

Optimal maintenance and use of uncertainty in visual working memory

Aspen H. Yoo, Wei Ji Ma

{aspen.yoo, weijima} @ nyu.edu

Department of Psychology, Center for Neural Science
New York University, 4 Washington Place
New York, NY 10003 USA

Abstract

Unlike in perceptual tasks, it is unclear whether humans near-optimally use uncertainty information in their visual working memory (VWM) decisions. Some circumstantial evidence is available: people can explicitly report their uncertainty after a delay and can near-optimally integrate knowledge of uncertainty with working memories. However, it is unclear whether people can do the conjunction: accurately store uncertainty information in VWM and use it in a subsequent decision. To investigate this, we collected data in two orientation change detection tasks. One task did not require the maintenance of uncertainty information and the other did. We factorially evaluate Bayesian observer models with different assumptions about the memory noise generating process, the observer's assumption of this process, and the observer's decision rule. For both experiments, the model that best fits human data assumes that memory precision varies as a function of stimulus reliability and other internal fluctuations, observers know their memory uncertainty on an individual-item basis, and observers optimally integrate information across items when making their decision. These results provide evidence that participants are able to maintain uncertainty information across a delay, and use it optimally in subsequent decisions.

Keywords: visual working memory, psychophysics, Bayesian modeling, variable precision

While traditional theories of visual working memory (VWM) have described remembering something as an all-or-nothing process (e.g., Zhang & Luck, 2008), more recent theories have described memories as noise-corrupted representations of memoranda. In this framework, VWM is a limited resource that can be flexibly allocated to any number of stimuli or stimulus features (e.g., Bays & Husain, 2008; Yoo, Klyszejko, Curtis, & Ma, 2018). The models in this framework, called resource models, are able to explain human data better than previous theories can (van den Berg, Shin, Chou, George, & Ma, 2012; Fougne, Suchow, & Alvarez, 2012; Bays & Husain, 2008; Ma, Husain, & Bays, 2014).

In a resource model framework, categorization tasks such as change detection are no longer trivial. Change detection becomes a signal detection problem, in which an observer must maintain the uncertainty information about the memory in order to perform optimally (Wilken & Ma, 2004). Do people maintain and use uncertainty? Do people perform optimally in VWM change detection tasks?

We know that people are able to explicitly report their confidence regarding a memory, (e.g., Rademaker, Tredway, & Tong, 2012; Yoo et al., 2018) and choose which memoranda are remembered better than others (e.g., Fougne, Cormiea, Kanabar, & Alvarez, 2016). This suggests that uncertainty information can be maintained during a delay and explicitly recalled, but does not investigate whether that information is used in later decisions. Keshvari and others showed the complement: they indicated that observers can use uncertainty cues optimally in a change detection task, but did not investigate whether participants maintained that uncertainty information in VWM (Keshvari, Berg, & Ma, 2012).

In this study, we replicate and extend the results of Keshvari et al., 2012. We replicate the experiment and models finding that people can use uncertainty information optimally when that information is available to them at decision. We then modify the experimental paradigm such that uncertainty information must be maintained in order to be used and ask if people are able to maintain and use memory certainty.

Experimental Methods

Seven participants completed two experiments (Figure 1). Both experiments were four-item orientation change detection tasks. Each trial began with a fixation cross. Four ellipses were presented for 100 ms, followed by a 1000 ms delay, then by the second stimulus set presentation for 100 ms. On one half of the trial, the orientation of one ellipse changed; this change was drawn from a uniform distribution and each item was equally probable to contain the change. The participant indicated with a button press whether they believed there was an orientation change between the two displays.

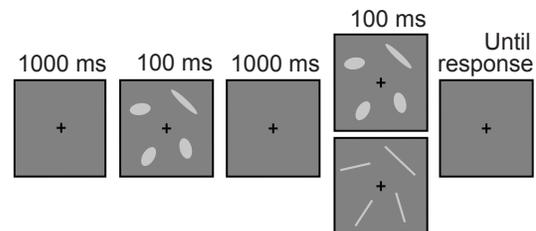


Figure 1: Trial sequence showing, in the second stimulus presentation, ellipses (top) as in Exp. 1 and lines (bottom) as in Exp. 2. Lines were presented instead of ellipses to avoid providing cues to the precision with which the first items were encoded.

Importantly, each ellipse presented could provide orienta-

tion information with either high or low reliability, manipulated through ellipse eccentricity. In order to perform optimally, the participant would have to use knowledge of each item's memory uncertainty when making decisions. In Experiment 1, the stimuli on both displays were ellipses (top example Figure 1). Ellipse reliability could be used as a cue of memory uncertainty, which would allow people to use uncertainty information without having to actually maintain it over a delay. In Experiment 2, the stimuli on the second display were lines (bottom example in Fig. 1), which does require the maintenance of memory uncertainty in order for it to be used. Participants completed 2000 trials of each experiment over six one-hour sessions.

Models: Fitting Experiment 1

In this first section, we replicate the models from Keshvari et al., 2012 and use them to fit Experiment 1 data.

Encoding stage

The probability of change occurring on each trial is set to 0.5, $p(C) = 0.5$. If $C = 1$, any item is equally probable to be changed. All the orientations of the items presented on the first display, ξ , are independently drawn from a uniform distribution over orientation space.

If there is a change, this change Δ is drawn from a uniform distribution. The orientations at the second display, ϕ , are the orientations at the first display plus Δ at the location of change. The noisy measurements of each item on each display, $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{y} = (y_1, \dots, y_N)$, respectively, is drawn from a Von Mises distribution centered on the actual orientation presentation,

$$p(\mathbf{x}|\xi, \kappa_x) = \prod_{i=1}^N p(x_i|\xi_i, \kappa_{x,i}) = \prod_{i=1}^N \frac{1}{2\pi I_0(\kappa_{x,i})} e^{\kappa_{x,i} \cos(x_i - \xi_i)}$$

$$p(\mathbf{y}|\phi, \kappa_y) = \prod_{i=1}^N p(y_i|\phi_i, \kappa_{y,i}) = \prod_{i=1}^N \frac{1}{2\pi I_0(\kappa_{y,i})} e^{\kappa_{y,i} \cos(y_i - \phi_i)}$$

The κ s are the concentration parameter of the Von Mises distribution, and are related to the precision with which each item is remembered. The subscript of each κ indicates which item it refers to (e.g., $\kappa_{x,i}$ is for the i^{th} item the first stimulus presentation). We consider Fixed Precision and Variable Precision encoding of items (van den Berg et al., 2012). With a Fixed Precision assumption of encoding noise, the κ for each item is determined only by its ellipse reliability; items with high ellipse reliability would be encoded with parameter κ_{high} , and the lower reliability ellipse with κ_{low} . In other words, $\kappa_{x,i}$ and $\kappa_{y,i}$ are equivalent.

With a Variable Precision encoding scheme, $\kappa_{x,i}$ and $\kappa_{y,i}$ are themselves random variables and thus can differ. The Fisher information of the Von Mises distribution, J , is drawn from a gamma distribution with mean precision \bar{J} and scale parameter τ . We assume that the precision of memory corresponding to low-reliability ellipses are drawn from a gamma distribution with mean \bar{J}_{low} , and those corresponding to high-reliability ellipses are with \bar{J}_{high} .

Decision stage

Decision variable. We assume that the observer calculates for each item, the ratio of the likelihood of there being a change and the likelihood of there being no change, $d_i = \frac{p(C=1|x_i, y_i)}{p(C=0|x_i, y_i)}$.

When calculating the decision variable, we consider that the observer has an assumption about their memory noise independent of the true generative process. We consider that the observer may have one of the three assumptions:

1. Variable precision (V): mean memory precision varies with ellipse shape, and there is additional noise for each item at each presentation.
2. Fixed precision (F): memory precision varies only with ellipse shape.
3. Same precision (S): memory precision is the same throughout the experiment, and does not vary with ellipse shape or anything else.

Decision rule. The observer uses this decision variable to decide whether they believe a change occurred. We consider two decision rules: the optimal (O) and max (M) rules.

The Bayes-optimal observer responds "change" whenever the probability of there being a change is greater than 0.5. This is equivalent to observer responding "change" if the ratio of the likelihood of there being a change and the likelihood of there being no change is greater than 1:

$$\frac{p_{\text{change}}}{1 - p_{\text{change}}} \frac{1}{N} \sum_{i=1}^N d_i > 1,$$

where p_{change} is the observer's belief of $p(C = 1)$. An observer using the max rule does not optimally combine evidence, but rather responds "change" whenever the maximum evidence of change is greater than some criterion, k .

$$\max_i d_i > k.$$

Parameters

There are two possible encoding schemes ((V)ariable, (F)ixed), three possible observer assumptions of noise ((V)ariable, (F)ixed, (S)ame), and two possible decision rules ((O)ptimal, (M)ax). Factorially combining each of these characteristics would yield 12 different models. We choose not to consider the two models in which the generative model is F but the observer assumes V, so we test a total of 10 models. We denote each model by the letters corresponding to their encoding scheme, assumption, and decision rule (e.g., VVO is the model with variable precision encoding, an observer that assumes variable precision encoding, and an optimal decision rule).

Observers with an V encoding scheme have parameters \bar{J}_{high} and \bar{J}_{low} corresponding to the mean precision of the high

and low reliability ellipses, respectively. The scale parameter, τ , of the gamma distribution from which item-wise precision is drawn is shared across the two ellipse values. If the observer incorrectly believes they are F, then they have precision $J_{\text{high}} = \bar{J}_{\text{high}}$ and $J_{\text{low}} = \bar{J}_{\text{low}}$ for high and low reliability ellipses, respectively. If the true generative model and observer assumption are both F, then the model does not have the τ parameter.

If the participant has the incorrect assumption that their precision is equal across reliabilities, items, and trials, then there is an additional parameter J_{assumed} , corresponding to the assumed precision of all items.

There is one additional parameter for the decision process. If the observer uses the optimal decision rule, there is parameter p_{change} corresponding to the observer's belief of the prior. While it is 0.5, we allow the observer to have an incorrect belief. If the observer uses the max rule, then we have parameter k , corresponding to the decision criterion.

Model fitting and comparison

We used Bayesian Adaptive Direct Search (BADs; Acerbi & Ma, 2017) to estimate, for each participant and model, the parameter combination θ that maximizes the likelihood of the data given the model. We compare models using AICc and BIC. The results are consistent for both measures, so we only show BIC results.

Results and Discussion

We find that the VVO model provides a good qualitative fit of the data (Figure 2A). Model comparison (Figure 2B) indicates that participants are best fit by the VVO model, which assumes encoding with variable precision, an observer who assumes variable precision encoding, and an optimal decision rule. This model fit better for every subject for almost all models; it fit better than VVM for 4 of 7 subjects (worse by 4, 8, and 26 AICc/BIC points, for the remaining three participants). These results suggest, like Keshvari et al., 2012, that participants are aware of how their memory noise is generated, and may use that information optimally when making change detection decisions.

Models: jointly fitting both experiments

In the first experiment and models, we asked if people could use uncertainty information optimally when it is presented to them during the time of the decision. This was a standard replication. In the second experiment and second set of models, we ask if people can maintain and use uncertainty information optimally. Here, we go over the modifications to the above models to test this prediction.

We have an additional parameter, \bar{J}_{line} , which corresponds to the mean precision with which each line on the second display is measured by the observer. The gamma function from which each item's precision is drawn shares the same τ as \bar{J}_{high} and \bar{J}_{low} . If the observer incorrectly assumes that they have the same precision across stimuli, we allow them to have different representations for ellipses, $J_{\text{assumed,ellipse}}$, and lines,

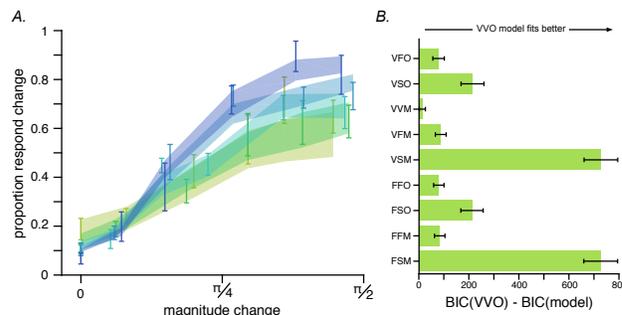


Figure 2: A. Experiment 1 model fits. $M \pm SEM$ of data (error bars) and model predictions (fill) across subjects for the VVO model. Colors indicate how many high reliability ellipses are presented on each display (going from 0 in green to 4 in blue). B. model comparison. $M \pm SEM$ BIC(VVO)-BIC(model) across subjects (greater value indicates worse fit in comparison to the VVO model).

$J_{\text{assumed,line}}$. This leads to one more parameter for all models than when fitting just Experiment 1, and two more if the observer assumes same precision, S. We estimated ML parameters and compared models as described in the first section.

Results and Discussion

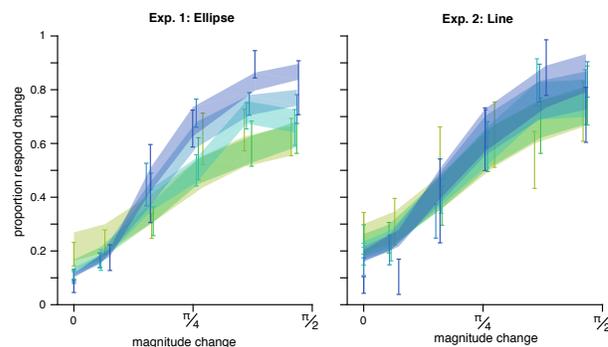


Figure 3: Joint model fits. $M \pm SEM$ of data (error bars) and model predictions (fill) across subjects for the VVO model for Exp. 1 (left) and Exp. 2 (right). Colors indicate how many high reliability ellipses are presented on each display (going from 0 in green to 4 in blue).

Again, we find the VVO model provides a good qualitative fit of the data (Figure 3). The model comparison results of the joint fit are also consistent with previous results. Participants are best fit by the VVO model, which assumes memory representations with variable precision, an observer who assumes variable precision encoding, and an optimal decision rule (Figure 4). This model fit better for every subject for almost all models; it fit better than VVM for 5 of 7 participants (worse by 17 and 19 AICc/BIC points for the remaining two participants).

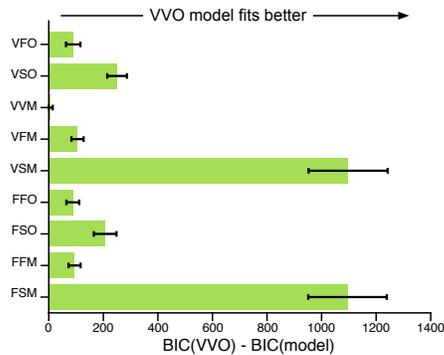


Figure 4: model comparison. $M \pm SEM$ BIC(VVO)-BIC(model) for all subjects (greater value indicates worse fit in comparison to the VVO model).

Conclusions

In this study, we investigated whether people can maintain and use uncertainty information in a VWM change detection task. We replicated the experiment and models finding that people can use uncertainty information optimally when that information is available to them at decision. We then modified the experimental paradigm such that this information was no longer available during the decision stage. Thus, uncertainty had to be maintained in VWM in order to be used.

We find that that people were for the most part best fit by a model that assumes variable precision in encoding, a variable precision assumption of encoding, and an optimal decision rule. For three (Exp.1) or two (both experiments jointly fit) of the subjects, a similar model with a max decision rule fit better. These results are consistent with previous findings that, in a variety of perceptual decision-making tasks including change detection, the optimal decision rule almost always fits data better than the max rule, except for the odd case in which both rules perform similarly well. (Ma, Shen, Dziugaite, & van den Berg, 2015).

While the results of the first experiment only indicate an optimal use of information, the results of the second experiment indicate an ability to maintain that uncertainty information over a delay. Overall, these results suggest that people maintain and use uncertainty information in working memory, and use them in tasks to maximize performance. This makes intuitive sense, because in naturalistic settings we are not always presented cues about how well our memories were maintained.

With decoding models of neural data, we have gained insight into how the brain encodes stimulus feature values (e.g., Harrison & Tong, 2009) and uncertainty (van Bergen, Ma, Pratte, & Jehee, 2015). Perhaps future research may be able to investigate how these representations are combined in the brain when doing decision-making tasks like change detection.

Acknowledgments

We thank Marissa Evans for collecting the majority of the data and Luigi Acerbi for being involved in previous iterations of this project. This work was funded by R01EY020958 to WJM.

References

- Acerbi, L., & Ma, J. W. (2017). *Practical bayesian optimization for model fitting with bayesian adaptive direct search* (Vol. 30). Advances in Neural Information Processing Systems.
- Bays, M. P., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854.
- Fougnie, D., Cormiea, M. S., Kanabar, A., & Alvarez, A. G. (2016). Strategic trade-offs between quantity and quality in working memory. *J of exp psychol: HPP*, *42*(8), 1231–1240.
- Fougnie, D., Suchow, W. J., & Alvarez, A. G. (2012). Variability in the quality of visual working memory. *Nature communications*, *3*.
- Harrison, A. S., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635.
- Keshvari, S., Berg, d. R. v., & Ma, J. W. (2012). Probabilistic computation in human perception under variability in encoding precision. *PLoS ONE*, *7*.
- Ma, J. W., Husain, M., & Bays, M. P. (2014). Changing concepts of working memory. *Nature neuroscience*, *17*, 347–356.
- Ma, J. W., Shen, S., Dziugaite, G., & van den Berg, R. (2015). Requiem for the max rule? *Vision research*, *116*, 179–193.
- Rademaker, L. R., Tredway, H. C., & Tong, F. (2012). Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision*, *12*, 21–21.
- van Bergen, S. R., Ma, J. W., Pratte, S. M., & Jehee, F. J. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*(18), 1728–1730.
- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, J. W. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *PNAS*, *109*, 8780–8785.
- Wilken, P., & Ma, J. W. (2004). A detection theory account of change detection. *Journal of Vision*, *4*, 1120–1135.
- Yoo, H. A., Klyszejko, Z., Curtis, E. C., & Ma, J. W. (2018). Strategic allocation of working memory resource. *Scientific reports*, *8*(1).
- Zhang, W., & Luck, J. S. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233–235.