prior to the selection cue. It is at least as conceivable that participants would initially retain the two items at their encoded (past) locations, then wait for the cue to select the target memory item, and only then consider the future location associated with the target memory item. After all, in every trial, there is only 1 relevant future location: the one associated with the cued memory item. The time-resolved nature of our gaze markers argues against such a scenario, by virtue of our observation of the joint (simultaneous) consideration of past and future memory attributes (as opposed to selection of past-before-future). These temporal dynamics are central to the insights provided by our study.

In our view, it is thus not obvious that the rule would be applied at encoding. In this sense, we do not assume that the future location is part of both memory objects from encoding, but rather ask whether this is the case – and, if so, whether the future location takes over the role of the past location, or whether past and future locations are retained jointly.

Our statements regarding what is “trivial” and what is “less trivial” regard exactly this point: it is trivial that the future is considered (after all, our task demanded it). However, it is less trivial that (1) the future location was already available at the time of initial item selection (as reflected in the simultaneous engagement of past and future locations), and (2) that in presence of the future location, the past location was still also present in the observed gaze biases.

Having said that, we agree that an interesting possibility is that participants remap both memory items to their future-relevant locations ahead of the cue, but that the past location is not yet fully “unbound” by the time of the cue. This may trigger a gaze bias not only to the new future location but also to the “sticky” (unbound) past location. We now acknowledge this possibility in our discussion (also in response to comment 3 below) where we also suggest how future work may be able to tap into this:

Page 6 (Discussion): “In our study, the past location of the memory items was technically irrelevant for the task and could thus, in principle, be dropped after encoding. One possibility is that participants remapped the two memory items to their future locations soon after encoding, and had started – but not finished – dropping the past location by the time the cue arrived. In such a scenario, the past signal is merely a residual trace of the memory items that serves no purpose but still pulls gaze. Alternatively, however, the past locations may be utilised by the brain to help individuate/separate the two memory items. Moreover, by storing items with regard to multiple spatial frames (cf. 37) – here with regard to both past and future visual locations – it is conceivable that memories may become more robust to decay and/or interference. Also, while in our task past locations were never probed, in everyday life it may be useful to remember where you last saw something before it disappeared behind an occluder. In future work, it will prove interesting to systematically vary to the delay between encoding and cue to assess whether the reliance on the past location gradually dissipates with time (consistent with dropping an irrelevant feature), or whether the past trace remains preserved despite longer delays (consistent with preserving utility for working memory).”

(2) The authors refer to "future" and "past" information in working memory and this makes sense at a surface level. However, once the retrocue is revealed, the "rule" is retrieved from long-term memory, and the feature (e.g. right/left, top/bottom) is maintained in memory like any other item representation. Consider the classic test of digit span. The digits are presented and then recalled. Are the digits of the past or future? The authors might say that one cannot know, because past and future are perfectly confounded. An alternative view is that some information in working memory is relevant and some is irrelevant. In the digit span task, all the digits are relevant. Relevant information is relevant precisely because it is thought to be necessary in the future. Irrelevant information is irrelevant precisely because it is not thought to be needed in the immediate future. In the current study, the orientation of the grating is relevant, but its location is irrelevant; and the location of the test probe is also relevant.

Thank you for this stimulating reflection. We agree that in our set-up, past location is technically “task-irrelevant” while future location is certainly “task-relevant”. At the same time, the engagement of the past location suggests to us that the brain uses past location for the selection – presumably because the brain uses spatial location to help individuate/separate the items, even if encoded locations are never asked about. Therefore, whether something is relevant or irrelevant ultimately depends on how one defines relevance (past location may be relevant/useful for the brain even if technically irrelevant from the perspective of the task). In comparison, the use of “past” and “future” may be less ambiguous.

It is also worth noting how we interpret our findings in relation to demands on visual working memory, inspired by dynamic situations whereby visual stimuli may be last seen at one location but expected to re-appear at another, such as a bird disappearing behind a building (the example in our introduction). Thus, past for us does not refer to the memory item *per se* (like in the digit span analogue) but, rather, quite specifically to the past location of a dynamic visual stimulus in memory (which, in our experiment, was operationalised by the future rule, for convenience).

(3) It is not clear how the authors interpret the "joint representation" of past and future. Put aside "future" and "past" for a moment. If there are two elements in memory, both of which are associated with spatial bindings, the attentional focus might be a spatial average of the associated spatial indices. One might also view this as an interference effect, such that the location of the encoded location attracts spatial attention since it has not been fully deleted/removed from working memory. Again, for the impact of the encoded location to be exactly zero after the retrieval cue, requires zero interference or instantaneous decay of the bound location information. It would be helpful for the authors to expand their discussion to further explain how the results fit within a broader theoretical framework and how it fits with empirical data on how quickly an irrelevant feature of an object can be deleted from working memory.

Thank you also for this point (that is related to the two points above). As we stated in our reply to comment 1 above, we agree that one possibility is that the past location is merely “sticky” and pulls the task-relevant future bias toward the past location. If so, our time courses suggest that such “pulling” occurs only until approximately 600 ms after cue onset, as the past bias is only transient. An alternative interpretation is that the past location may not be merely a residual irrelevant trace, but actually be useful and used by the brain.

For example, the encoded (past) item locations provide a coordinate system in which to individuate/separate the two memory items. While the future locations also provide such a coordinate system, the brain may benefit from holding onto both coordinate systems at the same time, rendering our observation of joint selection in both frames. Indeed, in a recent VR experiment in which we had participants (rather than the items) rotate, we also found evidence for the joint use of two spatial frames, even if neither was technically required for the upcoming task (see Draschkow, Nobre, van Ede, *Nature Human Behaviour*, 2022). Though highly speculative at this stage, such reliance on multiple spatial frames may make our memories more robust to decay and/or interference. Moreover, while past location was never explicitly probed in our task, in daily life the past location may sometimes (unexpectedly) become relevant, hence it may be useful to hold onto it, just in case. Thus, considering the past location merely as an “irrelevant feature” (that takes time to delete) may not do sufficient justice to the potential roles of retaining past locations of dynamic visual objects held in working memory.

As also stated in response to comment 1 above, we now added these relevant considerations to our Discussion:

Page 5 (Discussion): “In our study, the past location of the memory items was technically irrelevant for the task and could thus, in principle, be dropped after encoding. One possibility is that participants remapped the two memory items to their future locations soon after encoding, and had started – but not finished – dropping the past location by the time the cue arrived. In such a scenario, the past signal is merely a residual trace of the memory items that serves no purpose but still pulls gaze. Alternatively, however, the past locations may be utilised by the brain to help individuate/separate the two memory items. Moreover, by storing items with regard to multiple spatial frames (cf. 37) – here with regard to both past and future visual locations – it is conceivable that memories may become more robust to decay and/or interference. Also, while in our task past locations were never probed, in everyday life it may be useful to remember where you last saw something before it disappeared behind an occluder. In future work, it will prove interesting to systematically vary to the delay between encoding and cue to assess whether the reliance on the past location gradually dissipates with time (consistent with dropping an irrelevant feature), or whether the past trace remains preserved despite longer delays (consistent with preserving utility for working memory).”

Reviewer 3, Comments:

This study utilizes saccade metrics to explore, what the authors term the "past and future" of working memory. The study features an original design: in each trial, two pairs of stimuli are presented, first a vertical pair and then a horizontal one. Between these two pairs comes the cue that points the participant to one target of the first pair and another of the second pair. The task is to compare the two cued targets. The design is novel and original but it can be split into two known tasks - the first is a classic working memory task (a post-cue informs participants which of two memorized items is the target), which the authors have used before; and the second is a classic spatial attention task (a pre-cue signal that attention should be oriented left or right), which was used by numerous other studies in the past. The combination of these two tasks in one design is novel and important, as it enables the examination of the dynamics and overlapping processes of these tasks, and this has a lot of merit. However, each task separately is not new. There are quite a few studies on working memory and microsaccades and many on spatial attention and microsaccades. I am concerned that the interpretation of "past vs. future" could mislead readers to think that this is a new field of research, when in fact it is the (nice) extension of an existing one. Since there are so many studies that examined pre-cues and post-cues relative to microsaccades, I expected the interpretation here to rely more heavily on the existing knowledge base in this field. I believe this would have provided a better context of these findings, which are not only on "past" vs. "future" but also on "working memory" vs. "spatial attention".

Thank you for considering our findings novel and important, while at the same time reminding us of the parallels to prior tasks studying spatial attention in perception and working memory. We fully agree that our task likely engages both attention to the (past) memory item as well as spatial attention to the upcoming (future) test stimulus. At the same time, there is a critical difference in spatial attention for the future in our task compared with ample prior tasks engaging spatial cueing of attention for perception. In our task, the cue never directly cues the future location. Rather, it exclusively cues the relevant memory item. It is the memory item that is associated with the relevant future location, according to the future rule. This integration of the rule-based future location into the memory representation is distinct from classical spatial-attention tasks in which attention is cued directly to a specific location via, for example, a spatial cue such as an arrow.

Thus, if we wish to think about our task as engaging cueing of spatial attention for perception, we have to at least also invoke the process of cueing the relevant location via the appropriate memory item. We feel it is more parsimonious to think of this as attending to both the past and future location of a dynamic visual object in working memory.

If we return to our opening example, when we see a bird disappear behind a building, we can keep in working memory where we last saw it, while anticipating where it will re-appear to guide our external spatial attention. Here too, spatial attention is fully dependent on working-memory content (the bird itself) – mirroring the dynamic semng in our study. Thus, we believe our findings contribute a fresh perspective, while of course also extending established fields. We now contextualize our finding within the literature and clarify our unique contribution in our revised manuscript:

Page 5 (Discussion): "Building on the above, at face value, our task may appear like a study that simply combines two established tasks: tasks using retro-cues to study attention in working memory (e.g., 2,31-33) and tasks using pre-cues to study orienting of spatial attention to an upcoming external stimulus (e.g., 31,32,34-36). A critical difference with common pre-cue studies, however, is that the cue in our task never directly informed the relevant future location. Rather, as also stressed above, the future location was a feature of the cued memory

item (according to the future rule), and not of the cue itself. Note how this type of scenario may not be uncommon in everyday life, such as in our opening example of a bird flying behind a building. Here too, the future relevant location is determined by the bird – i.e. the memory content – itself.”

Reviewer 2, Recommendations:

It would be helpful to set up predictions based on existing working memory models. Otherwise, the claim that the joint coding of past/future is "not trivial" is simply asserted, rather than contradicting an existing model or prior empirical results. If the non-trivial aspect is simply the ability to demonstrate the joint coding empirically through a good experimental design, make it clear that this is the contribution. For example, it may be that prevailing models predict exactly this finding, but nobody has been able to demonstrate it cleanly, as the authors do here. So the non-triviality is not that the result contradicts working memory models, but rather relates to the methodological difficulty of revealing such an effect.

Thank you for your recommendation. First, please see our point-by-point responses to the individual comments above, where we also state relevant changes that we have made to our article, and where we clarify what we meant with “non trivial”. As we currently also state in our introduction, our work took as a starting point the framework that working memory is inherently about the past while being for the future (cf. van Ede & Nobre, Annual Review of Psychology, 2023). By virtue of our unique task design, we were able to empirically demonstrate that visual contents in working memory are selected via both their past and their future-relevant locations – with past and future memory attributes being engaged together in time. With “not trivial” we merely intend to make clear that there are viable alternatives than the findings we observed. For example, past could have been replaced by the future, or it could have been that item selection (through its past location) was required before its future-relevant location could be considered (i.e. past-before-future, rather than joint selection as we reported). We outline these alternatives in the second paragraph of our Discussion:

Page 5 (Discussion): “Our finding of joint utilisation of past and future memory attributes emerged from at least two alternative scenarios of how the brain may deal with dynamic everyday working memory demands in which memory content is encoded at one location but needed at another.

First, [...]”

Our work was not motivated from a particular theoretical debate and did not aim to challenge ongoing debates in the working-memory literature, such as: slot vs. resource, active vs. silent coding, decay vs. interference, and so on. To our knowledge, none of these debates makes specific claims about the retention and selection of past and future visual memory attributes – despite this being an important question for understanding working memory in dynamics everyday settings, as we hoped to make clear by our opening example.

Reviewer 3, Recommendations:

I recommend that the present findings be more clearly interpreted in the context of previous findings on working memory and attention. The task design includes two components - the first (post-cue) is a classic working memory task and the second (the pre-cue) is a classic spatial attention design. Both components were thoroughly studied in the past and this previous knowledge should be better integrated into the present conclusions. I specifically feel uncomfortable with the interpretation of past vs. future. I find this framework to be misleading because it reads like this paper is on a topic that is completely new and never studied before, when in fact this is a study on the interaction

between working memory and spatial attention. I recommend the authors minimize this past-future framing or be more explicit in explaining how this new framework relates to the more common terminology in the field and make sure that the findings are not presented in a vacuum, as another contribution to the vibrant field that they are part of.

Thank you for these recommendations. Please also see our point-by-point responses to the individual comments above. Here, we explained our logic behind using the terminology of past vs. future (in addition, see also our response to point 2 or reviewer 2). Here, we also stated relevant changes that we have made to our manuscript to explain how our findings complement – but are also distinct from – prior tasks that used pre-cues to direct spatial attention to an upcoming stimulus. As we explained above, in our task, the cue itself never contained information about the upcoming test location. Rather, the upcoming test location was a property of the memory item (given the future rule). Hence, we referred to this as a “future attribute” of the cued memory item, rather than as the “cued location” for external spatial attention. Still, we agree the future bias likely (also) reflects spatial allocation to the upcoming test array, and we explicitly acknowledge this in our discussion. For example:

Page 5 (Discussion): “This signal may reflect either of two situations: the selection of a future-copy of the cued memory content or anticipatory attention to its the anticipated location of its associated test-stimulus. Either way, by the nature of our experimental design, this future signal should be considered a content-specific memory attribute for two reasons. First, the two memory contents were always associated with opposite testing locations, hence the observed bias to the relevant future location must be attributed specifically to the cued memory content. Second, we cued which memory item would become tested based on its colour, but the to-be-tested location was dependent on the item’s encoding location, regardless of its colour. Hence, consideration of the item’s future-relevant location must have been mediated by selecting the memory item itself, as it could not have proceeded via cue colour directly.”

Page 6 (Discussion): “Building on the above, at face value, our task may appear like a study that simply combines two established tasks: tasks using retro-cues to study attention in working memory (e.g., 2, 31–33) and tasks using pre-cues to study orienting of spatial attention to an upcoming external stimulus (e.g., 31, 32, 34–36). A critical difference with common pre-cue studies, however, is that the cue in our task never directly informed the relevant future location. Rather, as also stressed above, the future location was a feature of the cued memory item (according to the future rule), and not of the cue itself. Note how this type of scenario may not be uncommon in everyday life, such as in our opening example of a bird flying behind a building. Here too, the future relevant location is determined by the bird – i.e. the memory content – itself.”