Title

VisualWorking Memory as Decision Making:
Compensation for Memory Uncertainty in Reach Planning

Permalink

https://escholarship.org/uc/item/72v9s1v8

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Lerch, Rachel A
Sims, Chris R

Publication Date

2015

Visual Working Memory as Decision Making: Compensation for Memory Uncertainty in Reach Planning

Rachel A. Lerch (rachel.lerch@drexel.edu)
Chris R. Sims (chris.sims@drexel.edu)

Applied Cognitive & Brain Sciences, Department of Psychology
Drexel University

Abstract

Limitations in visual working memory (VWM) have been extensively studied in psychophysical tasks, but not well understood in terms of how memory limits translate to performance in more natural domains. For example, in reaching to grasp an object based on a spatial memory representation, overshooting the intended target may be more costly than undershooting, such as when reaching for a cup of hot coffee. The current body of literature lacks a detailed account of how the physical consequences and costs of memory error influence what we encode in visual memory, and how we act on the basis of remembered information. Here, we study whether externally-imposed monetary costs influence behavior in a task that involves motor planning based on information recalled from VWM. Our results indicate that subjects accounted for the uncertainty in their visual memory, showing a significant difference in their motor planning when monetary costs were imposed for memory errors. However, our findings indicate that subjects' memory representations *per se* were not biased by the imposed costs, but rather subjects adopted a near-optimal post-mnemonic decision strategy.

Keywords: Visual working memory; decision making; motor planning

Introduction

Visual working memory (VWM) can be defined as a system that actively maintains visual information to serve the needs of ongoing tasks (Luck & Vogel, 2013). The limitations of this system have been the subject of numerous psychophysical studies, with particular interest in understanding possible limits in the number of items that can be sustained in memory, as well as the quality or precision of recalled representations, particularly as the set size increases (Luck & Vogel, 2013; Ma, Husain, & Bays, 2014). Building on a substantial body of behavioral results, recent work has also focused on the development of computational models that explain and predict limits in memory performance, on the basis of information theory (Sims, Jacobs, & Knill, 2012; Orhan, Sims, Jacobs, & Knill, 2014; Sims, 2015) or theories based on limits in neural coding (Franconeri, Alvarez, & Cavanagh, 2013; Bays, 2014).

Limits in visual memory, though extensively studied, are not equally well understood in terms of how they influence behavior in ecological tasks, such as in motor planning and execution. What we are still largely left asking is the following: How is VWM used in natural tasks, and how might the costs of misremembering influence how and what we remember?

Hollingworh, Richard, and Luck (2008) demonstrated that visual working memory is important for a range of natural tasks, including gaze correction following saccadic error. Brouwer and Knill (2007, 2009) demonstrated that VWM is similarly critical for online movement control, even when reaching for targets that are currently visible. This paper builds on the close connection between VWM and motor control, and examines how imposed monetary costs on VWM errors affect an individual's movement planning. We therefore examine motor planning and visual working memory from the perspective of decision theory (Körding, 2007).

One intuitive example that illustrates how errors in VWM can translate into relevant behavioral costs is the so-called 'wine-glass problem'. You might imagine yourself on a dinner date, and maintaining eye contact with your date while simultaneously reaching to pick up your glass of wine. In this example there are two sources of information available to the brain regarding the location of your wine glass: information from the visual periphery present at the time of planning, and remembered information from previous fixations on the glass. However, both sources of information are of limited fidelity (Brouwer & Knill, 2007, 2009), and in this situation memory error may lead to significant social costs. If you misremember the location of the wine glass as being further from you than it really is, you might overshoot and knock over the glass. In this case, it is less costly to misremember the target as being closer to you than it really is, since this would result in undershooting and having to make a slight additional reaching movement to adjust for your mistake (example adapted from Trommershäuser, Maloney, & Landy, 2008).

From this perspective, the study of VWM can be approached as a form of decision making under risk. This builds on other research which examines motor planning as a form of decision making (Trommershäuser et al., 2008; Wolpert & Landy, 2012) and which similarly asks: How do the costs of motor error influence motor planning? We add to this research by studying how imposed costs affect visual spatial memory as well as the planning of hand movements on the basis of remembered information.

An important and closely related question is whether external costs bias the contents of visual memory, or rather, whether costs influence how people act on the basis of uncertain memory information. Previous research in categorical perception (Goldstone & Hendrickson, 2010) has demonstrated that the categorical structure of visual information influences our ability to discriminate between objects. Categorical perception effects raise the possibility that the costs of memory error may similarly bias the contents of visual working memory. In the context of remembering spatial information, the sensitivity of VWM to the costs of memory error might lead to biases in the recall of spatial locations. This possibility is further bolstered by a number of findings which show that VWM is sensitive to the statistical structure of the visual environment (Orhan et al., 2014).

To explore this idea, we developed a task that required participants to remember an array of colored targets, and then after the stimuli were removed, touch the remembered location of a cued target using a stylus (Figure 1). In different conditions, monetary penalties were associated with different kinds of memory errors: overshooting vs. undershooting the intended target. Successfully touching a target (hitting anywhere within the target boundary) always earned the participant money, but depending on condition, either overshooting or undershooting the target could decrease the participant's total earnings.

We hypothesized the following: (1) Memory precision should deteriorate with increasing set size (the number of items stored in memory). This expectation is well supported in the body of previous research (Ma et al., 2014). (2) Mean aim point should differ between conditions, that is people will undershoot when there are costs for overshooting and vice versa. (3) If people are sensitive to the uncertainty in their memory then they should aim further away from penalty regions in the large set size conditions where memory uncertainty is greater. Thus a strategy of under- or overshooting the target may not represent a simple and fixed heuristic, but rather may be more intricately tied to the cost structure of the task and to the level of uncertainty in memory. (4) The contents of memory may also be biased by the costs associated with memory error. This latter hypothesis requires distinguishing between biases in memory representations, and participants adopting a post-mnemonic decision strategy.

Methods

Participants

Twelve individuals (8 female) participated in the experiment (age range 18 to 35 years, mean 22.42). All participants reported normal or corrected-to-normal vision and no diagnosed motor impairments. Participants completed two experimental sessions, and were compensated a minimum of $20 with additional monetary incentives based on performance. All subjects provided informed consent according to procedures approved by the Drexel University institutional review board.

Apparatus

Stimuli were presented on a custom built "smart table", consisting of a glass surface (101x64cm) backed by rear projection film (Figure 2). A digital projector and mirror mounted below the glass were used to render stimuli onto the surface. The table height was 105cm. Participants held a stylus (shown resting on the tabletop) in their dominant hand for indicating responses. The stylus had reflective markers attached to the end; these markers were tracked by a motion capture system (NaturalPoint OptiTrack) that recorded the spatial position of the tip of the stylus in real-time at 120Hz.

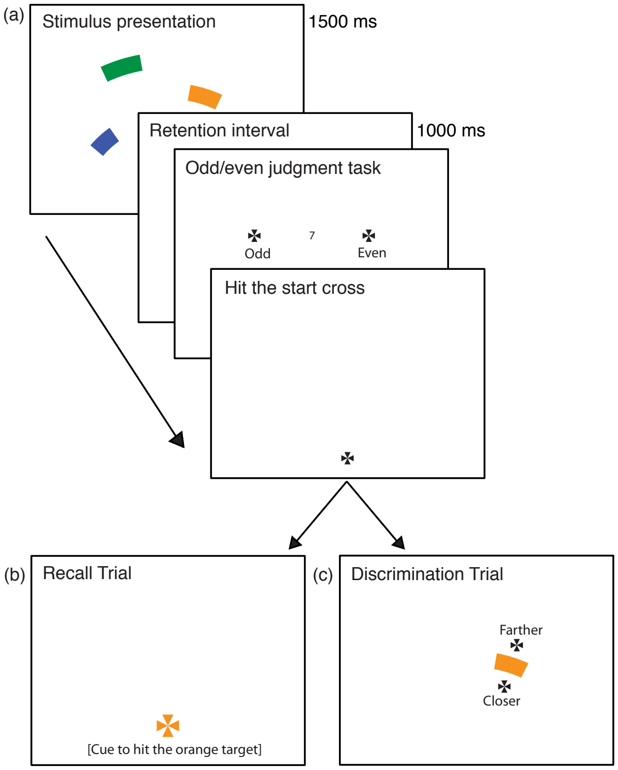

Figure 1: Sequence of events in the task. (a) Targets (one or three annular sectors) were presented for 1,500ms, followed by a blank retention interval. Subjects then completed an odd/even digit judgment task, and touched a start cross. Depending on the trial, subjects were then instructed to either (b) touch one of the targets, cued by the color of the start cross (recall trial), or else (c) complete a memory discrimination task and judge whether a probe stimulus was presented closer or farther than the original item.

Stimuli

The memory stimuli consisted of one or three colored targets (colors chosen randomly from the set blue, green, purple, and orange) that varied in angle and distance from the participant (see Figure 1). Each target was an annular sector (i.e., a section of a ring), with angular width = 10 degrees, and radial thickness = 6.35 cm. The target locations were defined in polar coordinates (angle and radial distance from the subject), with the angle to the target center sampled from the range (-45°, 45°), where 0°indicates straight in front of the subject. The radial distance to the targets varied from 12.7cm to 41.28cm. Target locations were randomly sampled on each trial subject to the constraint that targets did not overlap in angle.

The targets were rendered on top of a "white noise" pixel background. The pixel noise masked reflections on the glass table surface that could otherwise potentially be used to aid in

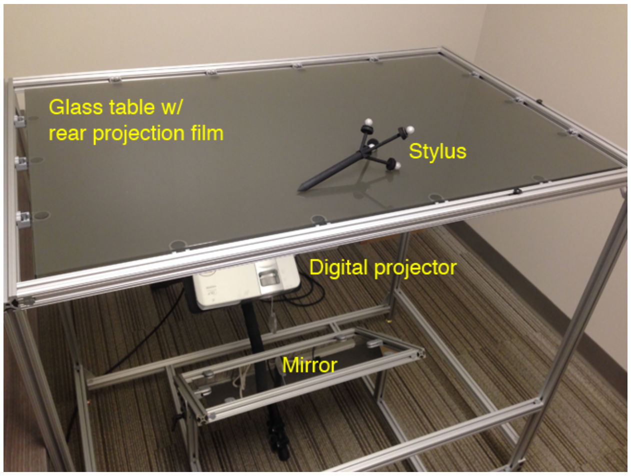

Figure 2: The apparatus used for the experiment. See text for description.

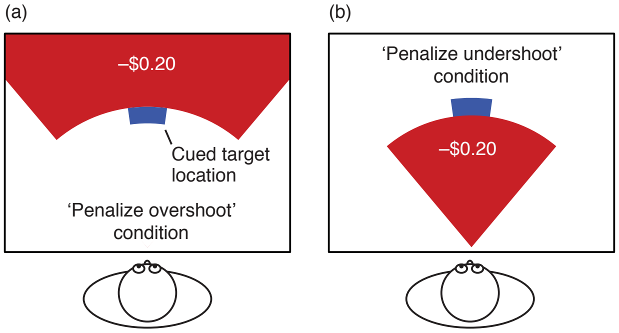

Figure 3: Recall trials incorporated monetary penalties, depending on the condition. (a) In the 'penalize overshoot' condition, touching an area in the red region (overshooting the target) resulted in a monetary penalty of $0.20. (b) The 'penalize undershoot condition. Touching an area in the white region had no penalty.

the localization of stimuli. Each trial began by having the subject touch a 'start cross' located at the proximal edge of the table. Stimuli were then presented for 1,500ms, followed by a 1,000ms retention interval. After the retention interval, subjects then completed an odd/even digit judgment task, where they were asked to indicate whether a single digit (randomly chosen from the set 1–9) was odd or even. The odd/even judgment task forced participants to make eye movements, and hence prevent using visual gaze as an 'external memory'. Then, subjects once again touched the start cross. Depending on the trial, subjects were cued to either complete a memory recall trial (Figure 1b) or a memory discrimination trial (Figure 1c). During recall trials, subjects were instructed to attempt to touch the location where the cued target had previously been displayed. After touching the display, subjects received visual feedback on whether they hit or missed the target, and whether it resulted in a monetary payoff or penalty (depending on condition, described in the Procedure section). During the discrimination trials, participants were asked to judge whether a probe stimulus was presented closer or farther than the original item. No feedback was given during discrimination trials.

Procedure

Each participant completed two experimental sessions, conducted on separate days. Each session consisted of two blocks that varied in terms of set size—the number of targets that needed to be remembered. In one of the blocks, subjects were shown a single target on each trial (set size = 1), and in the other block three targets were displayed (set size = 3). The order of the blocks was counterbalanced across participants. Each block consisted of 100 recall trials, and 50 discrimination trials. The two trial types were randomly interleaved, and subjects could not distinguish between the two trial types at the start of each trial. During recall trials, subjects gained or lost money depending on whether they successfully 'hit' the cued target. In particular, successfully hitting a target earned the participant 10 cents, in addition to a base pay of $5.00 for each block. The two experimental sessions differed in terms of the monetary penalty associated with missing a target (illustrated in Figure 3). In one of the sessions, 'overshoot errors'—touching a location farther than the target—cost the participant 20 cents (Figure 3a). In the other session, 'undershoot errors'—touching a location closer to the participant than the actual target—cost 20 cents (Figure 3b). Undershooting in the penalize overshoot condition resulted in zero cents (neither gain nor penalty), and vice versa for the penalize undershoot condition. Subjects were instructed at the start of each session on the relevant payouts and penalties for that condition. The order of the two penalty conditions was counterbalanced across participants.

During discrimination trials, a probe stimulus was displayed that differed from the location of one of the original stimulus items in terms of its radial distance. The participant completed a 2-alternative forced choice, deciding if the probe stimulus was closer or farther than the original target (Figure 1c). The distance of the probe stimulus relative to the true stimulus location was controlled using a one-up/one-down adaptive staircase procedure (Lu & Dosher, 2014), using two interleaved staircases. Thus, the discrimination trials were designed to determine the participants psychophysical threshold: the probe distance that was indistinguishable from the remembered location of the original target. This enabled a measure of whether memory for target distance was biased by the penalty condition.

To summarize, the experiment utilized a 2 × 2 within-subject design, manipulating set size (one or three items) and penalty condition (penalize undershoot vs. penalize overshoot). During each session, the penalty condition was held constant, but the two within-session blocks varied the set size. In total, each subject completed 400 recall trials, and 200 discrimination trials.

Table 1: Mean payoff in each condition of the experiment. Standard deviations given in parentheses.

Condition	Set Size 1	Set Size 3
Penalize Undershoot	\$12.70 (2.09)	\$6.12 (2.97)
Penalize Overshoot	\$13.08 (1.22)	\$6.19 (2.29)

Results

On average, participants earned \$38.08 across the two sessions of the experiment (highest earning participant = \$44.70; lowest = \$23.10). Average payoff for each condition is reported in Table 1. It is immediately apparent that participants found the task much harder in the set size = 3 condition, earning approximately half as much as they did in the set size = 1 condition.

Unlike most psychophysical studies of VWM, this experiment offered the possibility for participants to mitigate the negative consequences of memory error. In particular, we hypothesized that participants would exhibit a tendency to overshoot the target in the penalize undershoot condition, and vice versa.

To test this hypothesis, we computed the relative aiming position on each recall trial as the difference between the radial distance of the participant's response, and the radial distance to the center of the cued target. According to this measure, positive values indicate overshooting the center of the target, and negative values indicate undershooting. Mean relative aim is plotted in Figure 4.

A 2 × 2 within-subjects ANOVA was conducted on mean relative aim, with set size and penalty condition as factors. The results indicated that mean relative aim significantly differed between penalty conditions, $F(1,11) = 19.62, p = 0.001$, generalized $\eta^2 = 0.44$. The interaction between set size and penalty condition aslo reached statistical significance, $F(1,11) = 10.15, p = 0.009$, generalized $\eta^2 = 0.120$. Hence, subjects significantly shifted their mean aim location away from the penalty region, and the magnitude of this shift was larger in the set size = 3 condition.

Performance in the four conditions differed not just in the mean aim point, but also in the variability in aiming distance. The distributions of aim points are shown in Figure 5. This figure illustrates both the shift in mean aim point (as shown in Figure 4), as well as a substantial increase in variability in the set size = 3 conditions.

Given the observed shift in aim point between the two penalty conditions, we were interested in distinguishing between two possible explanations. One possibility is that prolonged exposure in each penalty condition resulted in a bias in the contents of VWM. That is to say, subjects' VWM systems consistently remembered the targets as closer than they really were in the penalty overshoot condition. An alternative explanation is that the memory representations were not biased; rather, subjects adopted a post-mnemonic decision strategy to shift their aim away from the penalty regions.

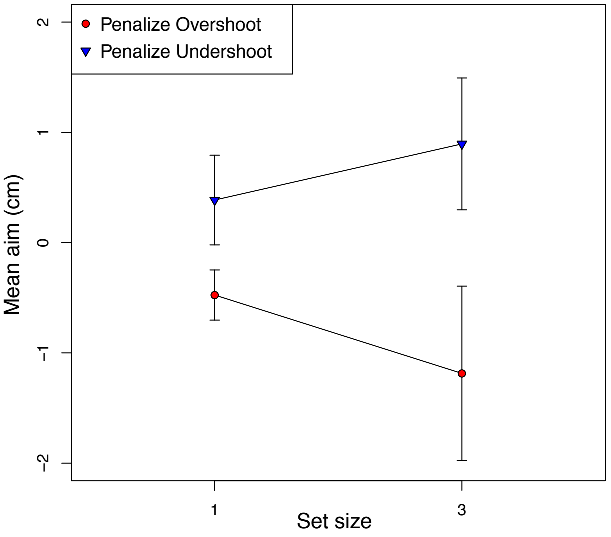

Figure 4: Mean aim point (averaged across subjects) for each penalty and set size condition. Aim point is defined as the radial distance of the participant's response relative to the center of the target. A value of zero indicates hitting the center of the target, positive values are representative of overshooting the target center, negative values indicate undershooting the target center. Error bars indicate 95% confidence intervals (computed across subjects).

The data from the discrimination trials allowed us to distinguish between these two possibilities. We fit a simple psychometric function (a Gaussian cumulative distribution) to the discrimination trials from each participant, using maximum likelihood estimation. The threshold parameter of the psychometric curve, μ, indicates the probe stimulus distance such that the subject was not able to reliably detect whether it was farther or closer than the true stimulus location. Hence, positive values of μ indicate that memory was biased towards remembering targets as further away than they really were, while values of μ close to zero indicate an absence of memory bias. Note that there were no monetary incentives for the discrimination trials, and no performance feedback was given on these trials. Hence, subjects had no incentive to apply a post-mnemonic decision strategy to the discrimination trials. The remaining parameter of the psychometric curve, σ, controlling the slope of the curve, offers an independent measure of memory uncertainty in each condition.

The mean parameter estimates from each condition are reported in Table 2. A 2 × 2 ANOVA was conducted on the parameter estimates from the psychometric curves. The results did not reveal any significant shift in μ across penalty conditions, $F(1,11) = 1.38, p = 0.27$, generalized $\eta^2 = 0.020$. Hence, given the current data, the observed shift in mean aim point during the recall trials (illustrated in Figure 4) is most parsimoniously explained as an adaptive decision strategy, rather than a bias in memory *per se*. Consistent with this

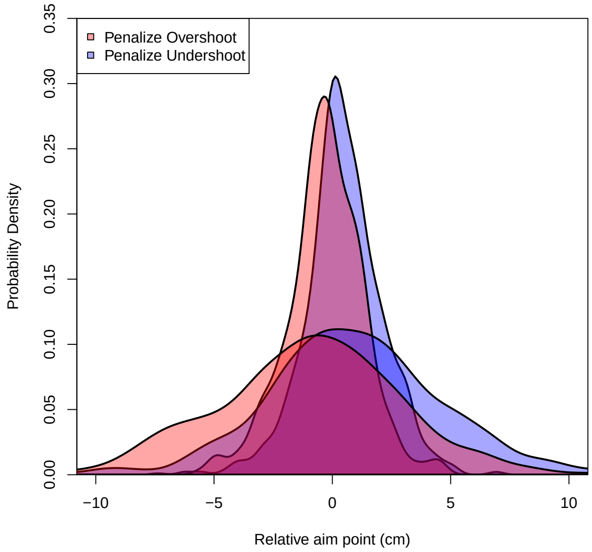

Figure 5: Distribution of responses on recall trials across all subjects and trials, relative to the center of the cued target. Continuous distributions were obtained using a kernel density estimate. In the penalize overshoot conditions, the response distributions are shifted towards negative values (undershooting the target) and vice versa in the penalize undershoot conditions.

Table 2: Mean parameters of the psychometric curve fit to the discrimination trials from each condition (SD in parentheses).

Set Size	**Penalty Cond.**	μ	σ
1	Undershoot	.018 (.042)	.052 (.030)
1	Overshoot	.009 (.030)	.070 (.034)
3	Undershoot	.045 (.044)	.106 (.052)
3	Overshoot	.031 (.042)	.126 (.042)

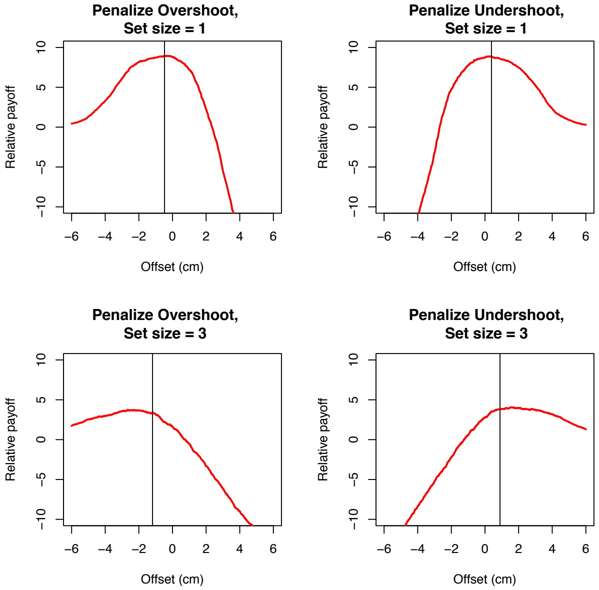

Figure 6: Predicted payoff as a function of mean offset—how much the subject attempts to undershoot or overshoot the targets. Each red curve shows the utility function for a given magnitude of offset, based on the distribution of aim points in the combined empirical data. The peak of the curve indicates the optimal magnitude of offset. The vertical black line in each figure indicates the empirically observed magnitude of offset in each condition, averaged across subjects.

interpretation, we also conducted a post-experiment survey with each participant. The survey indicated that 9/12 subjects reported adopting a deliberate strategy of aiming away from the penalty regions.

A separate ANOVA on the parameter σ from the psychometric curve found that the main effect of set size was significant, $F(1,11) = 32.06, p < 0.01$, generalized $\eta^2 = 0.340$. This result simply confirms that memory discriminability was poorer in the larger set size conditions.

Were subjects optimal or sub-optimal in their motor planning? The answer to this question is potentially informative as it may place constraints on the class of mechanisms offered as an explanation for performance. In particular, if subjects performed at a near-optimal level, it would suggest their decision strategies were adaptive to the actual costs of memory error defined by the task, and their level of memory uncertainty, rather than reflecting an invariant or approximate heuristic.

To answer this question, we examined how greater or smaller shifts in the distribution of responses (as shown in Figure 5) would influence expected payoff. There are several ways of approaching this analyis, such as fitting models of response variability to each participant. For the present paper we adopted perhaps the simplest approach: We subtracted the mean from the response distribution in each condition (centering all four distributions in Figure 5 around zero), and computed the predicted payoff as varying amounts of response bias were added back in. We performed the analysis on the aggregate data from all participants.

Figure 6 shows the resulting utility curves for each condition. The red curves show the predicted payoff as a function of the shift towards overshooting or undershooting the center of the target. The peak of these curves represent the optimal magnitude of offset. The vertical black lines show the empirical data—the mean offset actually observed in each condition. The results in Figure 6 show that, in aggregate, participants achieved close to optimal performance in the task. The empirical offset was nearly indistinguishable from the optimal offset for the set size one condition. Although the analysis suggests that subjects should have increased the magnitude of their offset in the set size three conditions, the net increase in payoff would be minimal. On the basis of these results, we argue that performance in the task does not simply reflect a fixed heuristic (e.g., 'aim away from the penalty regions'), but rather shows that motor planning was sensitive to both external costs and intrinsic memory uncertainty.

Discussion

The study of visual working memory has mostly focused on measuring capacity limits and studying its psychophysical properties. Largely lacking in the literature, however, are approaches that consider how these limitations translate to ecologically relevant tasks.

In this paper we sought to investigate how resource limitations and memory error influence the use of visual memory in motor planning. We developed an experimental paradigm that captured an important property of natural tasks, where memory errors may result in negative consequences or costs. We found that subjects—whether subconsciously or deliberately—offset their aim relative to penalty areas. This suggests that motor planning is sensitive to both the limitations of encoding visuospatial information, as well as the costs of memory error.

As hypothesized and supported by previous studies, memory precision deteriorated with increased set sizes. More interesting, subjects adaptively compensated for their memory uncertainty and the costs of memory error, by shifting their aim away from task-defined penalty regions. In aggregate, the mean direction and magnitude of shift in motor planning was near-optimal; this suggests that the observed overshooting and undershooting strategies are not fixed heuristics, but are rather more intricately linked to the cost structure of the task and to the level of uncertainty in memory.

We hypothesized that externally imposed costs might influence not only motor planning, but also the manner in which information is encoded in memory. However, we did not find a significant difference in performance across the discrimination trials from the two penalty conditions. The results of this study thus best attribute the significant difference in mean aim to an adaptive post-mnemonic decision strategy, rather than a bias in the participants' memory of target locations. This remains an area of interest for future studies, as the body of literature examining categorical perception shows that learned categories can exert a top-down influence on visual perception (Goldstone & Hendrickson, 2010). In the current study, we failed to observe a similar 'categorical memory' effect.

Lastly, the present results need to be integrated with existing work in developing computational models of visual working memory (Sims et al., 2012; Orhan et al., 2014; Sims, 2015). Existing computational models have largely focused on predicting the limits of VWM, but have not adequately addressed how visual working memory is used in natural tasks. Our results demonstrate that motor planning has access to both the contents of spatial information in VWM, but also the uncertainty or reliability of remembered information. This uncertain information is appropriately combined with the costs of memory error, demonstrating that VWM is an integral part of a larger adaptive biological control system.

Acknowledgments

The authors would like to thank Elizabeth Ellman for her assistance with data collection.

References

Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *The Journal of Neuroscience*, *34*(10), 3632–3645.

Brouwer, A.-M., & Knill, D. C. (2007). The role of memory in visually guided reaching. *Journal of Vision*, *7*(5), 6.

Brouwer, A.-M., & Knill, D. C. (2009). Humans use visual and remembered information about object location to plan pointing movements. *Journal of vision*, *9*(1), 24.

Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: competitive content maps for attention and memory. *Trends in cognitive sciences*, *17*(3), 134–141.

Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(1), 69–78.

Hollingworth, A., Richard, A. M., & Luck, S. J. (2008). Understanding the function of visual short-term memory: transsaccadic memory, object correspondence, and gaze correction. *Journal of Experimental Psychology: General*, *137*(1), 163.

Körding, K. (2007). Decision theory: what "should" the nervous system do? *Science*, *318*(5850), 606–610.

Lu, Z.-L., & Dosher, B. (2014). *Visual psychophysics: From laboratory to theory*. MIT Press.

Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in cognitive sciences*, *17*(8), 391–400.

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature neuroscience*, *17*(3), 347–356.

Orhan, A. E., Sims, C. R., Jacobs, R. A., & Knill, D. C. (2014). The adaptive nature of visual working memory. *Current Directions in Psychological Science*, *23*(3), 164–170.

Sims, C. R. (2015). The cost of misremembering: Inferring the loss function in visual working memory. *Journal of Vision*, *15*(3), 1–27.

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2012). An ideal observer analysis of visual working memory. *Psychological review*, *119*(4), 807.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in cognitive sciences*, *12*(8), 291–297.

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current opinion in neurobiology*, *22*(6), 996–1003.