# Journal of Experimental Psychology: Human Perception and Performance

**Latent Memory Traces for Prospective Items in Visual Working Memory**

Luzi Xu, Andre Sahakian, Surya Gayet, Chris L. E. Paffen, and Stefan Van der Stigchel

# Latent Memory Traces for Prospective Items in Visual Working Memory

Luzi Xu, Andre Sahakian, Surya Gayet, Chris L. E. Paffen, and Stefan Van der Stigchel
Department of Experimental Psychology, Helmholtz Institute, Utrecht University

Visual working memory (VWM) is a capacity-limited cognitive system that is utilized for enabling goal-directed actions. When sampling items for VWM storage, however, observers are often exposed to other items that are not selected for imminent action (hereafter: "prospective items"). Here, we asked whether such exposure leads to memory buildup of these prospective items, facilitating subsequent VWM encoding for imminent action. In a series of experiments, we addressed this question using a copying task, in which participants attempted to reproduce a model display by placing items in an adjacent empty grid. To investigate whether a memory is formed for prospective items, we swapped the position of unplaced items in the model and compared copying task performance to a condition in which items remained stable. The results show that, when prospective items remained stable, participants took less time inspecting the model when encoding these items later (compared to when they were swapped). This reduced inspection duration was not accompanied by a higher number of inspections or an increase in errors. We conclude that the memory system gradually builds up latent memory traces of items that are not selected for imminent action, thus increasing the efficiency of subsequent VWM encoding.

---

***Public Significance Statement***
The number of objects we can interact with, and the number of items we can memorize at any given time is very limited. When doing groceries, for example, we are often unable to memorize and find all items on the shopping list in one go. We are more likely to select a few items first (e.g., pears, pasta, and yogurt) and leave the remaining items (e.g., tomatoes and eggs) for prospective memorization and action. In this study, we reveal that, while processing some items first for imminent action, we form latent memory traces for the other items. These memory traces reduced the time needed to encode the prospective items into working memory when they were selected for action later. The present work shows one way in which the mnemonic system circumvents its capacity limitations to efficiently operate in a complex visual world.

---

*Keywords:* visual working memory, action, implicit memory, visual experience, copying task

*Supplemental materials:* https://doi.org/10.1037/xhp0001257.supp

---

Our visual environment provides us with an abundance of rich information, comprising elements that are relevant as well as elements that are irrelevant to our current goals. To engage in goal-directed behavior, we utilize visual working memory (VWM) to retain task-relevant information while avoiding interference from irrelevant information. Interference can be avoided by filtering out or suppressing irrelevant information in working memory (e.g., Feldmann-Wüstefeld & Vogel, 2019; Kuo et al., 2012; Zanto & Gazzaley, 2009) and related cognitive processes (e.g., attention; Arita et al., 2012; Gaspelin et al., 2015; Wang & Theeuwes, 2018). In naturalistic, visually rich real-world settings, however, some task-relevant items might not be encoded into VWM because the amount of task-relevant items exceeds

---

capacity limitations. Relatedly, the relevance of stimuli might change over time: stimuli that are irrelevant for immediate action might be relevant for upcoming behavior. In such cases, observers are visually exposed to items that are (prospectively) relevant, but are not yet encoded in VWM for immediate action. To illustrate this point, when memorizing items from a shopping list in a supermarket, we often select a few items as current targets (e.g., pears, pasta, and yogurt) while temporarily ignoring prospective targets (e.g., tomatoes and eggs) that will be selected later. In scenarios like this, observers are visually exposed to items that are not encoded into VWM for immediate use, but that will become relevant to behavior prospectively. Previous studies, however, have mainly focused on the processing of completely task-irrelevant items that will never be probed (e.g., Feldmann-Wüstefeld & Vogel, 2019; Zanto & Gazzaley, 2009) and the processing of items already encoded into memory that are not immediately probed (van Ede & Nobre, 2023). These studies do not address a more naturalistic scenario in which items that are accessible but not yet encoded into working memory might be needed for future tasks (e.g., tomatoes and eggs in the shopping example). To our knowledge, despite the prevalence of such situations in day-to-day VWM use, it remains unclear whether observers build memory traces for such prospective items through prior exposure (e.g., while sampling current targets—pears, pasta, and yogurt), which might benefit subsequent encoding of these items into VWM later.

Although this issue has not been addressed before, the existing literature points toward two competing hypotheses: One hypothesis is that memory traces are built up for prospective items (so that prior exposure benefits subsequent VWM encoding), and the other hypothesis is that this does not occur. The first hypothesis is derived from studies suggesting that visual memory can concurrently process surrounding items while maintaining currently selected items (Boduroglu & Shah, 2014; Hollingworth, 2007; O'Donnell et al., 2018). First, such studies have suggested that encoding items into working memory involves not only targets but also nontarget information such as surrounding objects (e.g., Hollingworth, 2007), for the encoding, maintenance, or retrieval of current targets. Previous studies indeed demonstrated that VWM performance regarding relevant targets is impaired by changes in irrelevant surrounding items (O'Donnell et al., 2018), and that observers develop memory traces for task-irrelevant features of multifeature items (Boduroglu & Shah, 2014). Based on these findings, it is conceivable that prospective items act like contextual cues for the processing of current targets, and are therefore actively processed in VWM. Another line of evidence supporting the buildup of memory traces comes from studies showing that prior visual experience facilitates visual perception (e.g., implicit learning of repeated contexts, e.g., Chun & Jiang, 1998; Fiser & Aslin, 2002; priming, e.g., Kristjánsson & Campana, 2010; familiarity, e.g., Krueger, 1975) and VWM (e.g., implicit learning of repeated items, e.g., Umemoto et al., 2010; familiarity, Blalock, 2015; Calmels et al., 2012; Jackson & Raymond, 2008; Lorenc et al., 2014; Ngiam et al., 2019; Reder et al., 2016; Scolari et al., 2008; but also see Chen et al., 2006). Finally, visual exposure to repeated stimuli can engage implicit learning (e.g., contextual learning, Chun & Jiang, 1998) or long-term memory (Woodman et al., 2013; Xu et al., 2024). Since engaging implicit learning mechanisms or long-term memory requires few cognitive resources (e.g., attention) (Brady et al., 2024), it can be hypothesized that the buildup of memory traces for prospective items can happen relatively effortlessly, and in parallel to VWM processing of current targets.

On the other hand, there is also work supporting the hypothesis that visual memory does not store prospectively relevant items in concurrence with items that are selected for imminent action. According to the selection-for-action hypothesis (Allport, 1989), for example, current targets for action compete with other items for cognitive resources (including VWM resources, see Oberauer & Lin, 2017). In line with this idea, previous studies showed that observers use top-down attentional control to filter out information irrelevant to their current tasks (e.g., Arita et al., 2012; Gaspelin et al., 2015; Wang & Theeuwes, 2018). Given that VWM is influenced by top-down attentional control (Gazzaley & Nobre, 2012), one might hypothesize that memory traces for prospective items do not accumulate. Supporting this, previous studies have shown that irrelevant features of memorized items cannot be retrieved from neural delay activity patterns (Serences et al., 2009). Moreover, some studies indicate that irrelevant information is actively suppressed during VWM tasks to optimize performance (Feldmann-Wüstefeld & Vogel, 2019; Zanto & Gazzaley, 2009). Filtering out irrelevant information from VWM can also reduce encoding errors (Emrich & Ferber, 2012) and decrease forgetting (Lewis-Peacock & Norman, 2014). Taken together, there is a substantial body of theoretical and empirical work supporting the possibility that items that are irrelevant for imminent action will not lead to the buildup of memory traces, and might even be actively suppressed.

To dissociate between the hypothesis that latent memory traces are created for prospective items, and the hypothesis that this does not occur, we compared memory use in a condition that does allow for the buildup of memory traces for prospective items, with a condition that does not. To do so, we used a so-called copying task (Ballard et al., 1995; Sahakian et al., 2023; Somai et al., 2020), in which participants are tasked to reproduce a configuration of items (the model grid) in an empty workspace, using building blocks (see Figure 1). To execute this task, observers need to inspect the model grid and maintain one or more individual items in VWM, select those items from the resources grid, and move them to the corresponding position in the workspace grid. This task allowed us to test whether repeated visual exposure to items in the model grid that are not selected for imminent action leads to the buildup of memory traces for those items, which can be shown in facilitated working memory encoding for those items when they are selected for imminent action later. The key manipulation in this study was that we either kept the positions of the unplaced items unchanged (i.e., the "stable" condition, which allows for the buildup of memory traces for these items), or we swapped the unplaced items in the model grid while the model grid was covered (i.e., the "shuffled" condition, hampering the buildup of memory traces for these items). We compared overall VWM performance in the stable versus shuffled condition when these items were subsequently selected as targets for action, to test whether prior exposure to these unplaced items benefited their later VWM encoding. Specifically, better VWM performance may involve encoding more items per inspection (i.e., fewer inspections for encoding a certain number of items), encoding items faster (i.e., shorter inspection durations), or making fewer recall errors. Importantly, better performance in one metric (e.g., shorter inspection durations) should not be accompanied by worse performance in another (e.g., more errors), which would indicate a tradeoff rather than an overall increase in performance. If prior exposure to unplaced items benefits subsequent VWM encoding of those items, we expect better overall VWM

**Figure 1**

*Procedure and Conditions*

**A.** *Original Model*

**B.** *Model with Aperture*

**C.** *Model with Cues*

**D.** *Procedure and Conditions*

Inspection

Selection

Placement

Inspection

...

Stable Model

Shuffled Model

*Timing of Shuffle*

*Exp. 1*

Shuffle unplaced items before each inspection

*Exp. 2*

Shuffle unplaced items after each correct placement

*Exp. 3*

Shuffle uncued items before each inspection when cued items are being placed
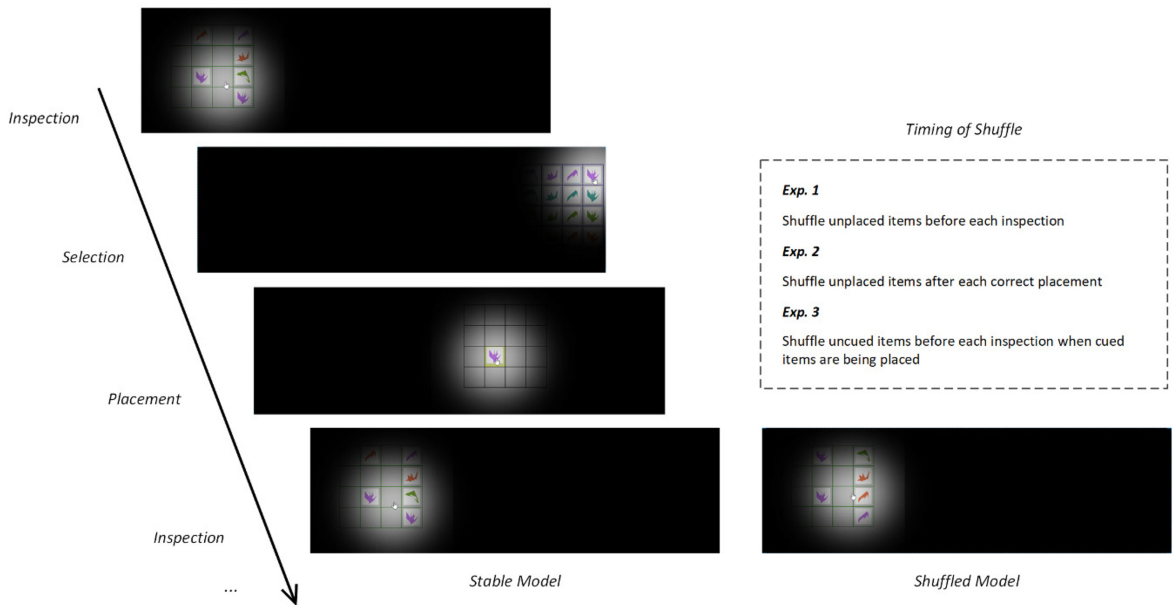
*Note.* (A) Overview of the copying task display. Participants are required to reproduce the model grid (left) by dragging items from the resources grid (right) to the correct position in the workspace (middle). (B) Experiment display with aperture, as used in Experiments 1 and 2. The aperture ensures that only the area around the cursor was visible to the participant. When the cursor was in the model grid, the model grid was visible; when the cursor was in other grids, the model grid was invisible. (C) Overview of the cued display, used in Experiment 3. Four items were cued by thick black outlines, and participants were instructed to finish these cued items first. Placing the uncued (i.e., not outlined) items was only allowed after all cued items were placed correctly. When the cursor was in the model grid, all grids were visible; when the cursor was in other grids, the model grid was covered by a gray square. (D) Example of a trial sequence and experimental conditions. Participants used their cursor to inspect the model grid, select items from the resources grid, and place items in the workspace grid. A trial was complete when all items were correctly placed in the workspace grid. In the stable condition, the model grid remained identical throughout the entire trial. In contrast, in the shuffled condition, the unplaced items in the model grid were shuffled while participants were in the resources/workspace area (see "timing of shuffle" for details). Exp. = experiment. See the online article for the color version of this figure.

performance in the stable condition than in the shuffled condition, because the shuffled condition hampers the buildup of memory traces for unplaced items.

## Experiment 1

## Method

### Participants

Forty-two participants (21 women and 21 men, $M_{age} = 27.98$, $SD = 4.54$) were recruited via the online platform Prolific (https://www.prolific.co). The sample size was based on a power analysis using G*Power 3.1 (Faul et al., 2009), yielding an estimated minimum of 30 participants for 80% power (at an α of 5%) to find an effect size of at least $f = .48$ (the smallest effect size of interest in Sahakian et al., 2023, who used a similar paradigm and similar outcome metrics). Following the same considerations (i.e., task duration, financial compensation for participants) as this previous study (Sahakian et al., 2023), we further increased the sample size to 40. Such combination of our main analysis of interest ($2 \times 2$ mixed analysis of variance [ANOVA]) and sample size ($N = 40$), according to power analysis for general ANOVA designs (Westfall, 2015), allows the detection of an effect size of Cohen's $d = 0.39$ for the main effect of model stability, given an α of 5% power and 80% power. Note that we included two participants more than planned, because the sample size was inadvertently unbalanced after 40 participants. We used the same criteria for participant recruitment as this previous study: We only included participants who (a) indicated to have normal or corrected-to-normal vision, (b) indicated to be fluent in English, (c) had an approval rate higher than 95%, and (d) had not taken part in earlier pilot versions of this experiment. We restricted the age range of participants to 18–35, because working memory performance is known to decline rapidly from the mid-thirties onwards (e.g., Salthouse, 2004, 2019), thus increasing variance in performance at the group level. The experiment complied with all ethical guidelines set out in the Declaration of Helsinki and was approved by the Ethics Committee of the Faculty of Social and Behavioral Sciences of Utrecht University. The approval is filed under number 21-0297. Participants received monetary compensation for their participation (6.25 British pounds).

### Apparatus and Stimuli

The experiment was programmed in the code editor Visual Studio Code (Version 1.75; https://code.visualstudio.com) using the JavaScript libraries jsPsych (Version 6.3.0) (de Leeuw, 2015) and Fabric.js (Version 4.3.1; https://www.fabricjs.com) and was hosted online on the web service Gorilla (https://app.gorilla.sc).

Participants completed the experiment using their own devices (laptop or desktop computers) with the use of a computer mouse. To account for the different sizes of the displays in different devices, we used a calibration procedure to equate the stimulus size prior to the formal experiment. In this calibration task, participants were asked to hold up a credit card (or any other standard sized card, commonly 8.56 cm wide) against the screen, and then resize a rectangle on the screen to match the size of the card. This procedure ensured that the light gray rectangle background (see Figure 1) was 25 cm

wide and 8.5 cm high, and that each cell of the copying task grid was 1 cm wide and 1 cm high.

A stimulus set was constructed for the experiment. The shapes of the items were randomly selected from an existing stimulus set (Arnoult, 1956) comprising 20 polygons (Figure 2). The colors of the items were randomly selected from the HSLuv (https://www.hsluv.org) color space (with 20 equidistant hues on the color wheel, and the saturation and luminance were set to 90% and 65%, respectively). Note that, because participants used different devices and performed the task under different lighting conditions, the actual luminance and hue of our stimuli inevitably varied between participants.

On each trial, a different subset of stimuli was chosen from the larger stimulus set, and positioned at random locations, so that participants were presented with a unique model grid on every trial. Accordingly, both the resources and model grids varied across trials for each participant. Specifically, to generate the resource area for each trial (Figure 1), we selected 16 items (combinations of 4 unique shapes × 4 unique colors) out of the stimulus set (400 possible items, combinations of 20 shapes × 20 colors). These items were required to encompass four distinct shapes and four distinct hues, ensuring that there was a minimum angular separation of 54° (equivalent to three steps) between any two hues. For constructing the model grid, six items were randomly selected from the resource area. To prevent multiple occurrences of identical items in the model grid (which would effectively decrease the amount of total information per model grid), we allowed for a maximum of two identical items per model grid (i.e., per trial). The occurrence of duplicate items did not differ between stable and shuffle conditions (see Supplementary Materials 2 in the online supplemental materials).

To track which part of the display participants were looking at, we used a cursor-directed aperture (Anwyl-Irvine et al., 2021); an aperture that followed the location of the cursor. To force participants to actively memorize the items in the model grid, only the part of the display within the aperture (the area around the cursor) was visible, rendering the rest of the display invisible (i.e., covered by a black overlay). The location of the aperture was recorded continually, and the events (e.g., number or duration of model inspections) were defined by the movements of the aperture across different predefined areas (i.e., model, workspace, and resources). The radius of the aperture was set to 9% of the width of the light gray experiment rectangle, and the standard deviation of the Gaussian function (for the transparency) was set to half the size of the aperture's radius. This aperture size enabled the whole model grid to be visible at once when the aperture entered the model grid, while keeping the workspace and resources grids invisible.

### Procedure

Before the formal experiment, participants received instructions and completed two practice trials; one without the overlay and aperture (as in Figure 1A) followed by one with the overlay and aperture (Figure 1B). On each trial, participants were required to reproduce the model grid by dragging items from the resources grid and dropping them onto the workspace grid. While an item was hovered over the workspace, the nearest square of the workspace was highlighted in yellow. Once participants released an item containing the correct features (shape and color) on the correct grid cell, it was pinned to

**Figure 2**

*The 20 Shapes and 20 Colors That Were Combined to Create the Stimuli in the Experiment*



*Note.* See the online article for the color version of this figure.

the center of that cell; if it was released in the wrong place, it would directly fly back to the original location in the resources grid.

In the stable condition, the positions of all items in the model grid remained unchanged throughout the entire trial, which allowed the buildup of memory traces for unplaced items. In the shuffled condition, the unplaced items (i.e., items that were not placed correctly) were shuffled every time the aperture reentered the model grid after at least one placement attempt (select or drop items) was made. The positions of items that were placed correctly remained unchanged. Because the unplaced items changed location in the shuffled condition, a potential buildup of memory traces for these items was hampered.

Participants finished 24 trials (one block) in the stable and shuffled blocks (48 trials in total). The order of these two blocks was counterbalanced between participants. Importantly, at the beginning of each block, participants were explicitly informed about the experimental condition of the upcoming trials.

### Transparency and Openness

Materials (including figures and demos) and data of the present study are openly available at the project's Open Science Framework page (https://osf.io/z75au/).

### Data Analysis

We assessed VWM encoding of items in terms of both encoding efficiency and encoding effectiveness. As proxies for VWM encoding efficiency, we analyzed (a) the number of model inspections (i.e., the number of mouse crossings toward the model grid) per correctly placed item, and (b) inspection duration (i.e., sampling time) per inspection. In addition, we analyzed (c) the number of incorrect placements (i.e., placement errors) per correctly placed item as a proxy for the effectiveness of VWM encoding. The relations between metrics are further discussed in Supplementary Materials 3 in the online supplemental materials.

For each output measure, we conducted a mixed ANOVA with model stability (stable vs. shuffled) as within-participant factor, and block order (stable blocks first vs. shuffled blocks first) as between-participant factor to account for learning effects. We

hypothesized that if exposure to prospective targets facilitates VWM encoding, we should observe fewer inspections, shorter inspection durations, or fewer errors when the position of prospective items remained constant (i.e., in the stable condition) compared to when their position changed (i.e., in the shuffled condition). We set out to investigate whether prior exposure would benefit subsequent VWM encoding at all and were initially agnostic whether such a benefit would express through (a) enhance VWM encoding efficiency, (b) encoding effectiveness, or (c) both. Since we are agnostic as to which metric would show an effect (if any), we additionally applied a Bonferroni correction (i.e., multiplying the *p* value by 3, the number of output measures) to the statistical outcomes of Experiment 1.

Notably, since we are interested in comparing copying task performance between stable and shuffled conditions we excluded from analysis: (a) all placements following the first inspection (because items can only be shuffled from the second inspection onwards), and (b) all placements after the fifth correct placement (because after this placement, there was only one item left, and a single item cannot be shuffled with itself). By doing so, we only analyzed phases of the experiment where the stable and shuffled conditions actually differed.

### Results

#### Inspections per Correctly Placed Item

The results of the 2 (model stability: stable vs. shuffled) × 2 (block order: stable blocks first vs. shuffled blocks first) ANOVA showed no significant main effects of model stability, $F(1, 40) = 0.13$, $p_{\text{Bonferroni}} = 1.00$, $\eta_p^2 < .01$, and block order, $F(1, 40) = 0.02$, $p_{\text{Bonferroni}} = 1.00$, $\eta_p^2 < .01$, on the number of inspections (Figure 3). We did find a significant interaction between model stability and block order, $F(1, 40) = 25.33$, $p_{\text{Bonferroni}} < .001$, $\eta_p^2 = .39$, which shows an improvement of performance (i.e., a reduced number of inspections) in later (vs. early) phases of the experiment. Thus, there is no evidence that the stability of items during previous exposure (when they were not yet correctly placed) influences the number of inspections for the subsequent placements of these items.

**Figure 3**
*Results of Experiment 1*



**A.** *Inspections*    **B.** *Inspection duration*    **C.** *Errors*
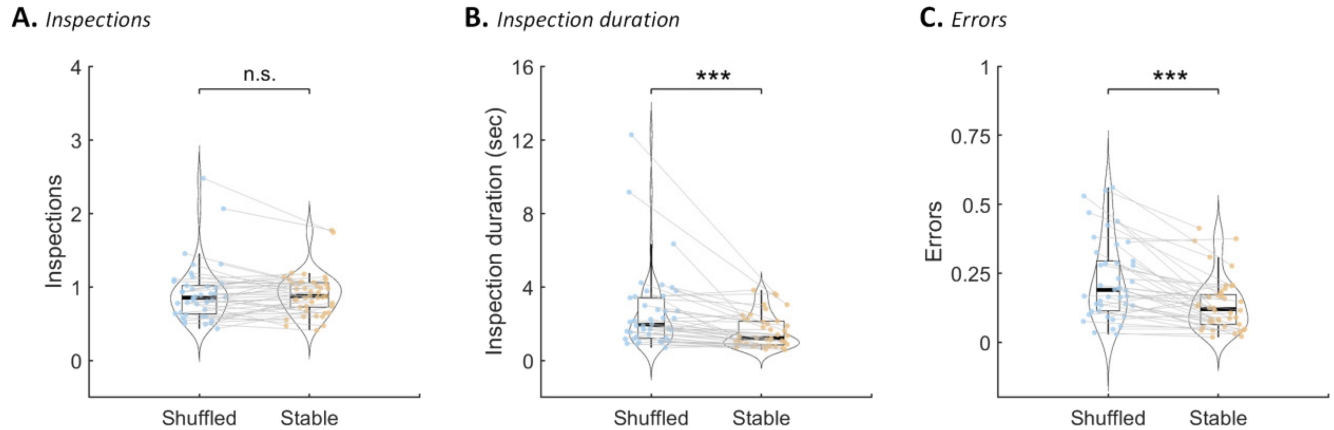
*Note.* (A) Number of model inspections per correctly placed item. (B) Inspection duration per model inspection. (C) Number of incorrect placements per correctly placed item. Box plots show the upper/lower quartiles, and the whiskers extend to the most extreme data point within 1.5 interquartile ranges above/below the upper/lower quartile. Dots show the mean data of individual participants. Curved contours show the probability density function of participant averages. n.s., $p > .05$. See the online article for the color version of this figure.
\*\*\* $p < .001$.

### Inspection Duration per Inspection

The same $2 \times 2$ ANOVA on inspection duration shows that the main effect of model stability (stable vs. shuffled) was significant, with a shorter inspection duration in stable blocks (1.59 s, $SD = 0.90$) compared to shuffled blocks (2.62 s, $SD = 2.21$), $F(1, 40) = 15.67$, $p_{\text{Bonferroni}} < .001$, $\eta_p^2 = .28$. The main effect of block order, $F(1, 40) = 0.30$, $p_{\text{Bonferroni}} = 1.00$, $\eta_p^2 = .01$, and the interaction of model stability and block order, $F(1, 40) = 3.81$, $p_{\text{Bonferroni}} = 0.174$, $\eta_p^2 = .09$, were not significant. Taken together, these results show that, given the same number of inspections per correctly placed item, participants spent less time encoding items in VWM when they had remained at the same location throughout a trial (i.e., in the stable condition) compared to when they were shuffled (i.e., in the shuffled condition).

### Errors per Correctly Placed Item

The same $2 \times 2$ ANOVA on the error rate shows that the main effect of model stability (stable vs. shuffled) was significant, with fewer placement errors in stable blocks (0.14, $SD = 0.10$) compared to shuffled blocks (0.22, $SD = 0.15$), $F(1, 40) = 32.92$, $p_{\text{Bonferroni}} < .001$, $\eta_p^2 = .45$. The main effect of block order, $F(1, 40) = 1.72$, $p_{\text{Bonferroni}} = .591$, $\eta_p^2 = .04$, and the interaction of model stability and block order were not significant, $F(1, 40) = 3.02$, $p_{\text{Bonferroni}} = .270$, $\eta_p^2 = .07$. Taken together, these results show that participants made fewer placement errors for items that were previously stable than for items that were previously shuffled.

### Interim Discussion

In Experiment 1, we found that when the model grid was stable (vs. shuffled), participants took less time inspecting the model grid, even though the number of inspections was the same. In addition, they made fewer errors when copying the items to the workspace. These results indicate that repeated exposure to items that had not been correctly placed improved subsequent VWM encoding efficiency and encoding effectiveness.

Given the goal of the study, we hypothesized that these improvements in performance were caused by prior exposure to items that were not selected for imminent action (i.e., prospective items). In this experiment, however, the shuffled items were arguably not all prospective items, but might also include items that were—in fact—selected for action. For example, participants might have attempted to place an item, but placed it incorrectly, and then went back to the model grid for further inspection; or, participants might have decided last-minute not to place an item that was initially memorized for action (e.g., because they were unsure about its location or feature), thus going back to the model grid. In these two situations, the item that was previously memorized for action already, would have been spatially swapped in the shuffled condition (along with the other unplaced items). As such, the performance cost in the shuffled condition of Experiment 1 is either caused by reduced exposure to prospective items (as hypothesized), or because items changed as participants encoded them during consecutive inspections. Since we are specifically interested in the possible effects of the exposure of items that have not been selected for action (i.e., prospective items), we should ensure to only disrupt the buildup of memory traces for items that were not actually acted upon in the shuffled condition (rather than shuffling all unplaced items). This is what we did in Experiment 2.

## Experiment 2

We conducted Experiment 2 to test whether the benefits in VWM encoding observed in Experiment 1 were solely due to prior exposure to items that were never acted upon (i.e., prospective items), rather than disruption of items that were selected but incorrectly placed. Instead of shuffling the model grid before each inspection (what we did in Experiment 1), in Experiment 2, we shuffled the model grid after each correct placement. Participants were forced

to go back to the model grid after each correct placement because they were only allowed to place one item. When participants placed an item incorrectly (or decided not to place an item) and then went back to the model grid, the model grid remained unchanged. This ensured that only those items that were shuffled had not been acted upon, and ensured that items were not shuffled when participants needed multiple glances to place it correctly. By doing so, any difference in performance between the shuffled condition and the stable condition can be attributed to the repeated exposure to items that were never acted upon. Again, if the previous exposure to prospective items facilitates VWM encoding, we would observe better VWM performance (i.e., shorter inspection duration, or fewer errors, as in Experiment 1) when prospective items were stable (vs. shuffled) compared to when they were shuffled.

## Method

### Participants

A new group of 40 participants (21 women and 19 men, $M_{age} = 28.30$, $SD = 4.50$) were recruited via Prolific, following the same recruitment criteria as in Experiment 1.

### Apparatus and Stimuli

The apparatus and stimuli in Experiment 2 were identical to those of Experiment 1. As in Experiment 1, six items were used for constructing the model grid.

### Procedure

The experimental design and procedure of Experiment 2 were similar to that of Experiment 1 except for the following changes: First, participants were allowed to place only one item in the workspace after each inspection. To achieve this, items could no longer be selected from the resources grid after a correct placement. As a result, they were forced to go back to the model grid before being able to place a new item. Second, in the shuffled condition, the model grid was shuffled after each correct placement of an item.

**Data Analysis.** The method of data analysis was identical to that of Experiment 1. As for Experiment 1, we needed to exclude data from those phases of the experiment where the stable and shuffled conditions did not differ. In this case, this entailed excluding from analysis all placements following inspection of the first item and following inspection of the last item.

## Results

### Inspections per Correctly Placed Item

The results of the 2 (model stability: stable vs. shuffled) × 2 (block order: stable blocks first vs. shuffled blocks first) ANOVA revealed no significant main effects of model stability, $F(1, 38) = 1.68$, $p = .202$, $\eta_p^2 = .04$, and block order, $F(1, 38) = 0.50$, $p = .483$, $\eta_p^2 = .01$, on the number of inspections (see Figure 4). However, there was a significant interaction between model stability and block order, $F(1, 38) = 14.79$, $p < .001$, $\eta_p^2 = .28$, which shows an improvement of performance (i.e., a reduced number of inspections) in later (vs. early) phases of the experiment. Thus, there is no evidence that the stability of items during previous exposure (when they were not selected for imminent action) influenced the number of inspections for the subsequent placements of these items.

### Inspection Duration per Inspection

The results show that the main effect of model stability (stable vs. shuffled) on inspection duration was significant, with shorter inspection duration in stable blocks (1.05 s, $SD = 0.50$) compared to shuffled blocks (1.34 s, $SD = 0.70$), $F(1, 38) = 20.68$, $p < .001$, $\eta_p^2 = .35$. The main effect of block order, $F(1, 38) = 0.54$, $p = .465$, $\eta_p^2 = .01$, was not significant. The interaction between model stability and block order was significant, $F(1, 38) = 6.87$, $p = .013$, $\eta_p^2 = .15$, which shows an improvement of performance (i.e., reduced inspection duration) in later (vs. early) phases of the experiment. Therefore, these results show that, given the same number of inspections per correctly placed item, participants spent less time encoding items in VWM when they had remained at the same location throughout a trial (i.e., in the stable condition) compared to when they were shuffled (i.e., in the shuffled condition).

### Errors per Correctly Placed Item

Results showed that the main effects of model stability, $F(1, 38) = 0.33$, $p = .569$, $\eta_p^2 = .01$, and block order, $F(1, 38) = 0.04$, $p = .853$, $\eta_p^2 < .01$, and the interaction between model stability and block order, $F(1, 38) = 3.65$, $p = .064$, $\eta_p^2 = .08$, on the amount of errors were all not significant. These results suggest that there is no evidence that the stability of prospective items during previous exposure (when they were not selected for imminent action) influenced the number of errors that participants made when subsequently placing these items.

## Interim Discussion

In Experiment 2, we found that when the unplaced items were stable (vs. shuffled), it took participants less inspection time to copy the model grid, despite similar numbers of inspections and errors. Since we ensured that the unplaced items that were shuffled were items that had not been utilized or even selected for action (i.e., prospective items), the difference between the stable and shuffled conditions can be attributed to the stability of prospective items. These results suggest that, through repeated exposure, memory traces were built up for items that were not selected for action. As a result of this, items were more efficiently encoded into VWM later.

We successfully replicated the effects on inspection durations observed in Experiment 1, but we did not replicate the effects on the number of errors. Note, however, that we cannot confirm whether the stability of prospective items in Experiment 2 affected the effectiveness of VWM encoding. This is because we allowed participants to place only one item after each inspection, which is likely to have led to floor effects in the number of errors. Therefore, it remains unclear whether the effectiveness of VWM encoding could also be affected by prior exposure to prospective items. We addressed this issue in Experiment 3.

## Experiment 3

In Experiment 3, we used a new approach that incorporated the strengths of the first two experiments into one design while also avoiding their weaknesses: First, in order to optimize the sensitivity for detecting a difference in the number of errors, we allowed participants

**Figure 4**
*Results of Experiment 2*



**A.** *Inspections*          **B.** *Inspection duration*          **C.** *Errors*
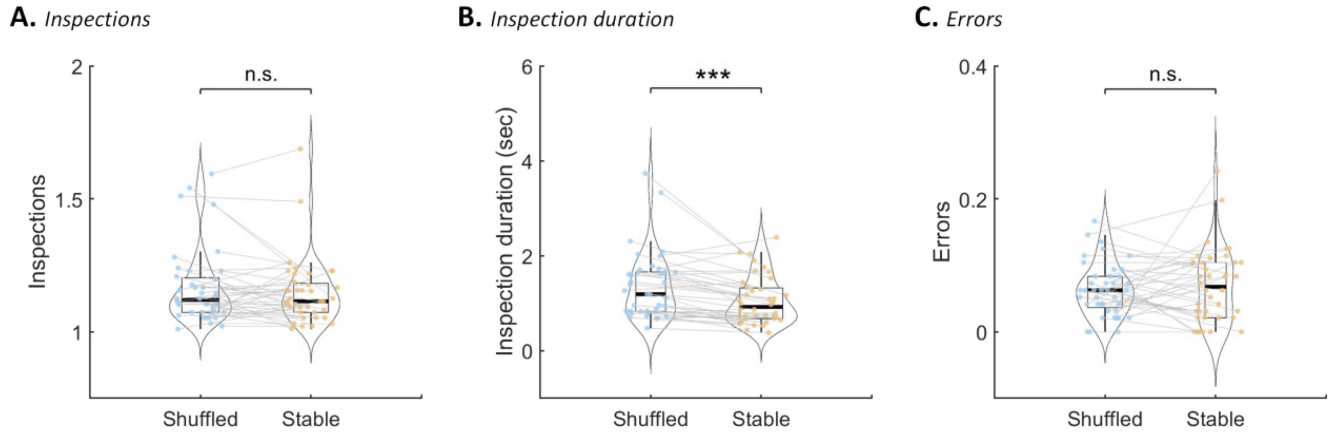
*Note.* (A) Number of inspections of the model grid per correctly placed item. (B) Inspection duration of the model grid per inspection. (C) Number of placement errors per correctly placed item. Box plots show the upper/lower quartiles, and the whiskers extend to the most extreme data point within 1.5 interquartile ranges above/below the upper/lower quartile. Dots show the mean data of individual participants. Curved contours show the probability density function of participant averages. n.s., $p > .05$. See the online article for the color version of this figure.
*** $p < .001$.

to place multiple items after each inspection (as in Experiment 1), ensuring a greater sensitivity for detecting difference in VWM encoding effectiveness between the stable and shuffled conditions. Second, we ensured that shuffled items were not previously selected for action (as in Experiment 2). To achieve this, we cued four items and instructed participants that these items needed to be placed first. Only after they finished placing these four cued items were they allowed to place the remaining four (uncued) items. Because participants were not allowed to place the uncued items before placing the cued items, we can assume that the uncued items were never selected for action when participants were placing the first four (cued) items. Thus, to manipulate the stability of prospective items, we shuffled the uncued items, while participants were placing the cued items (i.e., in the first phase of the trial). To measure whether exposure to stable prospective items facilitated VWM encoding, we then compared performance for placing uncued items (in the second phase of the trial), depending on whether these items were previously shuffled (shuffled condition) or not (stable condition).

## Method

### Participants

A new group of 40 participants (19 women and 21 men, $M_{age} = 28.55$, $SD = 4.76$) were recruited via Prolific, following the same recruitment criteria as in Experiments 1 and 2.

### Apparatus and Stimuli

The apparatus and stimuli in Experiment 3 were identical to those of Experiments 1 and 2. We divided the stimuli into cued items and uncued items. To ensure each category had an adequate number of items for measurement, we increased the total number of items used to construct the model grid from six (in Experiments 1 and 2) to eight in Experiment 3, comprising four cued items and four uncued items.

### Procedure

The experimental design and procedure of Experiment 3 were similar to those of Experiment 1, except for the following changes: First, there were four items cued by thickened black outlines (Figure 1). The cues remained on the screen throughout an entire trial. Participants were instructed to copy these four cued items first. If participants placed an uncued item correctly in this phase of the trial, the item would fly back to the resources grid as if it was incorrect. Only after participants finished placing all cued items, were they allowed to place the remaining items. Importantly, after having placed all cued items, participants had to reinspect the model grid (triggering one more swap in the shuffled condition) before being allowed to place the uncued items. Second, in the shuffled condition, we only shuffled the uncued items while participants were placing the cued items (i.e., in the first phase of a trial). When participants were placing the uncued items (in the second phase of a trial), the items in the model grid were always fixed (also in the shuffled condition). This allowed us to measure the influence of prior exposure on subsequent memory encoding, while keeping the visual stimulation identical between conditions.

Because the items that were shuffled (i.e., the uncued items) were in all likelihood never selected for action during the placements of the cued items, we did not need to incorporate any further restrictions for the swapping to occur (as in Experiment 1). Therefore, in Experiment 3, the uncued items in the model grid were shuffled every time participants went back to the model grid for inspection (in the first phase of a trial).

**Data Analysis.** The method of data analysis is similar to that of Experiments 1 and 2. Of note, we are mainly interested in participants' performance on the uncued items (i.e., the second phase of the trial), as a function of whether these uncued were shuffled or stable in the first phase of the trial (when they were not selected for imminent action). Therefore, we report the data for placing

uncued items first. We also report performance for placing cued items (first phase of the experiment), to investigate the influence of swapping (other items) on copying task performance, which might have contributed to the effects observed in Experiments 1 and 2.

## Results

### Inspections per Correctly Placed Item: Uncued Items (the Second Phase of the Trial)

The results of the 2 (model stability: stable vs. shuffled) × 2 (block order: stable blocks first vs. shuffled blocks first) ANOVA show no significant main effect of model stability, $F(1, 38) = 3.67$, $p = .063$, $\eta_p^2 = .09$, and block order, $F(1, 38) = 0.18$, $p = .677$, $\eta_p^2 < .01$, on the number of inspections, Figure 5. The interaction between model stability and block order is not

significant, $F(1, 38) = 4.04$, $p = .052$, $\eta_p^2 = .10$. Thus, there is no evidence that the stability of items during previous exposure (i.e., uncued items during the placements of cued items) influenced the number of inspections for subsequent placements of these (uncued) items.

### Inspections per Correctly Placed Item: Cued Items (the First Phase of the Trial)

The results of the 2 (model stability: stable vs. shuffled) × 2 (block order: stable blocks first vs. shuffled blocks first) ANOVA show no significant main effect of model stability, $F(1, 38) = 0.14$, $p = .710$, $\eta_p^2 < .01$, and block order, $F(1, 38) = 0.02$, $p = .882$, $\eta_p^2 < .01$, on the number of inspections. The interaction between model stability and block order, however, is significant, $F(1, 38) = 4.71$, $p = .036$, $\eta_p^2 = .11$, which reflects better performance (i.e., a reduced number of inspections) in later (vs. early)

**Figure 5**
*Results of Experiment 3*
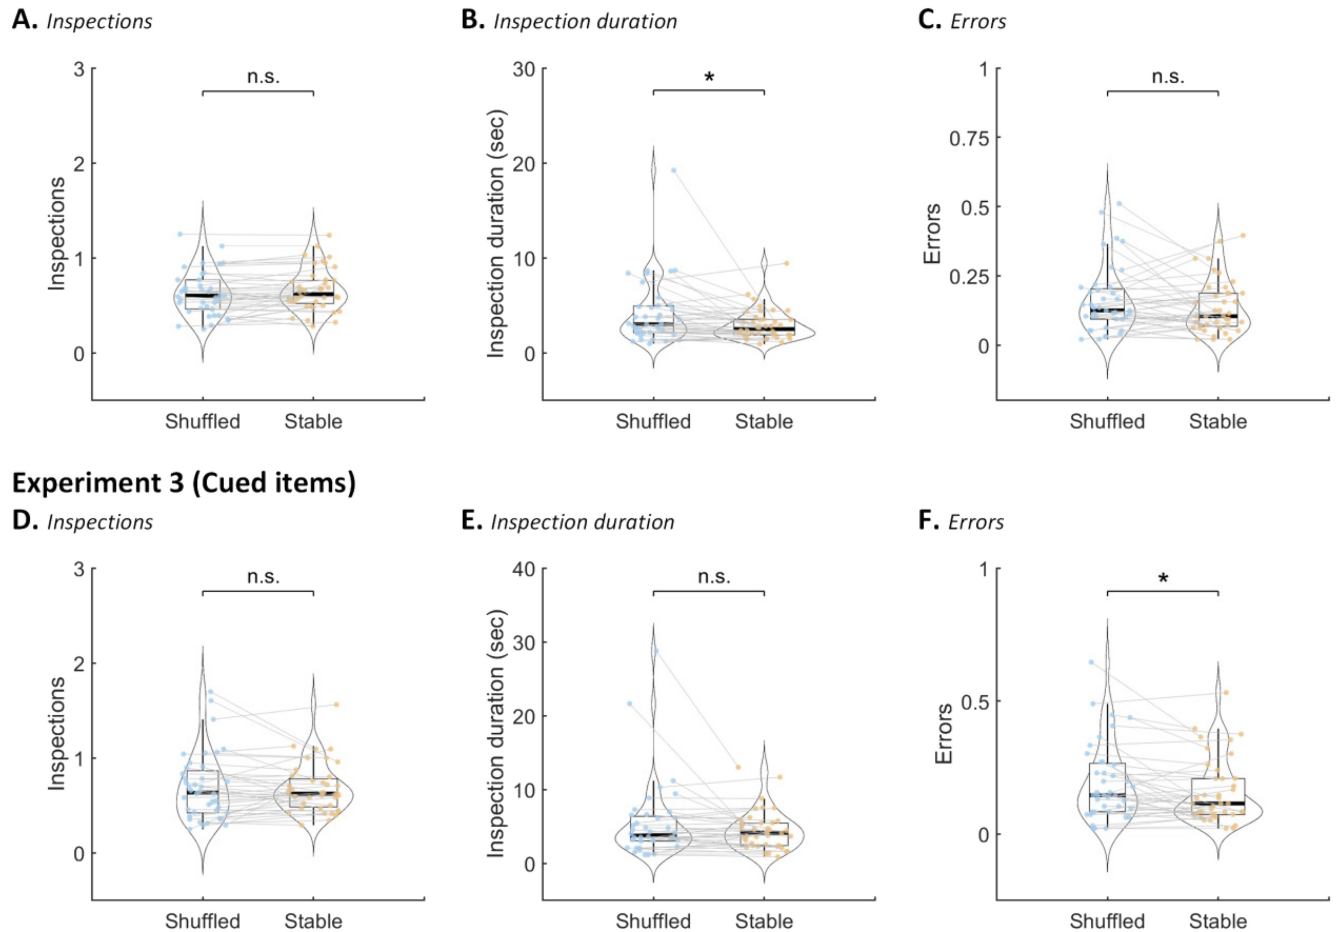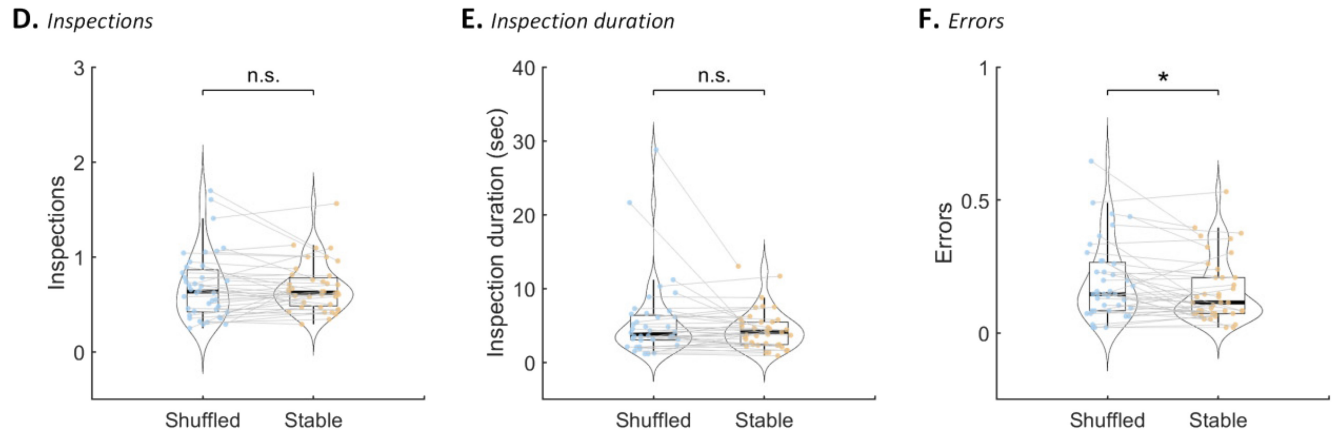


*Note.* (A) to (C) show the data for placing uncued items (second phase of the experiment). (D) to (F) show the data for placing cued items (first phase of the experiment). (A, D) Number of inspections of the model grid per correctly placed item. (B, E) Inspection duration of the model grid per inspection. (C, F) Number of erroneous placements per correctly placed item. Box plots show the upper/lower quartiles, and the whiskers extend to the most extreme data point within 1.5 interquartile ranges above/below the upper/lower quartile. Dots show the mean data of individual participants. Curved contours show the probability density function of participant averages. n.s., $p > .05$. See the online article for the color version of this figure.
* $p < .05$.

phases of the experiment. Thus, there is no evidence that the stability of prospective items (i.e., uncued items) influenced the number of inspections for placing the current targets (i.e., cued items).

### Inspection Duration per Inspection: Uncued Items (the Second Phase of the Trial)

Results show that the main effect of model stability (stable vs. shuffled) on inspection duration is significant, with shorter inspection duration in stable blocks (2.97 s, $SD = 1.63$) compared to shuffled blocks (4.14 s, $SD = 3.35$), $F(1, 38) = 5.72$, $p = .022$, $\eta_p^2 = .13$. The main effect of block order, $F(1, 38) = 1.07$, $p = .031$, $\eta_p^2 = .03$, and the interaction of model stability and block order, $F(1, 38) = 0.32$, $p = .576$, $\eta_p^2 = .01$, are not significant. These results show that, given the same number of inspections per correctly placed item, participants spent less time encoding uncued items in VWM when they had remained at the same location throughout a trial (i.e., in the stable condition) compared to when they were shuffled (i.e., in the shuffled condition) during previous exposure.

### Inspection Duration per Inspection: Cued Items (the First Phase of the Trial)

Results show that the main effects of model stability, $F(1, 38) = 2.86$, $p = .099$, $\eta_p^2 = .07$, and block order, $F(1, 38) = 0.01$, $p = .915$, $\eta_p^2 < .01$, on inspection duration were not significant. The interaction between model stability and block order was also not significant, $F(1, 38) = 0.98$, $p = .330$, $\eta_p^2 = .03$. These results suggest that there is no evidence that the stability of prospective items (i.e., uncued items) influenced inspection duration for placing the current targets (i.e., cued items).

### Errors per Correctly Placed Item: Uncued Items (the Second Phase of the Trial)

Results show that the main effects of model stability, $F(1, 38) = 1.68$, $p = .203$, $\eta_p^2 = .04$, and block order, $F(1, 38) = 0.08$, $p = .779$, $\eta_p^2 < .01$, on the number of errors, and the interaction between model stability and block order, $F(1, 38) < 0.01$, $p = .952$, $\eta_p^2 < .01$, are all not significant. These results show that there is no evidence that the stability of uncued items in previous exposure influenced the number of errors that participants made when they subsequently placed these items.

### Errors per Correctly Placed Item: Cued Items (the First Phase of the Trial)

Results show that the main effect of model stability on the amount of errors was significant, with fewer errors in the stable condition (0.16, $SD = 0.12$) compared with the shuffle condition (0.19, $SD = 0.14$), $F(1, 38) = 6.65$, $p = .014$, $\eta_p^2 = .15$. The main effect of block order, $F(1, 38) = 0.35$, $p = .558$, $\eta_p^2 = .01$, and the interaction between model stability and block order, $F(1, 38) = 3.60$, $p = .065$, $\eta_p^2 = .09$, were not significant. These results suggest that participants made more errors placing the current (cued) items when the surrounding prospective (uncued) items were being shuffled (i.e., in the shuffled condition), compared to when they were kept stable (i.e., in the stable condition).

**Table 1**
*An Overview of Conditions and Behavioral Results of Experiments 1–3*

| Experiment | Shuffled items | Timing of shuffling | No. of placements | Hypotheses for output measures | Rejections of H$_0$ | Support of H$_0$ |
|---|---|---|---|---|---|---|
| 1 | Items that are not yet correctly placed | Before each inspection | Unlimited (max. 6) | Inspections, inspection duration, and errors: H$_1$; stable < shuffle H$_0$: null results | Inspection duration and errors | Inspections |
| 2 | Items that were not yet used (either correctly or incorrectly) | After each correct placement | One placement at most after each inspection | Inspection duration and errors: H$_1$; stable < shuffle H$_0$: null results | Inspection duration | Inspections and errors |
| 3 | Items that were not yet cued for action | Before each inspection | Unlimited (max. 4) | Inspection duration and errors: H$_1$; stable < shuffle H$_0$: null results | Inspection duration | Inspections and errors |

*Note.* H = hypothesis; max. = maximum.

## Interim Discussion

In Experiment 3, we used a cue-based design to ensure that manipulated items were not selected for action, while at the same time optimizing the sensitivity for detecting a difference in the number of placement errors. We replicated the findings of Experiments 1 and 2 that when the position of prospective items (i.e., uncued items during the placement of cued items), participants took less time inspecting the model grid to place these items later. As such, even though participants knew they were not allowed to place the prospective (uncued) items before having placed all cued items, they still built up latent memory traces for those prospective (uncued) items. As in Experiment 2, we did not replicate the effects on the number of errors that we observed in Experiment 1 (Table 1). Thus, we confirmed that the effects of exposure to prospective items were mainly on encoding efficiency rather than effectiveness. Together, these findings show that prior exposure enhances subsequent VWM encoding by increasing VWM encoding efficiency, providing evidence for the buildup of memory traces for prospective items through repeated visual exposure.

## General Discussion

When encoding items into VWM for imminent action, observers are often simultaneously exposed to other items that are not selected for imminent action. Here, we asked whether observers build up memory traces of these prospective items through prior exposure. More specifically, we investigated whether prior exposure benefits subsequent encoding of items into VWM, when they are selected for action later. In this study, we addressed this question using a so-called copying task (Ballard et al., 1995; Sahakian et al., 2023; Somai et al., 2020), in which participants reproduced a six-item model grid, by placing items in an empty grid. Experiment 1 showed shorter inspection durations and fewer placement errors when the configuration of items that had not been correctly placed was stable rather than shuffled. In Experiments 2 and 3, in which we ensured that the manipulated items were not selected for action, we observed shorter inspection durations when the prospective items were stable (vs. shuffled). This shortened inspection duration was not accompanied by more inspections or more errors. Together, these results show that visual exposure to prospective items increases the efficiency (but not the effectiveness) of encoding these items into VWM when they are selected for action later.

The shorter inspection durations in the stable (vs. shuffled) condition provide evidence for the buildup of memory traces during visual exposure to prospective items. This memory trace then benefits subsequent encoding of these items into VWM. An alternative account, however, is that inspection durations were not reduced by prior exposure (in the stable condition), but increased in the shuffled condition due to interference. That is, the swapping of items might have hampered normal VWM encoding processes. Swapping prospective items might, for instance, cause proactive interference (Makovski & Jiang, 2008), where the memory of previously encoded items interferes with encoding of the current (swapped) items. Such changes in spatial configuration might increase the errors in binding specific features to certain spatial locations (Emrich & Ferber, 2012). Importantly, however, in our Experiment 3, items were no longer shuffled when participants encoded the (formerly prospective) items in VWM for imminent action. Nonetheless, under these circumstances, inspection durations were still shorter in the stable condition. This demonstrates that prior exposure was beneficial to subsequent VWM encoding, rather than swapping being detrimental to (normal) VWM encoding. Together, these results demonstrate that prior exposure to items that are not selected for imminent action expedites subsequent encoding of these items into VWM. We suggest that this facilitation effect may arise from several aspects: First, prior exposure to prospective items enables the processing of their spatial and featural information before they are selected as targets for action. Furthermore, the prior visual exposure to prospective items may enable faster binding of locations and features compared to scenarios where such bindings are not feasible (e.g., the shuffled condition).

Our findings support the hypothesis that—through repeated exposure—observers build up memory traces for items that are not yet selected for action. This runs counter to the view that currently irrelevant (i.e., prospectively relevant) items are not processed or even suppressed. Given the capacity limitations of VWM, it is essential to protect VWM content from interference. Since prospective items can potentially compete for cognitive resources with items that are selected for imminent action, it would make sense to inhibit processing of prospective items to benefit VWM processes for items that are encoded for imminent action. Contrary to this idea, analysis of the first phase of Experiment 3 (when prospective items were being shuffled or not) showed that swapping prospective items increased the number of errors for concurrently placed items. This finding shows that the prospective items are not (fully) suppressed, consistent with previous studies showing that processing of task-irrelevant surrounding items does not interfere with VWM performance (Tas et al., 2016) and may instead benefit current VWM processing (O'Donnell et al., 2018). Thus, (prospective) items surrounding the items that are encoded for imminent action may be simultaneously processed because they can contribute to current VWM processing (e.g., by providing contextual cues). The present findings indicate an additional advantage of encoding surrounding items: it benefits the subsequent encoding of these items into VWM. Future studies could further investigate this possibility by examining whether and how the task relevance of nontarget (i.e., contextual) stimuli influences VWM processing.

But how could observers build up memory traces for prospective items while concurrently encoding other items for imminent action? One possibility is that prospective items are encoded into VWM (alongside the current items) because they are relevant for the task. In this view, observers might have encoded the display as a whole in VWM (i.e., as a single complex object; Brady & Alvarez, 2015), prioritizing some individual items for imminent action, and other (prospective) items for later use (O'Donnell et al., 2018). Another possibility is that prospective items are encoded by a system that is independent of VWM, thereby bypassing the capacity limit of VWM storage. There are several possible (not mutually exclusive) mechanisms that could operate in parallel with VWM: First, observers may be able to allocate attentional resources to prospective items while concurrently processing current VWM targets (e.g., Einstein et al., 2005). In this way, the preallocation of attention may help facilitate the encoding of these items in VWM later. Second, observers might concurrently build up memory traces of these prospective items through implicit learning processes that are independent of VWM load, such as context learning (Vickery et al., 2010). Third, the repeated exposure of prospective items

might draw upon plasticity-based memory systems such as long-term memory (Heinen et al., 2022). Previous studies have indicated that several repetitions of visual stimuli suffice to engage long-term memory (e.g., Woodman et al., 2013). Therefore, prospective information may not necessarily be maintained in VWM, but rather stored in activated long-term memory (Cowan, 2019) where memory traces can accumulate at relatively low cost. All these mechanisms could benefit later VWM encoding of prospective items, while bypassing the capacity limits of the VWM system. Future work is needed to investigate which mechanism is responsible for the memory buildup of prospective items.

The present study has implications for a wide range of VWM studies: While previous studies mainly focused on VWM for items that are selected as targets for imminent goal-directed action (e.g., Zhang & Luck, 2008), our findings imply that, in VWM tasks, more items are processed than only the explicitly targeted items. Second, typical VWM studies require participants to memorize all presented stimuli, and provide a completely new set of stimuli on each trial. Although this approach is useful to investigate specific VWM processes in isolation (e.g., encoding and maintenance of novel visual input), it does not necessarily reflect how VWM operates in the more complex and often stable visual environments that we encounter in daily life. Therefore, the present study contributes to our understanding of how VWM might operate in daily life, where visual information remains externally available, and where the amount of task-relevant information might exceed the capacity limits of VWM (thus requiring multiple consecutive instances of VWM encoding). Third, our results support the idea that prior visual experience can enhance VWM processing of visual stimuli, which is in line with previous work (Blalock, 2015; Calmels et al., 2012; Jackson & Raymond, 2008; Lorenc et al., 2014; Moore et al., 2022; Ngiam et al., 2019; Scolari et al., 2008). The present work extends this body of literature, by showing that the benefit of prior exposure on VWM processes also occurs when observers are exposed to items that are not selected for imminent action.

The present work can be placed in a conceptual framework that describes the tradeoff between sampling information from the external world, and retaining information internally using VWM. That is, because of the limited capacity of VWM (Cowan, 2005, 2010; Miller, 1956), observers sometimes capitalize on the availability of information in the external world (using it as external memory). This strategy decreases the burden on internal VWM storage (Sahakian et al., 2023; Somai et al., 2020), thereby achieving cognitive offloading (Risko & Gilbert, 2016). According to these views, a tradeoff is made between the cost of internal VWM storage (e.g., attentional resources) and the costs of (re-)sampling the external world (e.g., the cost of eye movements or locomotion) (Van der Stigchel, 2020). Previous studies revealed that increasing the accessibility (Draschkow et al., 2021; Sahakian et al., 2023) and the stability of information in the external world (e.g., Chota et al., 2023) cause observers to rely less on internal VWM storage, since, in those situations, the cost of external sampling is decreased (e.g., when observers have faster access to reinspect visual stimuli). Our study reveals how an aspect of the external world—its stability—can reduce the cost of VWM storage; a stable outside world allows repeated exposure to prospective items, thereby rendering internal VWM storage more efficient. Although prior exposure reduced the cost of internal VWM storage (by reducing encoding time), this did not cause participants to rely more on internal VWM storage;

that is, participants did not sample the external world (i.e., the model grid) less frequently when a stable environment allowed for faster VWM encoding. One tentative explanation for this is that observers failed to monitor (subtle) differences in the cost of internal VWM storage, so that the tradeoff between internal VWM storage and external sampling depended mainly on the (perceived) cost of external sampling.

In conclusion, the present study shows that when sampling a subset of items for imminent goal-directed behavior, observers also form latent memory traces for the items that are not selected for imminent action. These memory traces subsequently increase encoding efficiency, by reducing the time needed to encode these items into VWM later. Together, these findings show that the human brain extracts more information from individual glances than what is actively selected for imminent behavior, thereby benefiting prospective VWM-guided behavior.

## Constraints on Generality

The stimuli of the present study included colored polygons. Due to the constraints of the materials we used, our conclusions are restricted to visual perception and cognition. However, when placing the findings in the fields of visual perception and cognition, we have no reason to believe that the results depend heavily on the (visual) feature characteristics. That being said, some variables might be potential moderators of the findings. For example, when the observer's cognitive load is extremely high, the benefits of prior visual exposure to prospective items on subsequent working memory encoding may be attenuated, or even eliminated. Besides, it remains to be seen whether our findings generalize to other (nonvisual) features (e.g., auditory processing).

## References

Allport, A. (1989). Visual attention. In M. I. Posner (Ed.), *Foundations of cognitive sciences* (pp. 631–682). MIT Press.

Anwyl-Irvine, A. L., Armstrong, T., & Dalmaijer, E. S. (2021). Mouseview.js: Reliable and valid attention tracking in web-based experiments using a cursor-directed aperture. *Behavior Research Methods*, *54*(4), 1663–1687. https://doi.org/10.3758/s13428-021-01703-5

Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(3), 580–584. https://doi.org/10.1037/a0027885

Arnoult, M. D. (1956). Familiarity and recognition of nonsense shapes. *Journal of Experimental Psychology*, *51*(4), 269–276. https://doi.org/10.1037/h0047772

Ballard, D. H., Hayhoe, M. M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*(1), 66–80. https://doi.org/10.1162/jocn.1995.7.1.66

Blalock, L. D. (2015). Stimulus familiarity improves consolidation of visual working memory representations. *Attention, Perception, & Psychophysics*, *77*(4), 1143–1158. https://doi.org/10.3758/s13414-014-0823-z

Boduroglu, A., & Shah, P. (2014). Configural representations in spatial working memory. *Visual Cognition*, *22*(1), 102–124. https://doi.org/10.1080/13506285.2013.875499

Brady, T. F., & Alvarez, G. A. (2015). No evidence for a fixed object limit in working memory: Spatial ensemble representations inflate estimates of working memory capacity for complex objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 921–929. https://doi.org/10.1037/xlm0000075

Brady, T. F., Robinson, M. M., & Williams, J. R. (2024). Noisy and hierarchical visual memory across timescales. *Nature Reviews Psychology*, *3*(3), 147–163. https://doi.org/10.1038/s44159-024-00276-2

Calmels, C., Foutren, M., & Stam, C. J. (2012). Beta functional connectivity modulation during the maintenance of motion information in working memory: Importance of the familiarity of the visual context. *Neuroscience*, *212*, 49–58. https://doi.org/10.1016/j.neuroscience.2012.03.045

Chen, D., Yee Eng, H., & Jiang, Y. (2006). Visual working memory for trained and novel polygons. *Visual Cognition*, *14*(1), 37–54. https://doi.org/10.1080/13506280544000282

Chota, S., Gayet, S., Kenemans, J. L., Olivers, C. N., & Van der Stigchel, S. (2023). A matter of availability: Sharper tuning for memorized than for perceived stimulus features. *Cerebral Cortex*, *33*(12), 7608–7618. https://doi.org/10.1093/cercor/bhad064

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*(1), 28–71. https://doi.org/10.1006/cogp.1998.0681

Cowan, N. (2005). *Working memory capacity*. Psychology Press.

Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, *19*(1), 51–57. https://doi.org/10.1177/0963721409359277

Cowan, N. (2019). Short-term memory based on activated long-term memory: A review in response to Norris (2017). *Psychological Bulletin*, *145*(8), 822–847. https://doi.org/10.1037/bul0000199

de Leeuw, J. R. (2015). Jspsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, *47*(1), 1–12. https://doi.org/10.3758/s13428-014-0458-y

Draschkow, D., Kallmayer, M., & Nobre, A. C. (2021). When natural behavior engages working memory. *Current Biology*, *31*(4), 869–874.e5. https://doi.org/10.1016/j.cub.2020.11.013

Einstein, G. O., McDaniel, M. A., Thomas, R., Mayfield, S., Shank, H., Morrisette, N., & Breneiser, J. (2005). Multiple processes in prospective memory retrieval: Factors determining monitoring versus spontaneous retrieval. *Journal of Experimental Psychology: General*, *134*(3), 327–342. https://doi.org/10.1037/0096-3445.134.3.327

Emrich, S. M., & Ferber, S. (2012). Competition increases binding errors in visual working memory. *Journal of Vision*, *12*(4), Article 12. https://doi.org/10.1167/12.4.12

Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*(4), 1149–1160. https://doi.org/10.3758/BRM.41.4.1149

Feldmann-Wüstefeld, T., & Vogel, E. K. (2019). Neural evidence for the contribution of active suppression during working memory filtering. *Cerebral Cortex*, *29*(2), 529–543. https://doi.org/10.1093/cercor/bhx336

Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, *99*(24), 15822–15826. https://doi.org/10.1073/pnas.232472899

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct evidence for active suppression of salient-but-irrelevant sensory inputs. *Psychological Science*, *26*(11), 1740–1750. https://doi.org/10.1177/0956797615597913

Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*(2), 129–135. https://doi.org/10.1016/j.tics.2011.11.014

Heinen, K. T., Kenemans, J. L., & van der Stigchel, S. (2022). Recruitment of a long-term memory supporting neural network during repeated maintenance of a multi-item abstract visual image in working memory. *Scientific Reports*, *12*(1), Article 575. https://doi.org/10.1038/s41598-021-04384-4

Hollingworth, A. (2007). Object-position binding in visual memory for natural scenes and object arrays. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(1), 31–47. https://doi.org/10.1037/0096-1523.33.1.31

Jackson, M. C., & Raymond, J. E. (2008). Familiarity enhances visual working memory for faces. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(3), 556–568. https://doi.org/10.1037/0096-1523.34.3.556

Kristjánsson, Á, & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, *72*(1), 5–18. https://doi.org/10.3758/APP.72.1.5

Krueger, L. E. (1975). Familiarity effects in visual information processing. *Psychological Bulletin*, *82*(6), 949–974. https://doi.org/10.1037/0033-2909.82.6.949

Kuo, B. C., Stokes, M. G., & Nobre, A. C. (2012). Attention modulates maintenance of representations in visual short-term memory. *Journal of Cognitive Neuroscience*, *24*(1), 51–60. https://doi.org/10.1162/jocn_a_00087

Lewis-Peacock, J. A., & Norman, K. A. (2014). Competition between items in working memory leads to forgetting. *Nature Communications*, *5*(1), Article 5768. https://doi.org/10.1038/ncomms6768

Lorenc, E. S., Pratte, M. S., Angeloni, C. F., & Tong, F. (2014). Expertise for upright faces improves the precision but not the capacity of visual memory. *Attention, Perception, & Psychophysics*, *76*(7), 1975–1984. https://doi.org/10.3758/s13414-014-0653-z

Makovski, T., & Jiang, Y. V. (2008). Proactive interference from items previously stored in visual working memory. *Memory & Cognition*, *36*(1), 43–52. https://doi.org/10.3758/MC.36.1.43

Miller, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81–97. https://doi.org/10.1037/h0043158

Moore, K. N., Lampinen, J. M., Adams, E. J., Nesmith, B. L., & Burch, P. (2022). Prior experience with target encounter affects attention allocation and prospective memory performance. *Cognitive Research: Principles and Implications*, *7*(1), Article 37. https://doi.org/10.1186/s41235-022-00385-7

Ngiam, W. X., Khaw, K. L., Holcombe, A. O., & Goodbourn, P. T. (2019). Visual working memory for letters varies with familiarity but not complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(10), 1761–1775. https://doi.org/10.1037/xlm0000682

Oberauer, K., & Lin, H.-Y. (2017). An interference model of visual working memory. *Psychological Review*, *124*(1), 21–59. https://doi.org/10.1037/rev0000044

O'Donnell, R. E., Clement, A., & Brockmole, J. R. (2018). Semantic and functional relationships among objects increase the capacity of visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(7), 1151–1158. https://doi.org/10.1037/xlm0000508

Reder, L. M., Liu, X. L., Keinath, A., & Popov, V. (2016). Building knowledge requires bricks, not sand: The critical role of familiar constituents in learning. *Psychonomic Bulletin & Review*, *23*(1), 271–277. https://doi.org/10.3758/s13423-015-0889-1

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, *20*(9), 676–688. https://doi.org/10.1016/j.tics.2016.07.002

Sahakian, A., Gayet, S., Paffen, C. L., & Van der Stigchel, S. (2023). Mountains of memory in a sea of uncertainty: Sampling the external world despite useful information in visual working memory. *Cognition*, *234*, Article 105381. https://doi.org/10.1016/j.cognition.2023.105381

Salthouse, T. A. (2004). What and when of cognitive aging. *Current Directions in Psychological Science*, *13*(4), 140–144. https://doi.org/10.1111/j.0963-7214.2004.00293.x

Salthouse, T. A. (2019). Trajectories of normal cognitive aging. *Psychology and Aging*, *34*(1), 17–24. https://doi.org/10.1037/pag0000288

Scolari, M., Vogel, E. K., & Awh, E. (2008). Perceptual expertise enhances the resolution but not the number of representations in working memory. *Psychonomic Bulletin & Review*, *15*(1), 215–222. https://doi.org/10.3758/PBR.15.1.215

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*(2), 207–214. https://doi.org/10.1111/j.1467-9280.2009.02276.x

Somai, R. S., Schut, M. J., & Van der Stigchel, S. (2020). Evidence for the world as an external memory: A trade-off between internal and external visual memory storage. *Cortex*, *122*, 108–114. https://doi.org/10.1016/j.cortex.2018.12.017

Tas, A. C., Luck, S. J., & Hollingworth, A. (2016). The relationship between visual attention and visual working memory encoding: A dissociation between covert and overt orienting. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(8), 1121–1138. https://doi.org/10.1037/xhp0000212

Umemoto, A., Scolari, M., Vogel, E. K., & Awh, E. (2010). Statistical learning induces discrete shifts in the allocation of working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(6), 1419–1429. https://doi.org/10.1037/a0019324

Van der Stigchel, S. (2020). An embodied account of visual working memory. *Visual Cognition*, *28*(5–8), 414–419. https://doi.org/10.1080/13506285.2020.1742827

van Ede, F., & Nobre, A. C. (2023). Turning attention inside out: How working memory serves behavior. *Annual Review of Psychology*, *74*(1), 137–165. https://doi.org/10.1146/annurev-psych-021422-041757

Vickery, T. J., Sussman, R. S., & Jiang, Y. V. (2010). Spatial context learning survives interference from working memory load. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(6), 1358–1371. https://doi.org/10.1037/a0020558

Wang, B., & Theeuwes, J. (2018). Statistical regularities modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(1), 13–17. https://doi.org/10.1037/xhp0000472

Westfall, J. (2015). *PANGEA: Power analysis for general ANOVA designs*. [Unpublished manuscript]. https://github.com/jake-westfall/pangea

Woodman, G. F., Carlisle, N. B., & Reinhart, R. M. (2013). Where do we store the memory representations that guide attention? *Journal of Vision*, *13*(3), Article 1. https://doi.org/10.1167/13.3.1

Xu, L., Gayet, S., Sahakian, A., Gottlieb, J., Van der Stigchel, S., & Paffen, C. (2024, April 30). *Trade-offs between visual sampling and memory in stable and changing worlds*. PsyArXiv. https://doi.org/10.31234/osf.io/e8wbs

Zanto, T. P., & Gazzaley, A. (2009). Neural suppression of irrelevant information underlies optimal working memory performance. *Journal of Neuroscience*, *29*(10), 3059–3066. https://doi.org/10.1523/JNEUROSCI.4621-08.2009

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860