

Setting and changing feature priorities in visual short-term memory

Zampeta Kalogeropoulou^{1,2} · Akshay V. Jagadeesh³ · Sven Ohl^{1,2} · Martin Rolfs^{1,2}

© Psychonomic Society, Inc. 2016

Abstract Many everyday tasks require prioritizing some visual features over competing ones, both during the selection from the rich sensory input and while maintaining information in visual short-term memory (VSTM). Here, we show that observers can change priorities in VSTM when, initially, they attended to a different feature. Observers reported from memory the orientation of one of two spatially interspersed groups of black and white gratings. Using colored pre-cues (presented before stimulus onset) and retro-cues (presented after stimulus offset) predicting the to-be-reported group, we manipulated observers' feature priorities independently during stimulus encoding and maintenance, respectively. Valid pre-cues reliably increased observers' performance (reduced guessing, increased report precision) as compared to neutral ones; invalid pre-cues had the opposite effect. Valid retro-cues also consistently improved performance (by reducing random guesses), even if the unexpected group suddenly became relevant (invalid-valid condition). Thus, feature-based attention can reshape priorities in VSTM protecting information that would otherwise be forgotten.

Electronic supplementary material The online version of this article (doi:10.3758/s13423-016-1094-6) contains supplementary material, which is available to authorized users.

✉ Zampeta Kalogeropoulou
zampeta@bccn-berlin.de

¹ Department of Psychology, Humboldt Universität zu Berlin, Berlin, Germany

² Bernstein Center for Computational Neuroscience, Humboldt Universität zu Berlin, Phillipstrasse 13, Haus 6, 10115 Berlin, Germany

³ Department of Psychology, Stanford University, Stanford, CA, USA

Keywords Feature-based attention · Visual short-term memory · Retro-cue · Mixture model

Introduction

We often rely on visual information that is no longer seen yet lingers in visual short-term memory (VSTM). While external attention (to locations, features, or objects) sets the priorities for which information will be encoded into VSTM in the first place (reviewed in Carrasco, 2011), internal attention can select information currently maintained therein (reviewed in Chun, Golomb, & Turk-Browne, 2011). In the present study, we shed new light onto the relation between these core selective mechanisms of human perception and cognition, providing evidence for feature-based selection in VSTM and revealing its independence of visual priorities during sensory encoding.

Many studies have approached the relationship between attention and memory using cueing procedures. A cue to a location, for instance, presented shortly after the offset of a visual array, retains memory of stimuli that had occupied the cued location (Sperling, 1960). Indeed, spatial retro-cues can generate behavioral results equivalent to those of pre-cues (Griffin & Nobre, 2003; Landman, Spekreijse, & Lamme, 2003), even when presented seconds after the disappearance of the stimulus array (e.g., Astle, Summerfield, Griffin, & Nobre, 2012; Sligte, Scholte, & Lamme, 2008).

Non-spatial retro-cues may also result in benefits for memory performance (Heuer & Schubö, 2016; Lepsien & Nobre, 2007; Li & Saiki, 2015; Pertzov, Bays, Joseph, & Husain, 2013), suggesting that attention can influence VSTM maintenance across many flavors of attention. Most of these studies (Heuer & Schubö, 2016; Li & Saiki, 2015; Pertzov et al., 2013) have shown that colored retro-cues can improve memory for items of the same color. However, these studies

invariably involved a spatial component, as the color cue identified isolated objects at particular locations. It remains unclear, therefore, whether selection in memory was based on spatial attention, drawn to the location of an object of a particular color (cf. Pertzov et al., 2013), or, alternatively, whether feature-based attention (FBA) – independently of space – affects memory maintenance. Moreover, few studies (if any) have addressed the interplay of external and internal attention.

We devised a novel double-cueing paradigm to fill these gaps. Observers reported from memory the orientation of one of two spatially interspersed groups of gratings (one white, one black). Using colored pre-cues (presented before stimulus onset) and retro-cues (presented after stimulus offset), predicting the group whose orientation would have to be reported, we manipulated observers' FBA independently during stimulus encoding and maintenance, respectively. This strategy allowed us to address the relationship between FBA and VSTM when selection by location was explicitly ruled out.

Methods

Participants

Based on the sample size in previous studies on FBA and VSTM (e.g., Herrmann, Heeger, & Carrasco, 2012; van Moorselaar, Olivers, Theeuwes, Lamme, & Sligte, 2015), we recruited ten observers (aged 23–30 years; three females, two left-eye dominant, all right-handed) for participation in the experiment. All observers were naïve regarding the purpose of the study and participated in exchange for a fixed monetary payment of 7€ per session (1 h each), and a bonus of 7€ after completion of all six sessions. The Ethics Committee of the German Society for Psychology (DGPs) approved the study and observers provided written consent before participation. The study was conducted in accordance with the Declaration of Helsinki.

Experimental setup

Observers sat in a dimly lit room with their head positioned on a chin rest, at a distance of 57 cm from a gamma linearized 22-in Sony GDM-F520 screen (1,280 × 960 pixels, 100 Hz vertical refresh rate). An Eyelink 1000 Desktop Mount (SR Research, Ottawa, ON, Canada) recorded the position of the dominant eye at 1,000 Hz. For stimulus presentation, online gaze control, and response collection, we used MATLAB (Mathworks, Natick, MA, USA), including the Psychophysics toolbox (Brainard, 1997; Pelli, 1997), the Eyelink toolbox (Cornelissen, Peters, & Palmer, 2002) running on a Mac mini (Apple Inc., Cupertino, CA, USA). Manual reports were collected using a PowerMate 3.0 (Griffin Technology, Nashville, TN, USA).

Stimuli and procedure

During each trial (Fig. 1a) observers had to maintain fixation at a centrally presented red filled circle (diameter 0.2° of visual angle, dva). After 500 ms of fixation, a pre-cue appeared for 500 ms (a ring with a diameter of 1.0 dva, width of 0.12 dva, and color defined by experimental condition), either informing observers about the color of the stimulus that would most likely be probed in the later memory test (black or white; 67 % validity), or not providing any prior information (red; neutral cue). After an inter-stimulus-interval (ISI) of 500 ms, the memory array, consisting of two sets of differently colored and oriented gratings (one black, one white), was briefly flashed for 150 ms. Observers were asked to keep the orientations in memory and report the orientation of one set of gratings at the end of the trial. Following an ISI of 1,000 ms, a retro-cue appeared for 500 ms (same visual features as the pre-cue), telling observers in 75 % of the trials (valid retro-cue) which set of gratings would be probed in the memory test (100 % validity). In the remaining 25 % of trials, a neutral retro-cue was uninformative with respect to the memory test and observers could rely only on the information provided by the pre-cue. Finally, following an ISI of 1,000 ms, one set of colored gratings (the positions of which did not coincide with the positions occupied in the memory array; see below) was presented and observers were asked to rotate (using a volume knob) the orientation of the gratings (initial orientation drawn from a uniform distribution ranging from 0° to 180°) until they matched the orientation of the stimulus kept in memory. Finally, by pressing the knob, the response was saved and participants received visual feedback about their error (in the range of 0–90°). Throughout the trial, we controlled fixation behavior using an online eye-tracking routine that aborted a trial automatically as soon as the distance between eye position and fixation point was larger than 2 dva. Observers repeated these aborted trials at the end of a block.

The memory array consisted of 32 (16 black and 16 white) gratings (spatial frequency = 3.8 cycles per dva) arranged on a 9 × 9 dva² regular grid centered at fixation (inspired by the stimuli used in Herrmann et al., 2012). Each grating was the positive (white) or the negative (black) lobe of a sine wave grating, fading out in a Gaussian window (SD = 0.16°). The minimum distance between the center of the screen and the center of its closest neighboring grating was 2.7 dva. We randomly jittered (uniform distributions, ranging from 0 to 0.6 dva) the horizontal and vertical position of each grating in each interval and trial. This spatial layout and small size of the stimuli was used to require observers to allocate attention globally over the whole visual display, to encode the orientation of the full set of gratings – and not the orientation of a single one (cf. Herrmann et al., 2012). The orientations of the target set and the non-target set of gratings varied independently with a distance drawn from a uniform distribution in the range from 60° to 120°.

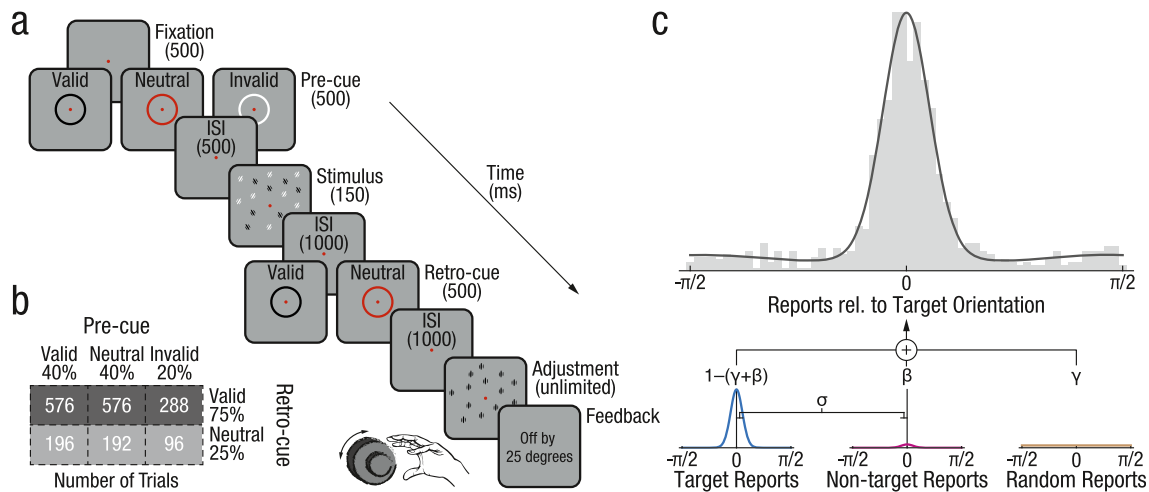


Fig. 1 (a) Trial sequence in the double-cueing paradigm. Observers fixated a red fixation point in the center of the screen. After 500 ms, a pre-cue was presented that was either informative (black or white; 67 % validity) or neutral (red). After an inter-stimulus-interval (ISI) of 500 ms, a memory array – consisting of two differently colored (black and white) and oriented sets of gratings – appeared for 150 ms. 1,000 ms after the disappearance of the memory array, a retro-cue (valid or neutral) appeared for 500 ms, followed by another retention interval of 1,000 ms. Finally, the probe array (either black or white; initial orientation drawn from a uniform distribution, ranging from 0° to 180°) appeared and observers were asked to adjust the orientation of gratings with a knob such that it matches the orientation of the same-colored group of gratings in the target array. By pressing the knob, the response was saved and participants

Each observer was tested in six experimental sessions (each session consisting of eight blocks with 40 trials per block) contributing a total of 1,920 trials to the final analysis (see Fig. 1b for trial numbers per condition), such that the least likely condition (invalid pre-cue and neutral retro-cue) had the minimum number of 96 reports.

Data analysis

Formal model comparison (see Supplemental Online Material) revealed that a mixture model (proposed by Bays, Catalao, & Husain, 2009) yielded the best description of individual report distributions across conditions. This model comprises target reports, non-target reports, and random guesses:

$$p(\theta) = (1-\gamma-\beta)\Phi_{\mu,\kappa} + \beta\Phi_{\delta,\kappa} + \gamma\frac{1}{2\pi},$$

where θ is the difference between the reported value and the target value, Φ is a von Mises probability density function (i.e., the circular analogue of a Gaussian) with mean μ (target orientation) or δ (non-target orientation) and the spread parameter κ (which relates to the standard deviation, $\sigma = \sqrt{1/\kappa}$). γ and β are weight coefficients, corresponding to the proportion of random guesses and non-target reports, respectively. Fits of this model to the report of each observer in each condition provided the basis for the results reported in this manuscript.

received feedback about their error. Note that informative cues (valid and invalid) could be either black or white. In the example shown here, the black cue correctly predicts the target group and is thus valid; the white cue incorrectly predicts the non-target group (invalid). (b) Number of trials in each experimental condition. (c) Example distribution of orientation reports, including the fit of a mixture model. The model comprised random guesses and non-target reports in addition to target reports, each associated with a certain probability (γ , β , and $1-(\gamma+\beta)$, respectively). Target and non-target reports were assumed to have the same precision, captured by its inverse, the standard deviation (σ). Note that in the overall fit (top panel), the non-target reports are distributed across a range of angles, due to the variable relation between target and non-target orientations

Note that for orientation reports the range of unique angles is 0°–180° (rather than 0°–360°); thus, we mapped orientation angles onto radians using the factor $\pi/90^\circ$ prior to fitting.

Statistics

To assess the impact of pre-cueing and retro-cueing on memory performance, we conducted two-way repeated-measures analyses of variance (rmANOVA) on each dependent variable. As proportions (of non-target or random reports) are not usually normally distributed, we used arcsine-transformed data in all rmANOVAs on these dependent variables. To further assess the evidence for the observed effects (or their absence), we computed Bayes Factors (BF), complementing each rmANOVA. BFs express the posterior probability for a model given the data (fixed effect of factor + random effect of participants) relative to a model including only participants as a random factor. BFs were computed using the R package BayesFactor (Morey & Rouder, 2015).

We calculated 95 % confidence intervals ($CI_{95\%}$) using a bootstrap technique (Efron & Tibshirani, 1993). To generate each of the 1,000 bootstrap data sets, we first resampled the number of reports, N, that each participant provided in a given condition exactly N times (with replacement), fitted the model, and obtained the parameters of the best fit. Then, we averaged these bootstrapped parameters across the ten subjects,

resulting in one bootstrap data set. We repeated this procedure 1,000 times and obtained 2.5 % and 97.5 % percentiles of a given dependent variable across the 1,000 bootstrap data sets. These 1,000 bootstrap data sets were also used to compute the $CI_{95\%}$ of differences between two conditions by simply subtracting the parameter values of the first condition's 1,000 bootstrap data sets from those of the second condition, and determining the 2.5 % and 97.5 % percentiles of the resulting distribution of differences.

Results

In our double-cueing paradigm, the pre-cue set observers' priorities during the initial encoding of the stimulus. It would either direct observers' attention correctly to the target orientation (valid condition), incorrectly to the non-target orientation (invalid condition), or to both orientations at the same time (neutral condition). If participants were able to change their priorities after the stimulus had disappeared, then a valid retro-cue should increase performance (relative to the neutral retro-cue) even if, initially, the wrong orientation had been attended. Indeed, the absence of an interaction between pre-cueing and retro-cueing would provide strong support for this hypothesis. We assessed this prediction on a number of dependent variables.

As a first step in our analysis, we calculated the mean angular error of observers' orientation reports in each condition (Fig. 2a). To assess the effect of the pre-cue and the retro-cue on this global measure of memory performance, we

conducted a two-way rmANOVA. We observed main effects for both types of cueing. Both during encoding (pre-cue: $F(2, 18) = 13.98$, $p < 0.001$; $BF = 18182$) and during maintenance (retro-cue: $F(1, 9) = 45.47$, $p < 0.001$; $BF = 7.9$), directing observers' FBA significantly affected report errors, and these main effects were largely independent of one another (interaction: $F(2, 18) = 0.14$, $p > 0.250$; $BF_{\text{Main effects}} = 2677153$ vs. $BF_{\text{Main effects} + \text{Interaction}} = 597676$, $BF\text{-Ratio} = 4.48$). Here, the ratio of BFs ($BF\text{-Ratio}$) indicates that a model including the two main effects only is preferable by a factor of 4.48 to a model containing an additional interaction of the pre-cue and the retro-cue. As compared to performance in the neutral pre-cue conditions, reports were more accurate following valid pre-cues (Mean difference, $\Delta = 2.88^\circ$, $CI_{95\%}$ 2.25 to 3.57) and less accurate following invalid pre-cues ($\Delta = -2.73^\circ$, $CI_{95\%}$ -3.68 to -1.84). Moreover, the presentation of a valid retro-cue further decreased the error in orientation reports ($\Delta = -2.69^\circ$ across all pre-cue conditions, $CI_{95\%}$ -3.31 to -2.10).

Next, to identify potential sources of the observed errors and their differences across conditions, we modeled orientation report distributions with a three-parameter mixture model (Bays et al., 2009; see Methods and Fig. 1c). For each individual and each condition, we obtained estimates for the standard deviation of target reports (σ) around the true stimulus orientation, as well as for the proportion of non-target reports (δ) and the proportion of random guesses (γ).

The model estimates for the standard deviation σ (Fig. 2b) provide a measure of the precision of the memory representation (the inverse of the standard deviation, $1/\sigma$). A rmANOVA revealed a significant effect of the pre-cue ($F(2, 18) = 8.26$, $p = 0.003$; $BF = 27.39$), but no effect of the retro-cue (valid vs. neutral: $\Delta\sigma = 0.44^\circ$, $CI_{95\%}$ -0.81 to 1.39 ; $F(1, 9) = 0.6$, $p > 0.250$; $BF = 0.36$) on σ . Importantly, there was no interaction between the pre-cue and retro-cue ($F(2, 18) = 0.04$, $p > 0.250$; $BF_{\text{Main effects}} = 10.9$ vs. $BF_{\text{Main effects} + \text{Interaction}} = 2.3$, $BF\text{-Ratio} = 4.74$). As compared to performance in the neutral pre-cue conditions, reports had a smaller σ (i.e., they were more precise) following valid pre-cues ($\Delta\sigma = -1.12^\circ$, $CI_{95\%}$ -0.48 to -1.80) and higher σ following invalid pre-cues ($\Delta\sigma = 1.56^\circ$, $CI_{95\%}$ 3.53 to 0.25). Thus, memory precision for orientation is limited by stimulus encoding, and does not improve (nor deteriorate) following a retrospective deployment of FBA to the memory representation.

Both pre-cue and retro-cue strongly affected the proportion of random guesses γ (Fig. 2c). As compared to neutral pre-cues, the proportion of random guesses was lower for valid ($\Delta\gamma = -0.05$, $CI_{95\%}$ -0.02 to -0.08) and higher for invalid pre-cues ($\Delta\gamma = 0.06$, $CI_{95\%}$ 0.02 to 0.10). The valid retro-cue further decreased the proportion of random guesses ($\Delta\gamma = -0.06$, $CI_{95\%}$ -0.03 to -0.09), irrespective of pre-cue type. A rmANOVA corroborated these results, yielding a significant main effect of pre-cue ($F(2, 18) = 4.83$, $p = 0.020$; $BF = 11.74$) and retro-cue ($F(1, 9) = 11.11$, $p = 0.009$; $BF = 2.22$). The

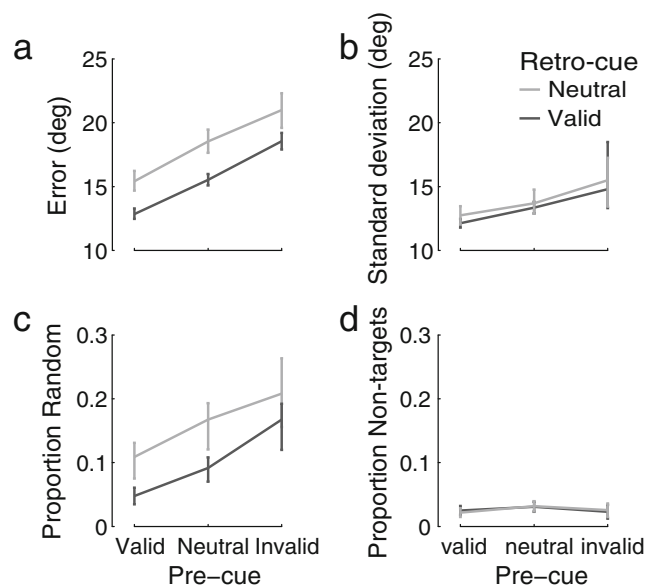


Fig. 2 Memory performance, averaged across observers, as a function of the pre-cue and retro-cue conditions. **(a)** Mean angular error of the orientation reports. **(b)** Standard deviation of target (and non-target) reports, σ . **(c)** Proportion of random guesses, γ . **(d)** Proportion of non-target reports, β . The error bars are bootstrapped $CI_{95\%}$

interaction between pre-cue and retro-cue again was not significant ($F(2, 18) = 0.71, p > 0.250$; $BF_{\text{Main effects}} = 43.7$ vs. $BF_{\text{Main effects} + \text{Interaction}} = 12.1$, BF-Ratio: 3.61). Thus, participants systematically rely on the information provided by both pre-cue and retro-cue in maintaining the memory of the orientations. More importantly, a switch in the priorities during maintenance (invalid pre-cue, valid retro-cue) leads to an improvement in performance that is as large as the retro-cueing effect in the valid or neutral pre-cue conditions.

Finally, the proportion of non-target reports β was consistently very small ($\delta \approx 0.03$) across all conditions (Fig. 2d). A rmANOVA revealed no main effects (pre-cue: $F(2, 18) = 2.57, p = 0.10$; $BF = 0.9$; retro-cue: $F(1, 9) = 0.75, p > 0.250$; $BF = 0.298$), and no significant interaction ($F(2, 18) = 0.16, p > 0.250$; $BF_{\text{Main effects}} = 0.3$ vs. $BF_{\text{Main effects} + \text{Interaction}} = 0.072$, BF-Ratio = 4.14).

In summary, valid pre-cues strongly reduced the mean angular error of observers' reports by increasing the precision of target reports and decreasing the proportion of random guesses. Valid retro-cues further reduced mean angular errors by decreasing the proportion of random guesses. For none of our four dependent variables did we see evidence for an interaction between pre-cues and retro-cues. In fact, statistical models comprising only the two main effects were favored over models with an additional interaction term by a factor of about 4.

Discussion

We demonstrated that the deployment of FBA – much like spatial attention – is a powerful selection mechanism that not only increases observers' sensitivity during sensory processing of visual information (reviewed in Carrasco, 2011), but also fortifies feature information held in VSTM, once the stimulus has disappeared from view. Whereas the initial encoding of a stimulus appears to set an upper limit for the precision of a memory representation, the deployment of FBA during maintenance effectively increases the likelihood that observers remember the target orientation (reduced guess rate). Our results suggest further that FBA prioritizes visual information during the maintenance of memoranda independently of the priorities the observer imposed during sensory encoding: reinforcing these initial priorities as well as changing priorities to a new feature in VSTM benefitted observers' performance.

Whereas spatial retro-cues are known to strongly affect the content of visual memory (Astle et al., 2012; Griffin & Nobre, 2003; Landman et al., 2003; Makovski, Sussman, & Jiang, 2008; Murray, Nobre, Clark, Cravo, & Stokes, 2013; Rerko & Oberauer, 2013; Rerko, Souza, & Oberauer, 2014; Slight et al., 2008), few studies have addressed the role of FBA during VSTM maintenance. In these studies, colored retro-cues identified one of several colored objects held in memory

and participants reproduced their orientation (Heuer & Schubö, 2016; Li & Saiki, 2015; Pertzov et al., 2013). The observed advantage of valid retro-cues was comparable to that of spatial retro-cues in terms of its time course and the size of the effect (Pertzov et al., 2013), although it had a different spatial profile (Heuer & Schubö, 2016). Nevertheless, observers may have used the feature retro-cue to guide *spatial* attention to the remembered item's location (cf. Pertzov et al., 2013). We meticulously designed our stimuli to pre-empt any spatial strategy to solve the task. Observers deployed attention globally to the orientation of one group of gratings, spatially interspersed with gratings of an uncued orientation, and displayed at sizable eccentricity. We thus isolated robust effects of FBA on the content of VSTM, despite a potentially privileged role of location in the maintenance of information in VSTM (Pertzov & Husain, 2014).

The literature on spatial attention in VSTM (see overview in Souza & Oberauer, 2016) may provide inspiration for the mechanisms underlying feature-based retro-cueing. One set of accounts suggests that retro-cues protect visual information from time-based decay or interference. Visual memory traces may become diffuse over time (Wimmer, Nykamp, Constantinidis, & Compte, 2014) or merge with neighbouring ones (Wei, Wang, & Wang, 2012). This view predicts a continuous decrease of the fidelity of memory representations and thus reduced precision in memory reports. A valid retro-cue could stop this diffusion, or reduce interference with distractor orientations (Pertzov et al., 2013). Alternatively, the entire representation may fade out (Wei et al., 2012), be suddenly lost (Zhang & Luck, 2009), or no longer available for recall, unless selected by retrospective attention (Murray et al., 2013). A proper evaluation of these scenarios calls for a manipulation of set-size or incentives to trade off quality and quantity of memory representations (Fougnie, Cormiea, Kanabar, & Alvarez, 2016). Moreover, if the retro-cue protected information from decay then memory performance should be the same in trials with valid retro-cues and trials with no retro-cue in which the report is required at the time the valid retro cue would normally appear (cf. Makovski et al., 2008; Pertzov et al., 2013).

Retro-cues may also protect memory representations from visual interference caused by the probe array (Makovski et al., 2008) or interference from a secondary task (Makovski & Pertzov, 2015). To account for the observed lack of interactions between pre-cues and retro-cues, this would imply that perceptual interference from the probe array affected all memory representations in a similar way, independent of the priority set by the pre-cue.

Finally, recent results show that repeated retro-cueing during the retention interval resulted in better memory performance for the repeatedly cued item as compared to single cues, suggesting that spatial retro-cues can strengthen memory representations (Rerko & Oberauer, 2013). Here, we observed a similar result: a valid retro-cues further added to the effect of

the pre-cue. In contrast to the study on spatial retro-cueing, we presented cues before and after the memory array, suggesting that a combination of pre- and retro-cues can equally strengthen representations in a cumulative fashion.

In summary, we have demonstrated the visual system's ability to flexibly re-set priorities for perception and memory. While future studies need to disentangle the mechanisms giving rise to these results, they highlight the role of FBA beyond initial sensory stages and add to the accumulating evidence for the versatility of fundamental selection processes in human vision and cognition.

Acknowledgments This research was supported by the Deutsche Forschungsgemeinschaft (DFG; grants RO3579/2–1, OH274/2-1, RO3579/6-1), GRK 1589/2 "Sensory Computation in Neural Systems" and a DAAD RISE fellowship for undergraduates to A.Y.

Author Contributions Z. Kalogeropoulou, S. Ohl, and M. Rolfs developed the study concept and the study design. Testing and data collection were performed by A. Jagadeesh. Z. Kalogeropoulou performed the data analysis and interpretation under the supervision of S. Ohl and M. Rolfs. Z. Kalogeropoulou drafted the manuscript, and S. Ohl and M. Rolfs provided critical revisions. All authors approved the final version of the manuscript for submission.

Compliance with ethical standards

Conflict of Interest The authors declare that they have no conflicts of interest with respect to their authorship or the publication of this article.

References

- Astle, D. E., Summerfield, J., Griffin, I., & Nobre, A. C. (2012). Orienting attention to locations in mental representations. *Attention, Perception, & Psychophysics*, *74*(1), 146–162.
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7.1–11.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*(4), 433–436.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525.
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*(1), 73–101.
- Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman and Hall/CRC.
- Fougnie, D., Cormiea, S. M., Kanabar, A., & Alvarez, G. A. (2016). Strategic trade-offs between quantity and quality in working memory. *Journal of Experimental Psychology: Human Perception and Performance*. doi:10.1037/xhp0000211
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194.
- Herrmann, K., Heeger, D. J., & Carrasco, M. (2012). Feature-based attention enhances performance by increasing response gain. *Vision Research*, *74*, 10–20.
- Heuer, A., & Schubö, A. (2016). Feature-based and spatial attentional selection in visual working memory. *Memory & Cognition*, *44*(4), 621–632. doi:10.3758/s13421-015-0584-5
- Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*(2), 149–164.
- Lepsien, J., & Nobre, A. C. (2007). Attentional modulation of object representations in working memory. *Cerebral Cortex*, *17*(9), 2072–2083.
- Li, Q., & Saiki, J. (2015). Different effects of color-based and location-based selection on visual working memory. *Attention, Perception, & Psychophysics*, *77*, 450–463.
- Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *The Quarterly Journal of Experimental Psychology*, *68*(9), 1735–1743.
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 369–380.
- Morey, R. D., & Rouder, J. N. (2015). BayesFactor: Computation of Bayes Factors for Common Designs. R package version 0.9.4. Retrieved from <http://CRAN.R-project.org/package=BayesFactor>
- Murray, A. M., Nobre, A. C., Clark, I. A., Cravo, A. M., & Stokes, M. G. (2013). Attention restores discrete items to visual short-term memory. *Psychological Science*, *24*(4), 550–556.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
- Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1224–1231.
- Pertzov, Y., & Husain, M. (2014). The privileged role of location in visual working memory. *Attention, Perception, & Psychophysics*, *76*(7), 1914–1924.
- Reerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1075–1096.
- Reerko, L., Souza, A. S., & Oberauer, K. (2014). Retro-cue benefits in working memory without sustained focal attention. *Memory & Cognition*, *42*(5), 712–728.
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *PLoS ONE*, *3*(2), e1699.
- Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*. doi:10.3758/s13414-016-1108-5
- Sperling, G. (1960). The information available in brief visual presentations. *Science*, *345*(6193), 212–215.
- van Moorselaar, D., Olivers, C. N. L., Theeuwes, J., Lamme, V. A. F., & Sligte, I. G. (2015). Forgotten but not gone: Retro-cue costs and benefits in a double-cueing paradigm suggest multiple states in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(6), 1755–1763.
- Wei, Z., Wang, X.-J., & Wang, D.-H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience*, *32*(33), 11228–11240.
- Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience*, *17*(3), 431–439.
- Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*(4), 423–428.