Journal of Experimental Psychology: Learning, Memory, and Cognition

Action Consequences Guide the Use of Visual Working Memory

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

Online First Publication, February 8, 2024. https://dx.doi.org/10.1037/xlm0001326

AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: Learning, Memory, and Cognition

ISSN: 0278-7393

https://doi.org/10.1037/xlm0001326

Action Consequences Guide the Use of Visual Working Memory

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

Visual working memory (VWM) is a store for temporary maintenance of visual information. It is often disregarded, though, that information is typically stored to enable actions. Therefore, the context of these actions is of great importance for how VWM is used. Here, we questioned whether the severity of the consequence of an action might affect how well information is memorized, and how cautiously it is utilized. We employed an (online) copying task, in which participants reproduced an example display comprised of six items in a grid, using a pool of items. Crucially, we manipulated the severity of penalties: participants had to wait 5 (high) or 0.5 (low error cost) s after an error. Additionally, we manipulated the accessibility of task-relevant information (a well-studied manipulation in this paradigm): participants had to wait 5 (high) or 0.5 (low sampling cost) s to inspect the example. Our results show that with higher error cost the number of inspections remained comparable, but the number of errors decreased. Furthermore, they show that with higher sampling cost the number of inspections halved, and the number of errors increased. Thus, more severe action consequences increase the reluctance to act on uncertain information in VWM, but do not lead to more attempts to store information in VWM. We conclude that, in contrast to the effect of the accessibility of information, action consequences do not affect how well information is memorized, but affect how cautiously this stored information is utilized.

Keywords: visual working memory, action consequences, copying task, naturalistic task contexts

Supplemental materials: https://doi.org/10.1037/xlm0001326.supp

We humans, such as many other species, make extensive use of our sight to guide our actions. The sight has enabled us to successfully approach food and to avoid predators, lest we would become food. Therefore, vision has become one of our most important senses for survival. Yet, our sight has limitations as well: we cannot—in one glance—perceive our complete surroundings. Luckily, we can turn our head and shift our gaze with great precision to see more of our environment. With every shift of gaze and attention toward new information, though, we may also lose sight of relevant visual information currently in view. To deal with this limitation, we have a cognitive system in place, called visual working memory (VWM), that is able to temporarily store visual information from our environment (Baddeley & Hitch, 1974). To illustrate, imagine you are about to pick some berries from a bush, but before doing so check behind your back to ensure there is no predator in sight. It would be very useful to maintain a visual image of the bush in your mind's eye as you look out for the predator (to avoid having to start your search from scratch, each time you look over your shoulder). Having a memory of the appearance and locations of the berries will save time and effort in finding them, and thus enable easy picking. VWM takes care of this: it can maintain relevant visual information to enable future actions. Shortly maintaining action-relevant visual features around us is therefore commonplace in daily life. While many aspects of VWM have been studied in the past five decades, one aspect has received relatively little attention: namely the potential consequences of the action. To illustrate, imagine not picking berries, but hunting mushrooms. A novice mushroom hunter will likely have a guidebook to counsel when in doubt. The mushroom hunter must ensure that the mushroom in their hand looks identical to the one they have seen in the book, or else it might be their last mushroom hunt. Compare that to the situation of a (less adventurous) person who must find an exotic mushroom in the supermarket, of which they received a picture on their phone. The worst that can happen there is a ruined meal. The central question here is whether action-relevant information is memorized better or used more cautiously when the consequences of the actions are more severe.

In the last few decades, many studies have focused on the format in which information is stored in VWM (Luck & Vogel, 1997;

Andre Sahakian https://orcid.org/0000-0003-0106-1182
This study was not preregistered. All data, data processing scripts, and statistical analyses are uploaded on the Open Science Framework (OSF) platform and are accessible via the link: https://osf.io/e6vpx/. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (Grant Agreement 863732). This project was supported by a Veni grant from the Netherlands Organisation for Scientific Research (NWO; Grant Vl.Veni.191G.085) to Surya Gayet. The authors declare no competing financial interests.

Andre Sahakian served as lead for data curation, formal analysis, visualization, and writing–original draft. Andre Sahakian, Surya Gayet, Chris L. E. Paffen, and Stefan Van der Stigchel contributed equally to conceptualization, methodology, investigation, and writing–review and editing. Surya Gayet and Stefan Van der Stigchel contributed equally to funding acquisition. Surya Gayet, Chris L. E. Paffen, and Stefan Van der Stigchel contributed equally to supervision.

Correspondence concerning this article should be addressed to Andre Sahakian, Experimental Psychology, Helmholtz Institute, Utrecht University, Heidelberglaan 1, 3584 CS Utrecht, The Netherlands. Email: a.sahakian@uu.nl

Ma et al., 2014; Schurgin et al., 2020). Fewer studies have investigated the influence of the context in which information is put into memory and subsequently utilized from memory. In this latter category, Ballard et al. (1995), Draschkow et al. (2021), Sahakian et al. (2023), and Somai et al. (2020) have investigated how the accessibility of information determines how much of what is memorized is put to use. These studies showed that when accessing relevant information is relatively easy, only a little information from VWM is used before referring to the available information. However, when access to information is made more difficult, information storage increases. Recently, we showed that while little information was utilized from VWM, some information in VWM remained unused (Sahakian et al., 2023). Moreover, we showed that the increase in the amount of information utilized (when access became harder) is a consequence of both memorizing more information, and utilizing a larger portion of the (otherwise residual unused) information in VWM. The latter strategy can be interpreted as a more liberal tendency to try out actions, even if the actions are based on uncertain information. Indeed, relying on uncertain information in memory can be an efficient strategy when retrieving information from the external world is hard. But there are limitations to such a strategy: actions based on weak memories will inevitably be wrong more often. In our previous work, incorrect actions had negligible consequences—incorrect actions were immediately and automatically undone. Negligible consequences make acting upon weak memories (i.e., educated guessing rather than confident actions) an attractive strategy. In many situations (such as hunting mushrooms in a forest), however, acting upon weak memories is not desirable. In situations where one cannot afford to err, memory-guided behavior will likely be adapted to minimize errors. But how is memory-guided behavior adapted? Our aim in the current study was to investigate what strategy changes occur when adverse consequences are tied to incorrect actions.

Many studies have confirmed that working memory performance can be improved when there is an incentive to do so. For instance, visual information is memorized better if it is more likely to be probed (Gorgoraptis et al., 2011; Klyszejko et al., 2014). Likewise, studies have shown that visual information which is associated with higher rewards receives more cognitive resources and is memorized better (Cho et al., 2022; Klyszejko et al., 2014; van den Berg & Ma, 2018; but see van den Berg et al., 2023). Thus, the consequences of actions that are based on VWM content (e.g., correctly reporting the memorized item) seem to influence how well information is memorized or recalled. Yet, to the best of our knowledge, no studies have investigated the effects of the consequences of actions in the context of naturalistic VWM tasks. Specifically, we mean VWM tasks in which relevant visual information remains available for reinspection, and the information is memorized to enable meaningful actions (not merely to be recalled on request). In the present study, we manipulate the consequences of actions, to investigate how the cost of making an error affects VWM-guided behavior. To do so, we build on our previous study (Sahakian et al., 2023), which had proved fruitful in revealing strategies of memory use in different contexts. The paradigm in question is a copying task: a task in which an arrangement of items (typically called the "Model") must be recreated elsewhere using a pool of items. Critically, the arrangement of items remains available for inspection throughout the task. Participants recreate the "Model" in an initially empty area called the "Workspace," by picking up items from the "Resources" (a pool of items) and placing the items in the correct position in the "Workspace" until the "Model" is perfectly copied. Tracking for how long and how often participants inspect the "Model," or how many correct and incorrect item placements follow each inspection, gives insights into to use of VWM. In our previous study, we manipulated the cost to inspect the Model. The cost could be low (i.e., easy to reach the Model for inspection) or high (i.e., hard to reach the Model for inspection). We observed several strategy changes between these conditions. When the accessibility of relevant information was worse, participants (re)inspected it less often, but for a longer duration and they were able to correctly apply more information after each inspection. These findings were robust within our study (Sahakian et al., 2023) and in the three other studies with a comparable paradigm and manipulation (Ballard et al., 1995; Draschkow et al., 2021; Somai et al., 2020). Our study also revealed that when the accessibility of information was worse, participants made more errors, which we interpreted as an increased tendency to use relatively uncertain VWM content (e.g., deciding to place an item in the workspace, despite doubts about the item's identity or location). Our key question is whether this tendency to place uncertain VWM content will persist when the consequence of making errors becomes more severe. The well-established paradigm of the copying task allows us to answer these questions.

In the current study, we were mainly interested in strategy changes in the use of VWM caused by more severe consequences for incorrect actions. To this aim, we employed a copying task paradigm hosted on the internet. The primary manipulation was that of the severity of the punishment for incorrect actions: a short or long wait time after an erroneous placement. The secondary manipulation was that of the accessibility of the "Model": a short or long wait time before the Model can be inspected. We included the accessibility manipulation as a reference, since we know from previous studies what effects to expect. Our main outcome measures of interest were the amount of information loaded up into VWM with each inspection, and the amount of information used from VWM with each inspection. Based on findings from previous studies, we expected that the amount of information loaded up into VWM and used from VWM would be larger when accessibility of the "Model" was worse. Regarding the severity of punishments for incorrect actions, we expected that more information would be loaded up into VWM when punishments were severe to minimize the risk of making an error. Furthermore, we expected that less information (i.e., only very certain information) would be used from VWM when punishments were severe, again to minimize the risk of errors. To preface the main findings of the current study: (a) We replicate the findings of previous studies regarding the manipulation of accessibility of information: there was more information loaded up into VWM and more information used from VWM when the accessibility of the "Model" was worse. (b) We find, however, no evidence that more information was loaded into VWM when the severity of punishments for errors was higher. (c) We show that there was indeed less information used from VWM when the severity of punishments for errors was higher. In sum, it seems that when the consequences of incorrect actions become more severe, humans do not necessarily try to memorize more (or better) but only become more cautious in what memories they rely on for action. Evidently, VWM-guided behavior in naturalistic contexts is affected by various factors in various ways. Put in a different way: humans adapt the use of VWM in various ways to best suit the current demands.

Method

Participants

We recruited participants via Prolific (www.prolific.co), a platform for recruiting participants for online studies. Using Prolific's screening tool, we made our study available to participants who (a) had normal or corrected-to-normal vision, (b) were fluent in English, (c) had an approval rate higher of at least 95%, and (d) had not taken part in studies (with similar tasks) we had run via Prolific. We aimed to include 25 participants per between-observer condition (there were four experimental conditions), based on a power analysis using data from a previous copying task study (see Supplemental Material 1 in the online supplemental material). As participants conducted the online study in parallel and allocation to experimental conditions was randomized, we ended up including a few more participants than 25 in some conditions. A total of 110 participants completed the study. We excluded six of them, who indicated that they had used an aid to complete the task (e.g., made a photograph of the stimuli they had to remember). In total, we included 104 participants for the formal analyses. The experiment was designed and conducted in compliance with the Declaration of Helsinki, and it was approved by the Ethics Committee of the Faculty of Social and Behavioural Sciences of Utrecht University. The approval is filed under 21-0297. All participants gave informed consent before beginning the online study. Upon successful completion of the experiment, participants were rewarded with 3.75 GBP.

Apparatus and Stimuli

The experiment was scripted using the JavaScript libraries jsPsych (Version 7.3.0; de Leeuw, 2015) and Fabric.js (Version 5.2.4; fabricjs.com). We used the web service Gorilla to build and host the online experiment (Anwyl-Irvine et al., 2020).

We enabled the option on Prolific to make the study only available to participants on a desktop device (i.e., a laptop or a personal computer). Still, displays of varying types and sizes will most likely have been used. We tried to keep the size of the experiment display constant across participants by implementing a calibration procedure before the task: participants held a standard-sized (credit) card against their monitor and resized a displayed rectangle to match the card. If the instructions were followed correctly, the experiment display would be contained in a light gray rectangle 25 cm wide × 8.5 cm high, on a white background (see Figure 2). After size calibration, the items were 1 cm × 1 cm in size.

The items used here were the same as used in previous experiments in our lab (Sahakian et al., 2023). The 20 shapes of (the polygons in) the items were adapted from stimuli used by Arnoult (1956). The 20 colors were selected from the HSLuv color space (https://www.hsluv.org). Specifically, we selected 20 perceptually equidistant hues (starting at a hue value of 3.59°, and adding 17.95° to get the next hue value) on the color wheel, with the saturation set to 89.1% and luminance to 64.35% (see Figure 1). We assume that there was a range of different monitor types, screen configurations, and room lighting conditions across participants. Thus, each participant will probably have been presented with slightly different colors.

Procedure

After giving informed consent and going through instructions and practice trials, participants started the main task, which consisted of 24 experimental trials. Participants took about a minute to complete each trial. The goal in every trial was to recreate an arrangement of six items in a 4 × 4 grid, called the Model grid. Participants did this

Figure 1

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

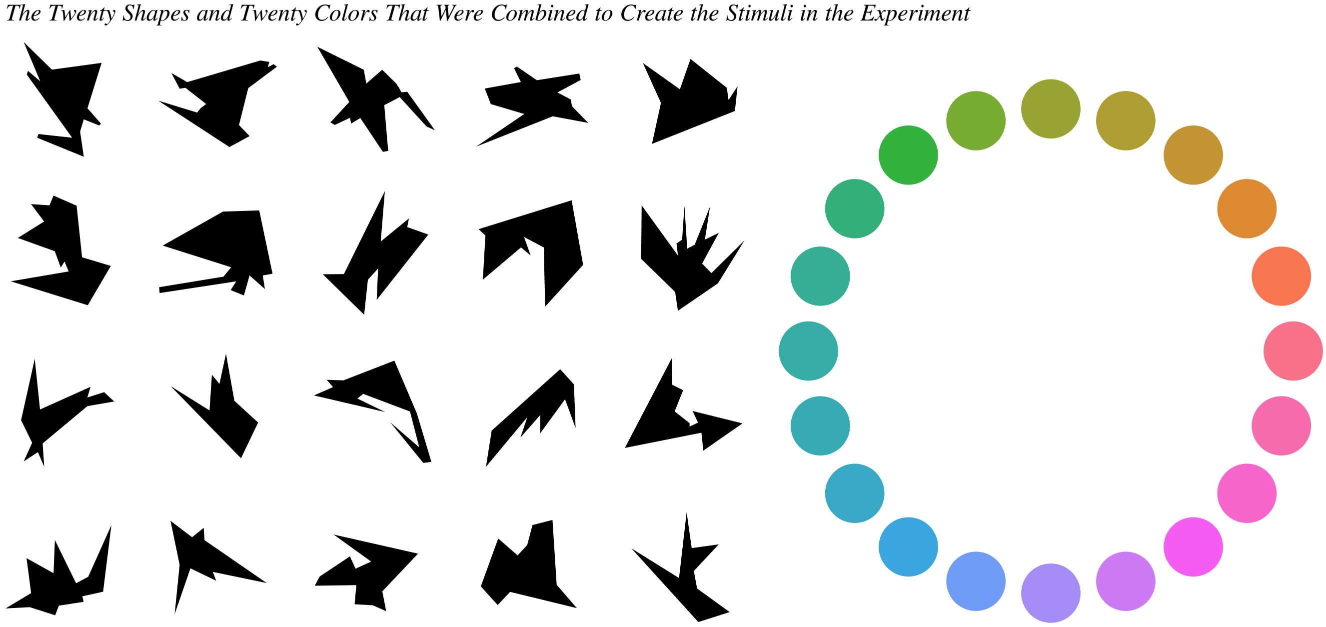

Note. Given 20 shapes and 20 colors, we could create 400 unique stimuli. For each trial, a random selection (without replacement) of four shapes and four colors was used to create 16 unique stimuli (with the additional restriction that no two neighboring colors were ever selected). From this pool of 16 stimuli, six were randomly selected (with replacement), and randomly positioned in the Model grid for each trial. See the online article for the color version of this figure.

Figure 2
Overview of a Trial

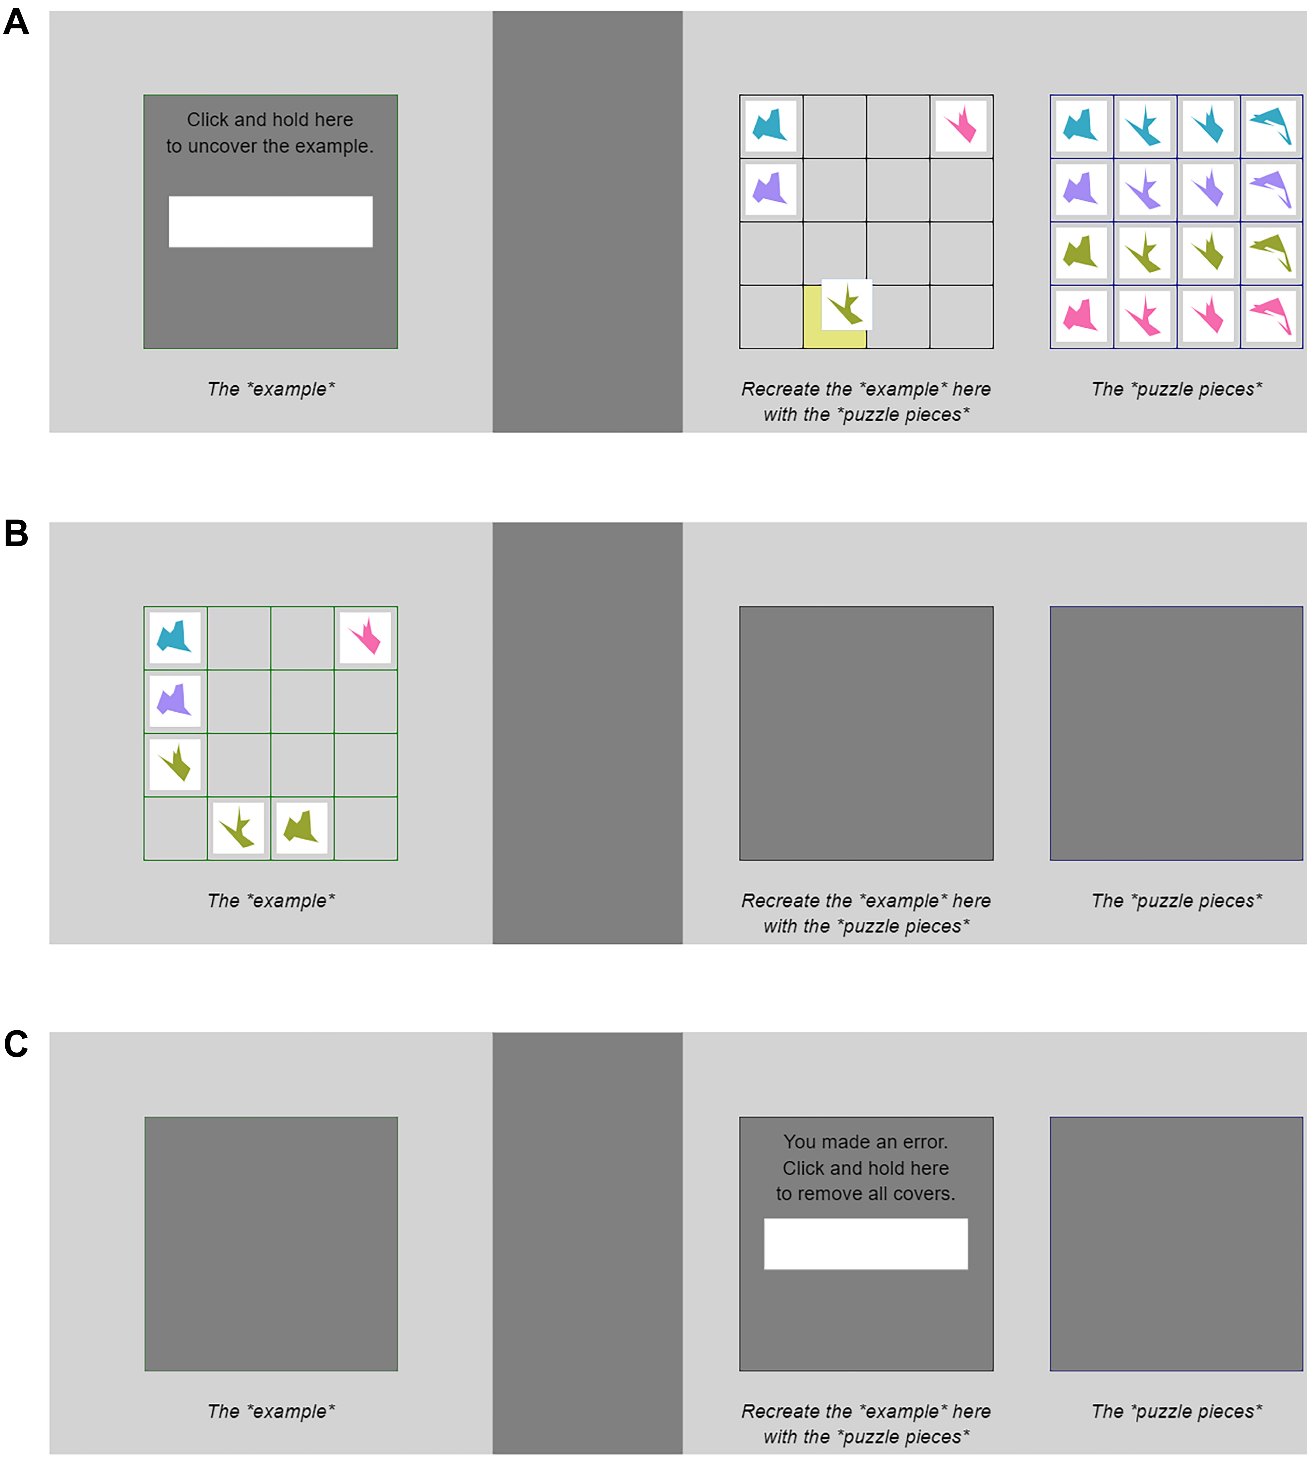

Note. There were three relevant 4 × 4 grids. From left to right they are the Model, Workspace, and Resource grid. The Model needs to be recreated in the Workspace using the items from the Resource. Panel A shows the task view when items are being placed from the Resource in the Workspace. When placing items the Model is covered. Panel B shows the task view when the Model is open (after clicking and holding the cover for either 0.5 or 5 s). While the Model is open, the Workspace and Resources are covered. Panel C shows the task view after an item is placed incorrectly in the workspace, which results in all grids being covered. After clicking and holding the Workspace grid (for either 0.5 or 5 s), the covers are opened and the view switches to the view in Panel A. The horizontal white bars would shrink in width and disappear in 0.5 or 5 s (depending on the condition) to visually convey how long to hold the cursor on the grid. See the online article for the color version of this figure.

by dragging and dropping items from the Resource grid onto the Workspace grid (see Figure 2). Once the six items were correctly placed in the Workspace, the trial was finished and the participant could move on to the next trial.

For each trial, a random sample of four different shapes and a random sample of four different colors (with the restriction that no neighboring colors were chosen) were used to generate 16 unique items. These 16 items comprised the Resource grid. The Model grid in the trials was created by selecting six items (with replacement) from the Resource grid, and placing them in six different locations. Selecting items with replacement entailed that the same item might have occurred more than once in a Model grid.

The cost of sampling was manipulated by imposing a long or short wait time to reveal the Model grid. During the task, the Model grid was completely covered by a dark gray square. To inspect the Model, the cover could be lifted by clicking and holding the cursor on it (for 0.5 s in the low sampling cost condition and 5 s in the high sampling cost condition). After the Model was uncovered, it stayed uncovered until the cursor was moved away from the Model grid. The cost of making an error was manipulated in a similar way. After an incorrect placement was made, all grids (Model, Workspace, and Resource) were covered by dark gray squares. Specifically, each time an item was released on a cell of the Workspace grid in which it did not belong, all grids were immediately covered. The incorrectly placed item was directly put back in the Resource grid. Only after clicking and holding the cursor on the workspace grid cover (for 0.5 s in the low error cost function and 5 s in the high error cost condition), all covers were removed and the task could be resumed.

Experimental Design

The main factors of interest were the cost of errors and the cost of sampling. We manipulated both the cost of sampling (high or low) and the cost of errors (high or low) between participants, in a full factorial design. Each participant was assigned to one of the four conditions (e.g., low sampling cost and high error cost) and completed all 24 trials in this one condition. We opted for a between-participant design for two reasons. First, participants would be more likely to settle on a consistent copying strategy within a given condition, when not having to switch between blocks of varying error cost and sampling cost. Second, condition order effects might be asymmetrical: An unpublished analysis of a pilot study shows that the effects of sampling cost and error cost are strongly depend on condition order; specifically, participants are more reluctant to change strategies when costs decrease between blocks, and are more willing to change strategies when costs increase between conditions. The only other factors varying between participants were the selection of (the shapes and colors of) the items for the Resource grid and the layout of the Model grids. These factors were not balanced between participants, but were generated randomly on a trial-by-trial basis.

Analysis

We conducted the statistical analyses using Bayesian statistics with the Jeffreys's Amazing Statistics Program (JASP) software using the default priors wherever relevant for conducting Bayesian statistics, and always setting the seed value to 1 for reproducibility (JASP Team, 2023). We conform to the labels suggested by Kass and Raftery (1995) for the interpretation of Bayes factors.

Our main analyses were Bayesian analyses of variance and for all of them, we performed an analysis of effects across matched models (Mathôt, 2017). This analysis compares models that contain the effect in question to equivalent models stripped of the effect. Here, the inclusion Bayes factor (BF_{incl}) reflects the amount of evidence specifically for the (main or interaction) effect in question, compared to the absence thereof.

Measures

The current task is open-ended in nature. Participants are free to copy the items in the way they wish (e.g., in any order, with as many Model inspections and incorrect placements as needed). Therefore a large number of measures could be obtained. To keep the current article streamlined, we focus on and report three direct outcome measures and two derived measures. We believe these measures give the necessary insight into the behavior and are essential to answer the questions in the current study. The three direct outcome measures were:

- Number of inspections: The number of times a participant inspected the Model grid in each trial.
- Duration of an inspection: The time for which the Model remained uncovered (a proxy for how long it was looked at).
- Number of errors: The number of times an item was placed incorrectly (note that accidentally placing an item on top of another item or outside of the grid did not constitute an error).

In the supplements (see Supplemental Material 2 in the online supplemental material), we also included the analysis of a fourth measure, namely the building duration: the time participants took to pick up and place items in the Workspace grid (excluding any wait times).

The two derived measures were:

- First correct streak: This is the number of consecutive correct placements immediately after an inspection and before the next inspection. This measure is used as a proxy for the amount of information that was maintained in VWM (i.e., amount of information soaked up).
- Subsequent attempts: This is the number of placement attempts after the first correct streak and before the next inspection. This measure is used as a proxy for the willingness to utilize information from VWM (i.e., the amount of information squeezed out). Note that the "squeeze" and "soak" measures are effectively independent.

For an elaborate discussion on why our two indirect measures are better proxies for the amount of information memorized, and the amount of information utilized from VWM than the total number of correct placements and the total number of errors per inspection, we refer to our previous work using the same paradigm (Sahakian et al., 2023).

Results

Number of Inspections

The analysis of effects showed decisive evidence ($BF_{incl} = 1.33 \times 10^{10}$) in favor of the main effect of sampling cost (see Figure 3). This means that participants inspected the Model less often when sampling was more costly. However, there was no conclusive evidence ($BF_{incl} = 0.67$) in favor of the main effect of error

Figure 3

The Three Direct Outcome Measures Across the Four (2 Sampling Cost × 2 Error Cost) Between-Observer Conditions

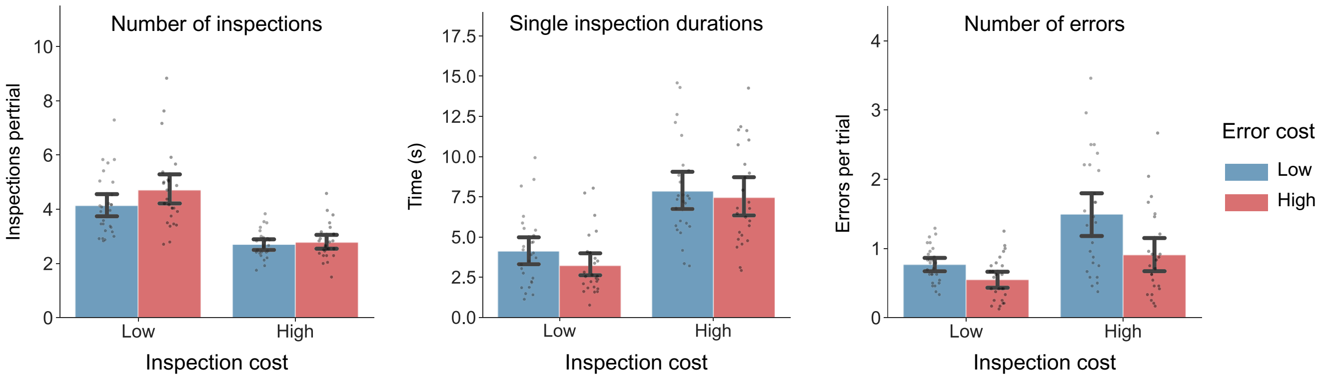

Note. On the left, the mean number of Model inspections per trial is shown. More inspections per trial suggest that fewer items per inspection were memorized, as there were always six items to be copied in a trial. In the middle, the mean durations of Model inspections are shown. Longer inspection times suggest more effort is put into memorizing information. On the right, the number of errors per trial is shown. More errors suggest attempts to place items are based increasingly on less certain information (i.e., a more liberal tendency to utilize uncertain information). In all plots, the bars represent group means, the gray dots represent individual participants' means, and the error bars represent (bootstrapped) 95% confidence intervals. See the online article for the color version of this figure.

cost. This suggests that we cannot make claims based on the data on whether or not the cost of making errors affected the number of inspections. Finally, there was also no conclusive evidence ($BF_{incl} = 0.50$) in favor of an interaction effect between sampling and error cost. This means that the data do not allow us to tell whether the cost of sampling affects the cost of errors (or vice versa).

Duration of an Inspection

The analysis of effects showed decisive evidence ($BF_{incl} = 1.09 \times 10^9$) in favor of the main effect of sampling cost (see Figure 3). This implies that when sampling was costly, participants inspected the Model for longer. Furthermore, we found no conclusive evidence ($BF_{incl} = 0.43$) in favor of the main effect of error cost. This means that the data do not allow claims about the effect of the cost of errors on the inspection duration. Finally, we found substantial evidence ($BF_{incl} = 0.30$) against an interaction effect between the effects of sampling and error cost. This means that the effects of sampling cost and error cost do not influence each other.

Number of Errors

The analysis of effects showed decisive evidence ($BF_{incl} = 3,816$) in favor of the main effect of sampling cost (see Figure 3). This implies that when sampling was costly, more errors were made. Furthermore, we found strong evidence ($BF_{incl} = 48$) in favor of the main effect of error cost. This means that when errors were costly, participants made fewer errors. Lastly, we did not find conclusive evidence ($BF_{incl} = 0.84$) for an interaction effect between the effects of sampling and error cost. Inconclusive evidence here means that the data do not warrant claims about whether an interaction between error cost and sampling cost is present or not.

First Correct Streak (Soak)

The analysis of effects showed decisive evidence ($BF_{incl} = 1.45 \times 10^8$) in favor of the main effect of sampling cost on the length of the first correct streak of placements (see Figure 4). This means that when sampling was costly, participants had more consecutive correct placements after an inspection. We interpret this finding as participants memorizing more information per inspection when inspections were costly. We found substantial evidence ($BF_{incl} = 0.24$) against the main effect of error cost on the length of the first correct streak of placements. This means that participants had the same number of consecutive correct placements after an inspection in the low- and high-error-cost conditions. This demonstrates that participants memorized the same amount of information per inspection regardless of the severity of punishments. Finally, there was no conclusive evidence ($BF_{incl} = 2.44$) in favor of (or against) an interaction effect between sampling cost and error cost on the length of the streak of correct placements.

Subsequent Attempts (Squeeze)

The analysis of effects showed decisive evidence ($BF_{incl} = 5.82 \times 10^6$) for the main effect of sampling cost on the number of placing attempts after the first correct streak (see Figure 4). This means that when inspections are costlier, participants keep attempting to place items for longer. We interpret this finding to mean that when information is hard to access, participants utilize more content from their memory. Furthermore, we found strong evidence ($BF_{incl} = 88$) for the main effect of error cost on the number of placing attempts after the first correct streak. When errors are costlier, participants have fewer placement attempts after the first streak of correct placements. We interpret this as a decreased tendency to utilize information from working memory. Lastly, there was no conclusive evidence ($BF_{incl} = 2.09$) for or against again an interaction effect between sampling cost and error cost.

Discussion

In the current study, we sought to understand what influence the consequences of actions have on the way we memorize and use

Figure 4

The Two Derived Outcome Measures Across the Four (2 Sampling Cost × 2 Error Cost) Between-Observer Conditions

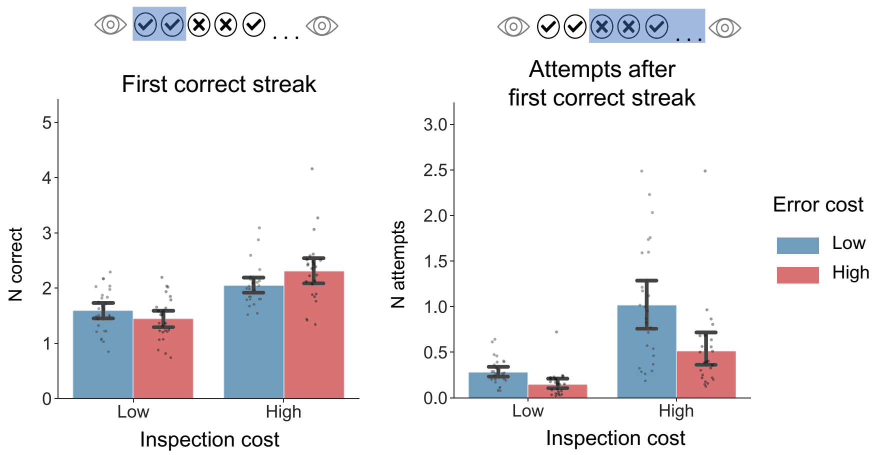

Note. On the left, the mean length of the first correct streaks is shown (the "soak" measure). As visualized above the figure title, this is the number of successive correct placements immediately after a Model inspection (the eye icon). We take this as a proxy for how much information is memorized (or soaked in VWM). Being able to correctly place many items without errors suggests a strong memory of the Model. On the right, the mean number of attempts after the first correct streak is shown (the "squeeze" measure). These are all remaining attempts (correct or erroneous) before the next Model inspection (the eye icon). We take this measure as a proxy for the willingness to utilize information (or squeezed out of VWM). Attempting to place more items (even after making an error) suggests a stronger tendency to utilize any (even uncertain) information that is in VWM, accepting the risk of making more errors. In all plots, the bars represent group means, the gray dots represent individual participants' means, and the error bars represent (bootstrapped) 95% confidence intervals. VWM = visual working memory. See the online article for the color version of this figure.

visual information. We employed a copying task in which participants had to recreate an arrangement of items as shown in an example. This example always remained available for reinspection. Such a paradigm allowed us to investigate the memorization and use of visual information in a task context which naturally engages visual working memory (VWM; Ballard et al., 1995; Draschkow et al., 2021; Sahakian et al., 2023; Somai et al., 2020). One typical finding, across all studies using such a paradigm, is that observers use more information from VWM when task-relevant information becomes harder to access: Observers look at the "Model" less often but for a longer time, and place more items after each inspection. The accessibility manipulation in the current study—which we included as a reference—yielded the same result. The novel manipulation in the current study was the extent to which incorrect actions (e.g., placing an item in an incorrect location) were penalized. We found that when incorrect actions were penalized more, participants did not memorize the required information better. Instead, their actions became more cautious. They only attempted to put information to use of which they were very certain that it would be correct.

The current findings fit well within the theoretical model we have proposed in a previous article (Sahakian et al., 2023). In this model, the information in VWM is represented as a (multidimensional) landscape in feature space, where the height of the landscape corresponds to the certainty of information. (see Figure 5 for a visualization with two feature dimensions). Figures that require more than three dimensions are hard to visualize (in a static figure and in one's mind), so in Figure 5 only the color and one-dimensional location are used to create a feature space. But conceptually this feature space can be extended with more spatial dimensions and visual feature dimensions. For example, a high peak in feature space around "coordinate" (third row, second column, square, red) implies that you are very certain that on the spatial location (third row, second column) there is a square-shaped item that is red. To this landscape of certainty over feature space, we add a so-called "Action Threshold." This action threshold represents the minimum amount of certainty about an item that will lead to an action (rather than a reinspection). If the certainty peak in feature space at (third row, second column, square, red) rises above the threshold an action will follow: a red square will be placed at (third row, second column) in the workspace. How high (or low) the threshold is set depends on the context of the task. A high threshold entails that only very certain information is used. A low threshold entails that even uncertain information will be put to use. Our results suggest that with higher error costs, the certainty of the memory content (the landscape) remains comparable, but the "Action threshold" is increased. Thus, only the most certain pieces of information from memory are used, while the uncertain ones (below the "Action threshold") are left untouched. In essence, this theoretical model combines the continuous nature of working memory content (Bays, 2018; Ma et al., 2014; Schurgin et al., 2020; but see Zhang & Luck, 2008), with the concept of a criterion from signal detection theory (SDT; Banks, 1970; Pastore & Scheirer, 1974). To elaborate on the parallel between our model and SDT, the representational strength in our model (i.e., the certainty

Figure 5

Three-Dimensional Visualization of the Theoretical Model We Use to Describe VWM Use, Based on Certainty of Information and an "Action Threshold"

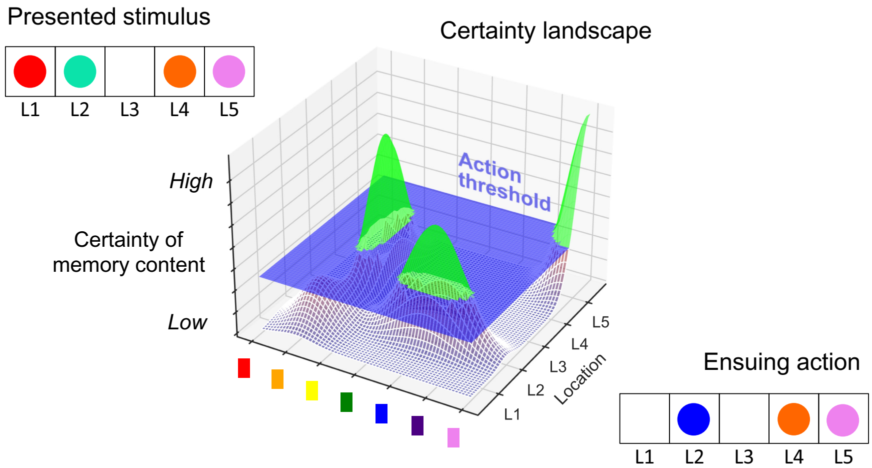

Note. This model dictates how information in mind will be put to use given the memory content and a context-dependent threshold. The presented stimulus has items comprised of two feature dimensions: 1D space (locations L1–L5) and color (red–pink). After an inspection of the stimulus, the memory representation can be visualized by a certainty level on the 2D feature space. On the right in the figure, there is a high peak at (L5, pink) suggesting high certainty that there is a pink item at L5. Peaks can be narrow in one dimension, but a bit wider in the other: the leftmost green peak in the figure is narrow in the color dimension but a bit wider in the spatial dimension suggesting: "There is a certainly orange item, somewhere around L4." The central green peak suggests: "At L2 there is some blueish item." As these three peaks rise above the "Action threshold," they are put to use. Note that while the left and right green peaks resulted in a correct placement, the central peak—which had enough certainty to rise above the threshold and thus be acted on—resulted in an incorrect action: the color should have been turquoise. Finally, note that there is a small peak on the left below the "Action threshold," suggesting "There might be some orange–red item around L1 or L2." As this peak does not rise above the current threshold it is not put to use now. After the next inspection in which the item at L1 is inspected better, this peak might rise above the threshold can be acted on. VWM = visual working memory; 1D = one-dimensional; 2D = two-dimensional. See the online article for the color version of this figure .

of memory content) can be thought of as the sensitivity in SDT terms; while the "Action Threshold" in our model (i.e., the required certainty to act) can be thought of as the criterion in SDT terms. The marriage of continuous resource models of VWM and SDT constitutes a powerful model to describe how and when memories are put into action. Therefore, it lends itself to exploring the strong links between (visual) working memory and action: an idea which has received much attention in recent years (Heuer et al., 2020; Olivers & Roelfsema, 2020; van Ede, 2020; van Ede et al., 2019).

Interestingly, our results show that the amount of memorized information is unaffected by the severity of action consequences. Nevertheless, it is possible that with other design choices (e.g., with more severe consequences, more participants, or more trials per participants), we would have found an effect of error cost on the amount of information that is memorized. Other studies have shown reward-based effects on VWM performance (Cho et al., 2022; Klyszejko et al., 2014; van den Berg & Ma, 2018), which means that in certain situations humans can indeed memorize information better. Regardless, the key point here is the relative difference in effect size (expressed as the partial η^2) between employed strategies. When punishments are more severe, there is a large decrease in utilization: $\eta_p^2 = 0.131$ (i.e., only the very certain memory content is used); but a negligible change in memorization: $\eta_p^2 = 0.004$ (i.e., the total amount of memory content barely changes). A similar point can be made for the relative difference of effect size on memory content between the two equally severe cost manipulations. While a higher sampling cost—which was a 5 (vs. 0.5) s wait to inspect the Model—results in a large change in the amount of memory content ($\eta_p^2 = 0.362$); a higher error cost—which was again a 5 (vs. 0.5) s wait after an error—barely changes the amount of encoded information (as stated before $\eta_p^2 = 0.004$). Even if an experiment with enough participants and very severe consequences would produce a (small) effect of error cost on how much is memorized, it would not change our conclusion that observers—by a huge margin—choose to act more cautiously rather than to memorize more information. Therefore, we argue that our interpretation of the data is warranted in light of these stark differences in effect sizes. Thus, put in context, more severe consequences of incorrect actions lead to more cautious actions, rather than leading to memorizing more information.

As we mentioned in the previous paragraph and the introductory paragraph, there is quite some literature showing that VWM performance is better when observers are incentivized to perform better. Increased performance has been found for items that are associated with higher rewards or punishments (Cho et al., 2022; Klyszejko et al., 2014; van den Berg & Ma, 2018) or higher probability to be probed (Gorgoraptis et al., 2011; Yoo et al., 2018). At first sight, this body of literature might seem at odds with the findings of the present study, as we show that items are not memorized better when errors are punished harder. However, our study differs in a crucial aspect from these studies: in the current paradigm there are two possible strategy changes a participant could adopt in order to avoid punishments. First, such as in all referenced studies, observers could selectively memorize prioritized items better. Second, in contrast to other studies (in which selectively memorizing items better was the only option), the current paradigm also offered a second option: observers could alter their decision criterion, and choose to not use their VWM content, but rather reinspect the relevant information. In doing so punishments could be avoided by acting only on the memory content with strong representations, which decreased the chance of making an error. Interestingly, we find that (when given the option) observers highly favor the second strategy: acting more cautiously, versus the first: memorizing more (or better). This finding shows the importance of considering what options to change one's strategy are available. Depending on the context, some strategy changes are favored over other strategy changes.

To conclude, in the current study, we investigated how the severity of the consequences of incorrect actions affected the way humans approach a task that naturally engages working memory. Much like in everyday life, task-relevant information remained externally available for reinspection. As such, participants had to continuously decide whether to act on the information they currently had in memory or to refresh their memory by reinspecting the external world. When we introduced adverse consequences for incorrect actions, we observed a striking change of strategy: instead of memorizing the task-relevant information better, participants rather chose to act more cautiously and only rely on their best memory representations. So, where does this leave the mushroom hunter? Our findings suggest that a novice mushroom hunter will return from the forest with fewer mushrooms: only the ones of which he was certain that they were safe. The person who went grocery shopping, in contrast, will return from the supermarket with many mushrooms, hoping that one of them is the mushroom he needed. With this study, we highlight the importance of investigation working memory in light of broader and more naturalistic task contexts. Doing so will yield deeper insights and a more applicable understanding of working memory.

References

Anwyl-Irvine, A. L., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. K. (2020). Gorilla in our midst: An online behavioral experiment builder. *Behavior Research Methods*, *52*(1), 388–407. https://doi.org/10.3758/s13428-019-01237-x

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

Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of Learning and Motivation–Advances in Research and Theory*, *8*(C), 47–89. https://doi.org/10.1016/S0079-7421(08)60452-1

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

Banks, W. P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, *74*(2), 81–99. https://doi.org/10.1037/h0029531

Bays, P. M. (2018). Failure of self-consistency in the discrete resource model of visual working memory. *Cognitive Psychology*, *105*, 1–8. https://doi.org/10.1016/j.cogpsych.2018.05.002

Cho, Y. T., Moujaes, F., Schleifer, C. H., Starc, M., Ji, J. L., Santamauro, N., Adkinson, B., Kolobaric, A., Flynn, M., Krystal, J. H., Murray, J. D., Repovs, G., & Anticevic, A. (2022). Reward and loss incentives improve spatial working memory by shaping trial-by-trial posterior frontoparietal signals. *NeuroImage*, *254*, Article 119139. https://doi.org/10.1016/j.neuroimage.2022.119139

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

Gorgoraptis, N., Catalao, R. F., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *Journal of Neuroscience*, *31*(23), 8502–8511. https://doi.org/10.1523/JNEUROSCI.0208-11.2011

Heuer, A., Ohl, S., & Rolfs, M. (2020). Memory for action: A functional view of selection in visual working memory. *Visual Cognition*, *28*(5–8), 388–400. https://doi.org/10.1080/13506285.2020.1764156

JASP Team. (2023). *JASP* (Version 0.17.3) [Computer software].

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795. https://doi.org/10.1080/01621459.1995.10476572

Klyszejko, Z., Rahmati, M., & Curtis, C. E. (2014). Attentional priority determines working memory precision. *Vision Research*, *105*, 70–76. https://doi.org/10.1016/j.visres.2014.09.002

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. https://doi.org/10.1038/36846

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Mathôt, S. (2017). *Bayes like a baws: Interpreting Bayesian repeated measures in JASP*. https://www.cogsci.nl/blog/interpreting-bayesian-repeated-measures-in-jasp

Olivers, C. N., & Roelfsema, P. R. (2020). Attention for action in visual working memory. *Cortex*, *131*, 179–194. https://doi.org/10.1016/j.cortex.2020.07.011

Pastore, R. E., & Scheirer, C. J. (1974). Signal detection theory: Considerations for general application. *Psychological Bulletin*, *81*(12), 945–958. https://doi.org/10.1037/h0037357

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

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a unified theory of visual memory strength. *Nature Human Behaviour*, *4*(11), 1156–1172. https://doi.org/10.1038/s41562-020-00938-0

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

van den Berg, R., & Ma, W. J. (2018). A resource-rational theory of set size effects in human visual working memory. *eLife*, *7*, Article e34963. https://doi.org/10.7554/eLife.34963

van den Berg, R., Zou, Q., Li, Y., & Ma, W. J. (2023). No effect of monetary reward in a visual working memory task. *PLoS ONE*, *18*(1), Article e0280257. https://doi.org/10.1371/journal.pone.0280257

van Ede, F. (2020). Visual working memory and action: Functional links and bi-directional influences. *Visual Cognition*, *28*(5–8), 401–413. https://doi.org/10.1080/13506285.2020.1759744 van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2019). Concurrent visual and motor selection during visual working memory guided action. *Nature Neuroscience*, *22*(3), 477–483. https://doi.org/10.1038/s41593-018-0335-6

Yoo, A. H., Klyszejko, Z., Curtis, C. E., & Ma, W. J. (2018). Strategic allocation of working memory resource. *Scientific Reports*, *8*(1), Article 16162. https://doi.org/10.1038/s41598-018-34282-1

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

Received May 16, 2023
Revision received September 4, 2023
Accepted November 30, 2023 ■