Trade-off between Capacity and Precision in Visuospatial Working Memory

Chantal Roggeman[1], Torkel Klingberg[1], Heleen E. M. Feenstra[1], Albert Compte[2], and Rita Almeida[1,2]

Abstract

■ Limitations in the performance of working memory (WM) tasks have been characterized in terms of the number of items retained (capacity) and in terms of the precision with which the information is retained. The neural mechanisms behind these limitations are still unclear. Here we used a biological constrained computational model to study the capacity and precision of visuospatial WM. The model consists of two connected networks of spiking neurons. One network is responsible for storage of information. The other provides a nonselective excitatory input to the storage network. Simulations showed that this excitation boost could temporarily increase storage capacity but also predicted that this would be associated with a decrease in precision of the memory. This prediction was subsequently tested in a behavioral (38 participants) and fMRI (22 participants) experiment. The behavioral results confirmed the trade-off effect, and the fMRI results suggest that a frontal region might be engaged in the trial-by-trial control of WM performance. The average effects were small, but individuals differed in the amount of trade-off, and these differences correlated with the frontal activation. These results support a two-module model of WM where performance is determined both by storage capacity and by top–down influence, which can vary on a trial-by-trial basis, affecting both the capacity and precision of WM. ■

INTRODUCTION

Working memory (WM) is the ability to keep information in mind during a short period of time and is a fundamental component of many cognitive functions (Conway, Kane, & Engle, 2003; Cowan, 2001; Baddeley, 1986). Limitations in the amount of information that one can retain is a key characteristic of WM, determining individual differences in a wide range of cognitive tasks (Cowan, 2001, 2010; Kyllonen & Christal, 1990). These limitations have been described in terms of the maximum number of items about which one can simultaneously retain information (capacity) and in terms of the exactness of the remembered information (precision). A great amount of experimental and theoretical work has dealt with characterizing WM, in particular the limitations in visual WM. One theory, the "slots model," suggests that there is a discrete fixed number of items or "slots" that each individual can remember (Anderson, Vogel, & Awh, 2011; Zhang & Luck, 2008; Luck & Vogel, 1997). Another theory, the "resources model" (van den Berg, Shin, Chou, George, & Ma, 2012; Huang, 2010; Bays, Catalao, & Husain, 2009; Bays & Husain, 2008; Wilken & Ma, 2004), suggests that WM is better described as a resource that can be divided among the objects to retain in memory, so that the precision of the memories decreases with the number of items, without an upper limit on this number. Finally, others (Buschman, Siegel, Roy, & Miller, 2011; Xu & Chun, 2006; Alvarez & Cavanagh, 2004) have proposed hybrids of these two more extreme theoretical models suggesting that WM resources can be divided among items, that this sharing of resources has an impact on the precision of memories, but that there is still an upper limit of items that can be held in WM. Thus, the relation between capacity and precision as well as the biological origins of these limitations remain unclear.

Mechanistic hypotheses for visuospatial WM (vsWM) are specified in a computational network model in which short-term visuospatial memories are stored by selective sustained elevated neuronal activity (Edin et al., 2009; Compte, Brunel, Goldman-Rakic, & Wang, 2000). The model consists of a network of integrate-and-fire excitatory and inhibitory neurons connected so that selective activity can be maintained after stimulus offset. Rather than modeling WM in a single storage area, Edin et al. (2009) proposed a model with two modules: a storage area (e.g., parietal cortex) and a boost area (e.g., pFC) responsible for an excitatory nonselective input evenly received by the cells of the storage network. Analysis of this model suggested that the storage capacity can be increased by a top–down excitatory input or "boost signal" to the storage area (Edin et al., 2009). This hypothesis was shown to be consistent with behavioral and fMRI data: Interindividual

[1]Karolinska Institute, Stockholm, Sweden, [2]Institut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Barcelona, Spain

Journal of Cognitive Neuroscience 26:2, pp. 211–222
doi:10.1162/jocn_a_00485

differences in WM capacity correlated with the amount of activation of a prefrontal boosting area and with fronto-parietal functional connectivity (Edin et al., 2009). Other works have also addressed context-dependent adjustments of performance on visual WM tasks, some showing that memory precision can be regulated for a small number of items (Machizawa, Goh, & Driver, 2012; Zhang & Luck, 2008) whereas others failed to show regulation for higher number of items (Zhang & Luck, 2011). In particular, it has been recently shown (Machizawa et al., 2012) that participants are able to adjust at the trial-by-trial level the precision with which they store a few items in WM. Here, we hypothesized that capacity boosting (Edin et al., 2009) and precision enhancement (Machizawa et al., 2012) could share a common mechanism in the top–down boost signal proposed by Edin et al. (2009). We further hypothesized that an increase of capacity through a boost signal could be engaged with the time resolution of single trials. This would imply that manipulations of precision would affect WM capacity, and conversely similar trial-by-trial adjustments of capacity should result in changes of WM precision. We first tested these hypotheses in a computational model and derived a concrete prediction about how manipulating vsWM capacity through a boost signal should affect the precision of the memories. To test the prediction, we designed a behavioral protocol based on the assumption that a cue about the number of items (load) to be expected in a given trial could result in selective engagement of a boost signal on a trial-by-trial basis. This behavioral protocol was used in one behavioral and one fMRI experiment to validate the model's prediction. We thus establish a within-trial interdependence of capacity and precision in vsWM, and we propose a plausible mechanistic substrate based on a boost, possibly top–down, input to a storage area. Our computational model then describes a possible biological mechanism behind a trade-off effect in WM where more items can be stored at a cost of reduced precision. Such a trade-off is hypothesized by a resource view of WM and appears as a prediction in our biologically constrained model, which however has an upper limit on the number of memories that can be retained. Hence, our model offers a biologically plausible implementation of a hybrid WM model, thus helping to reconcile alternative views of WM as a resource or a limited number of slots.

METHODS

Model

The computational model consists of two connected networks of spiking neurons: one corresponding to a storage area and one corresponding to a boost area responsible for a boost signal to the storage area (Figure 1). The neurons in the storage area are modeled to encode positions (in angle) on a circle. Presentation of an item at a given angle is simulated by an input to the corresponding cells. Connections between neurons are tuned so that neurons encoding similar angles have stronger connections than neurons encoding far apart angles. This tuning allows

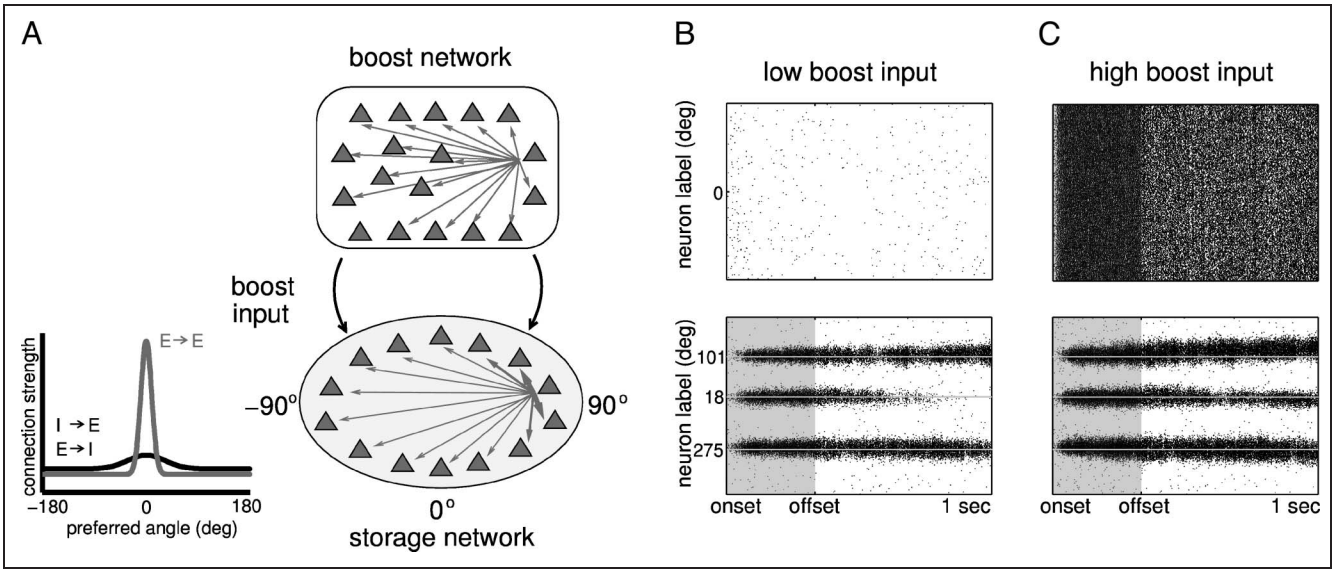

Figure 1. Computational model. (A) Schematic representation of the computational model formed by a boost and a storage network. The ring organization of the storage network and the connectivity structure between its excitatory neurons (represented as triangles) is illustrated. The profile of connectivity strengths for the storage network is shown in the plot on the left. (B and C) Example of the neuronal activity of excitatory neurons for a simulation with a stimulus with Load 3 (three positions to keep in memory). Each dot represents an action potential. Stimuli presentation of 500 msec, indicated with a gray background. (B) Simulation with low input from the boost network (top) to the storage network (bottom): The capacity of the model is lower (some items are forgotten), but precision is higher (the positions remembered are closer to the target positions). (C) Simulation with high input from the boost network (top) to the storage network (bottom): The capacity of the model is higher (no items are forgotten), but precision is lower (the memory traces deviate more from the target locations, and thus, items are remembered with less precision).

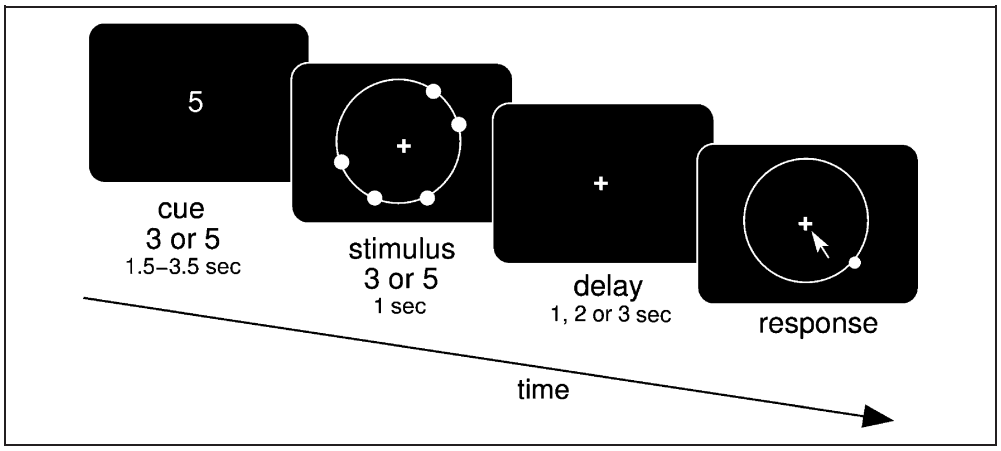

Figure 2. Schematic representation of the sequence of screens in the experimental procedure of the behavioral experiment.

the network to maintain localized bumps of high activity upon stimulation of a given angle. Furthermore, it allows that these localized bumps are maintained after stimulus offset, hence corresponding to maintenance of WM.

For simplicity, the boost network is equal to the storage network with the exception that the connectivity between neurons is not tuned. Because of this, the network is not able to sustain localized bumps of high activity and hence selective memory. The boost network is either globally activated and giving a nonspecific boost signal to the storage area or globally inactivated. The boost signal increases the capacity of the storage network (Edin et al., 2009). A mathematical description of the model and parameters used are given in the Appendix.

Each model simulation started with 100 msec of baseline activity, followed by onset of stimulus-specific stimulation for 500 msec. For trials with high boost, the excitatory neurons in the boost network received an external unspecific input for 100 msec at stimulus onset. This input caused the boost network to enter a state of global persistent high activity. Figure 1 shows examples of activity in the boost and storage networks in the case of low input (Figure 1B) and high input (Figure 1C) from the boost area.

To study the effect of the boost signal to the storage network, we ran 300 simulations for Loads 3 and 4, as the model was originally tuned to. The stimulations were made to model items positioned at random positions around a circle, with the only restriction that they could not be closer than 33°. The results were analyzed in terms of capacity and precision. Capacity was measured as the relative number of memories that persisted 1 sec after stimulus offset. Precision was measured as the inverse of the standard deviation across trials of the location of memories that persisted after 1 sec. The location of the memories was read out from the population vector (Lee, Reis, Seung, & Tank, 1997; Georgopoulos, Schwartz, & Kettner, 1986) of the neurons belonging to a bump of activity, using action potentials from 900 msec to 1 sec after stimulus offset. The memories were considered to persist if the norm of the population vector was higher than 25. The population vector is an average of the preferred direction of the neurons, weighted by the number of action potentials they emit. To determine if a neuron was contributing to a bump, we used a recursive algorithm. We started by attributing each neuron to one and only one memory, that of the item closest to its selectivity. Then we calculated the population vector for each memory bump using the neurons attributed to that memory that had preferred direction less than 35° away from the stimulus. The population vectors give estimations of the centers of the memory bumps. These estimations were refined repeating the procedure described above with the differences that neurons were now attributed to the nearest center of memory bump and that a new population vector was now calculated using neurons that had preferred directions less than 12° away from the estimated center of the memory bump. Further repeating this recursive procedure did not change the estimated centers of the memory traces.

Behavioral Experiment

Participants

Thirty-eight volunteers (two left-handed; average age and standard deviation, 25 ± 5 years) participated in the experiment. All had normal or corrected-to-normal vision. The study was approved by the ethical review board of Karolinska Institutet, Sweden, and all participants gave written informed consent before the experiment.

Apparatus

Stimuli were presented on a computer screen, and the experimental procedure was controlled by Eprime software (version 2.0, Psychology Software Tools, Inc., Sharpsburg, PA). Responses were given using a standard mouse. Participants were in a quiet room looking at a screen positioned 60 cm away at eye level.

Experimental Paradigm

The experimental task was a visuospatial task, requiring the participants to recall the positions of three or five items simultaneously presented along a circle (Figure 2). A trial

consisted of the following sequence: First, a blue fixation cross appeared in the middle of a black screen for 500 msec. Then, a cue indicating the number of items to be expected on the current trial was presented for a variable time (1500, 2000, 2500, 3000, or 3500 msec, counterbalanced over conditions). Immediately after the cue, the stimulus was presented for 1000 msec. Then, there was a variable delay (1000, 2000, or 3000 msec) during which only the fixation cross remained on the screen. Finally, the response screen appeared, and the participants indicated the remembered locations of the items with the mouse. Participants were asked to fixate on the cross during the trial. There was an intertrial interval of 2000 msec, during which the screen was completely black.

The cue was a number (3 or 5) in blue indicating how many items should be expected in the current trial. The stimulus consisted of three or five items (blue dots with radius of 1.85 visual degrees), randomly positioned on a visible blue circumference (radius of 25.5 visual degrees) centered on the visible fixation cross. The minimum angle between two neighboring items was 45°. On the response screen, the circumference and the fixation cross remained on the screen. A small red indication mark was displayed at a random location on the circle as well as the mouse pointer. Participants were required to click on the remembered positions of the items, in a counter clockwise order, starting from a red indication mark. The response phase was terminated by a right mouse button click.

An important point is that the cue did not correspond to the actual number of items presented on all trials but only matched on a fraction of trials (true cue trials). On the rest of trials (less than 23%), the cue did not match the number of items (false cue trials). Additionally, for some participants a trial type with Cue 0 was inserted: In these trials, the stimulus presentation was the same (three or five items), but participants did not have to remember them: In the response screen, the items were shown again and the participant had to click the indicated positions. These Cue 0 trials were not analyzed.

The experiment consisted of at least 150 Cue 3/Load 3 trials, at least 150 Cue 5/Load 5 trials, 45 Cue 3/Load 5 trials, 45 Cue 5/Load 3 trials, and 60 Cue 0 trials (for participants who had those trials). The experiment was divided in a minimum of eight blocks of 50 trials and took around 90 min.

Data Analysis

Each response (mouse-click) shows some deviation from the original true item location. This deviation was modeled using a probabilistic model previously introduced to account for performance on a recall task where both stimuli and responses are chosen from a circular parameter space (Bays et al., 2009; Zhang & Luck, 2008). The model assumes that the experimental distribution of errors $f_{\text{exp}}(\Delta\theta)$ can be described as a mixture of a von Mises distribution $f_{\text{VM}}(\Delta\theta|s)$ with dispersion parameter s and a uniform distribution $f_{\text{U}}(\Delta\theta)$: $f_{\text{exp}}(\Delta\theta) = (1 - u)f_{\text{VM}}(\Delta\theta|s) + uf_{\text{U}}(\Delta\theta)$. The idea is that when there is a memory of the item the responses over trials will be distributed around the true position following a von Mises distribution with some dispersion s. However, when the location of the item has been forgotten, the participant will guess and hence the responses will follow a uniform distribution. One can then estimate the parameters of the mixture model u, s and obtain estimates for the number of items retained in memory $(1 - u)$ and precision $(1/s)$ of the responses. The quantity $(1 - u)$ relates to WM capacity. It has been proposed that a measure of capacity can be obtained by multiplying $(1 - u)$ with the load (Zhang & Luck, 2008). In our work, we are using $(1 - u)$ directly because our aim is to characterize differences in $(1 - u)$ for fixed load (see below). The parameters u and s were estimated using the functions described in Zhang and Luck (2008) and Bays et al. (2009) and available at www.sobell.ion.ucl.ac.uk/pbays/resources.htm. Trials with an incorrect number of responses and trials where a click was given too far from the background circle (arbitrary threshold of 7.65 visual degrees) were removed. In some cases, it could be that because of imprecision in the remembered locations a participant mistakes the identity of the first item and therefore reports the remembered location of another item. Similar types of mistakes that result in the participant reporting remembered information for a nontargeted item have been modeled using a mixture model with three components (Bays et al., 2009). The third component models the response to a nontarget item. In the current experiment, we did not have enough trials to estimate one more parameter from the model. Instead, we checked the data for evidence of the described type of mistakes. In particular, we checked if for a given trial the positions indicated by the responses were better matched with the stimulus assuming that the participant had started from the last or second item. This was the case on only 1% of the trials. When this happened, we relabeled the clicks according to what the participant probably thought was the first item.

Estimates of capacity and precision were obtained using data from the first response (mouse-click) on each trial to ensure that responses were as independent as possible within a participant. However, the use of just one click per trial led to less than 50 responses per participant for conditions with false cues. This number of responses is too low for a reliable estimation of the parameters using the mixture model. To overcome this problem, we pooled data across participants and estimated an average s and u for all participants for each of the four trial types: Cue 3/Load 3, Cue 5/Load 3, Cue 3/Load 5, Cue 5/Load 5. Then we calculated the effect of cue on s and u for each load. Note that pooling data across participants might inflate the estimates of s because of possible inconsistent recalling bias across participants, which would mean that responses from different participants reflect von Mises distributions with unequal means. However, this eventual inflation should be similar for all trial types and hence should not

affect differences of s and u between conditions. To test the hypothesis that there was no trade-off effect of cue on precision and capacity for fixed load, we used a bootstrap method. First, we drew a random sample from participants, with repetitions allowed and with the same size as the original data set. Then, for each participant in the sample, we randomly drew two samples with the same number of trials as in the original samples under the null hypothesis that cue has no influence (that is, from a pool of trials with the same load value, but randomized across cue). For each of the two samples, s and u were estimated across participants. We then calculated the differences across samples of these estimates. This procedure was repeated 10,000 times. Finally, we determined the number of times M that the trade-off effect of cue was larger in the bootstrap draws than in the original sample, that is, the number of times the effects were larger for both s and u. Results were considered significant at $p < .05$ if $M/10{,}000 < 0.05$.

To relate brain activity to behavioral performance, we sought to estimate participant-specific capacity and precision. To have enough data for estimation, we used all responses (mouse-clicks) from a trial. However, using all the responses is problematic, leading to estimates less reliable than those calculated from first-click data (see Discussion). The problems are that the responses are not independent within each trial (in fact they will probably be guided by the location of the previous mouse-clicks) and that the delay to recall (likely affecting performance) is increasing with the number of clicks. To partly address one of these issues, we used a modification of the probabilistic mixture model where we assume that guesses are made uniformly between the last clicked position on that trial and the red indication mark.

fMRI Experiment

Participants

A subset of 22 right-handed participants (24 ± 3 years) participated in the fMRI experiment.

Stimuli and Procedure

The experiment was adapted from the behavioral experiment, with the following differences: The fixation cross at the beginning of the trial was made variable in length (500, 1500, or 2500 msec) to introduce jitter. The cue was always 1000 msec and was followed by another fixation cross for a variable time (1000, 2000, or 3000 msec). The stimulus was presented for 1000 msec and was followed by a delay that lasted for 2000, 3000, or 4000 msec. The maximum RT was set to 4000 msec. The intertrial interval was at least 2000 msec and was calculated so that the next trial always started together with the scanner pulse. Participants indicated the remembered positions with a scanner compatible trackball. Because this took longer than pointing with a classical mouse, participants only indicated the remembered position of the first item. Note that participants still had to remember all items because the red indication mark determining the first item was placed at random in the circle.

Participants completed three runs in which true cue trials (Cue 3/Load 3 and Cue 5/Load 5) and false cue trials (Cue 3/Load 5 and Cue 5/Load 3) were shown, intermixed with null events. There were 42 trials in each of these runs. Intermixed with these runs, participants completed three more runs in which we showed a different type of trials (control trials). These runs were not used in the analysis and are therefore not further described. Order of runs was counterbalanced between participants.

Stimuli were presented on a back-projection screen at the head of the scanner bore. Participants viewed the screen through a mirror mounted on the head coil. Images were collected with a 3T General Electrics Discovery scanner (MR750, General Electrics, USA). First, a high-resolution anatomical image was acquired using a T1-weighted 3-D anatomical sequence. Whole-brain functional images were collected using a T2*-weighted EPI sequence (repetition time = 2500 msec, echo time = 30 msec, frequency field of view = 288 mm, phase field of view = 1.0, frequency = 96, flip angle = 90°, slice thickness = 3.0 mm, voxel size 2.25 × 2.25 × 3, 40 axial slices). Runs started with six dummy scans, and 260 images were taken in each run. The experiment lasted around 1 hr.

Data Analysis

Runs were excluded from analyses when the realignment procedure showed excessive motion or when more than one third of the trials were not responded to. Only one run from one participant met these exclusion criteria.

Data analysis was performed using the SPM8 toolbox (www.fil.ion.ucl.ac.uk/spm/). Motion parameters (three translations, three rotations) were used to realign the functional volumes to the first image of each run, and all functional images were coregistered with the anatomical scan of the participant. Functional images were high-pass (140 sec) filtered, and first level analyses were done on these images. Next, we used the Dartel toolbox to create a template from all anatomical images of all participants. Finally, the contrast images and mask images created in the first-level analyses were normalized to MNI space using the Dartel template. The contrast images were further smoothed with a Gaussian kernel of 8 mm FWHM.

The signal was modeled using a general linear model analysis with seven predictors. Two predictors modeled the cues: Cue 3 and Cue 5. The duration of these events covered both the actual cue period (1000 msec) and the variable fixation time between cue and stimulus. Next, there were four predictors modeling the stimulus periods of the four different trial types. The duration of these events covered both the actual stimulus period (1000 msec) and the variable delay between stimulus and response phase. Finally, the response predictor included the response

phase of all trials. The predictors were built by convolving the actual presentation times of the events in the experiment with the standard hemodynamic response function. Time and dispersion derivatives as well as the motion parameters were added in the design matrix. Finally, three run-specific predictors were added to model differences between mean activation of the runs.

In the computational model, the boost signal is active during the stimulus and delay period in trials where boosting occurs. We expected boosting to occur in the trials where a cue for a high load was given. Therefore, we contrasted the stimulus and delay period from Cue 5/Load 3 trials with Cue 3/Load 3 trials and from Cue 5/Load 5 trials with Cue 3/Load 5 trials. These contrast images (two for each participant) were then used in a random effects conjunction analysis, with global null hypothesis. This analysis aimed to identify significant consistent effects among the two contrasts. The Marsbar toolbox (Brett, Anton, Valabregue, & Poline, 2002) was used to extract contrast values from ROIs for subsequent analysis. To investigate the possible functional role of the clusters found with this analysis, we correlated across participants the level of activation with the behavioral effects on precision and capacity. Note that the behavioral effects were estimated from data collected outside the scanner. The behavioral data acquired during scanning did not have enough responses to allow estimation of the parameters of the mixture model.

RESULTS

Model

In a computational model of vsWM, a nonspecific excitatory signal (boosting signal) to a storage network results in an increase of the capacity of the network (Edin et al., 2009). We used the same model (Figure 1) to study how the precision of memories is affected by such capacity boosting. Figure 1 shows an example of the activity of excitatory neurons of the storage network when the same stimulus is presented in the case of low input from the boost area (Figure 1B) and high input from the boost area. In the case of low input from the boost area, one of the memory traces is lost before the end of the delay period. This memory loss does not occur for the case with high boost input, consistent with the capacity boosting effect reported before (Edin et al., 2009). However, this ability of retaining more items in memory may come at the cost of a lack of precision of the remembered locations, as suggested. This is shown by larger deviations of the memory traces from the horizontal lines indicating the locations of the items to memorize in the high compared with the low boost input case. To quantify this apparent trade-off effect, we ran multiple simulations, with items presented at random locations on a circle, using low or high input from the boost network and loads of three or four items (300 simulations per combination of boost and load). For each load, we estimated the capacity by calculating the percentage of memory traces, which persisted after the delay period. The precision was estimated by the inverse of the standard deviation of the population vectors of the memory traces in the collection of trials (see Methods). We found (Figure 3) that a high top–down input for randomly located items resulted in an increase in capacity (binomial proportion test, $p < .0001$ for Loads 3 and 4), as described previously for evenly located items (Edin et al., 2009). Furthermore, as shown in Figure 3, we found that a high top–down signal resulted in a decrease in precision of the memories (F test for equality of variances, $p < .0001$ for Load 3 and $p = .003$ for Load 4). This trade-off effect of increasing capacity at the cost of precision was observed for both loads. In general, performance deteriorated with an increase in memory load, but the changes induced by the boost remained significant for the two loads tested (Figure 3). From these modeling results, we predicted that induction of a boost signal while keeping the load constant should result in a trade-off between capacity (increased) and precision (decreased) of vsWM.

Behavioral Experiment

We tested the prediction of the computational model using a behavioral experiment (Figure 2). We hypothesized

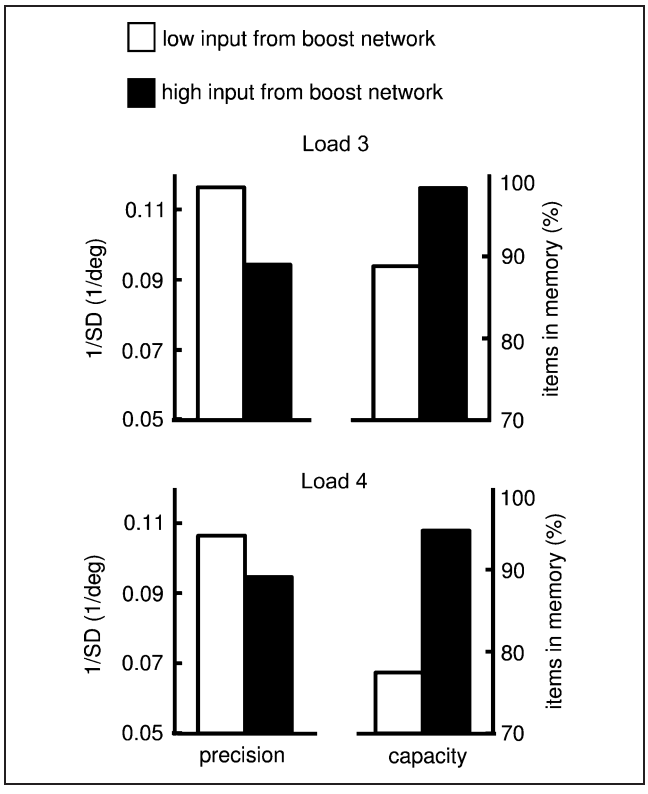

Figure 3. Model results. For both Load 3 (top) and Load 4 (bottom), the effect of high top–down as compared with low top–down input is a decrease in precision (inverse of standard deviation) and increase in capacity (percentage of items in memory). Results are averaged over 300 simulations for each load and level of top–down input.

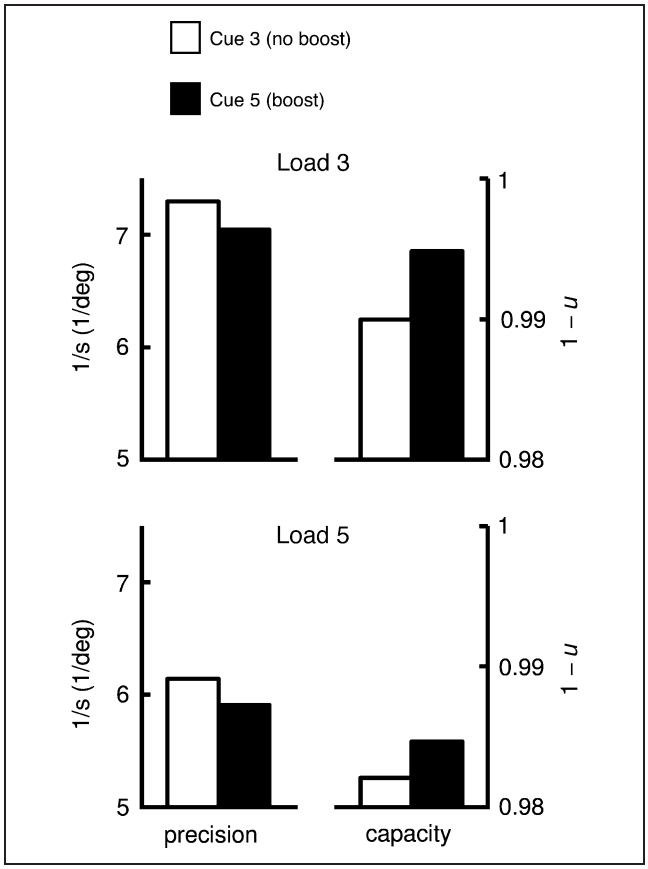

Figure 4. Behavioral results. Mean values (over participants) of precision (calculated as 1/*s*; see Methods) and capacity (calculated as $1 - u$). For both Load 3 (top) and Load 5 (bottom), the effect of Cue 5 as compared with Cue 3 (boost vs. nonboost) is a decrease in precision and increase in capacity.

that engagement of a boost signal could be controlled on a trial-by-trial basis and triggered by a cue indicating the number of items (three or five) to be expected on the current trial. Because three is below the capacity limit for many individuals, we expected that Cue 3 would not elicit a boost signal. We expected Cue 5, on the other hand, to elicit a boost signal. Crucially, the cue was false in around 20% of trials. This allowed us to detect the effect of cue on trials with identical memory load. Assuming that the cue has an effect on boosting, we could test our hypothesis that boosting, accessed comparing Cue 5 and Cue 3 trials with identical load, increases the capacity but decreases the precision of WM.

We used the probabilistic mixture model described in Zhang and Luck (2008) and Bays et al. (2009) to estimate capacity and precision for the different trial types using the first response of each trial and pooling data across participants for all conditions in Figure 4 (see Methods). We discarded 1% of the trials (maximum percentage discarded for a single participant was 5%) according to the criteria described in Methods. The bootstrap analysis showed that, for both loads, vsWM performance in Cue 5 (also labeled as boost condition in Figure 4) relative to Cue 3 (also labeled as no boost condition) trials revealed a concomitant increase in capacity and decrease in precision ($p = .019$ for Load 3 and $p = .047$ for Load 5). A decrease in performance with memory load was observed both for precision and capacity measures, but the trade-off effect was independent of load (Figure 4), as in the model results. Thus, the experimental evidence was in agreement with the hypothesized trade-off effect.

We sought to determine individual estimates of the effect of cue on capacity and precision. To have enough false cue trials for a single participant, we used data from all responses (mouse clicks) in each trial. This approach has several problems (see Methods and Discussion) leading to estimates less reliable than those calculated from the pooled data. Despite this, we determined individual estimates with the goal of using them for analysis with the fMRI data (see below). However, for completeness, we also analyzed the individual estimates per se. We did this using a 2×2 ANOVA for s and u, with Cue and Load as factors. For u (accessing capacity), there was an effect of Load ($p = .00016$), but there was no effect of Cue and no significant interaction between Load and Cue. For s (accessing precision), both the effects of Load and Cue were significant ($p < .0001$ and $p = .013$, respectively), but there was no significant interaction. One of the possible reasons for not finding a significant effect of Cue on s might be that not all participants engage a boost signal differentially for Cue 5 versus Cue 3 trials (see Discussion). In this situation, the trade-off effect would still predict that individual changes in capacity and precision should be correlated across participants for each load. We calculated the individual effects on capacity and precision of having Cue 5 as compared with Cue 3 for each load, and we found that the two effects were significantly correlated across participants for both loads, that is, individuals exhibiting a large increase in capacity for Cue 5 relative to Cue 3 trials also tended to show a large decrease in precision (Pearson correlation .52 for Load 3 and .51 for Load 5, $p = .001$ for both loads).

fMRI Results

We used fMRI data acquired using the same behavioral paradigm to investigate possible brain activity correlates of the trade-off effect hypothesize to result from the engagement of a boost signal. The behavioral data acquired during scanning indicated that the participants were alert and engaged in the task. Participants responded on average to 97% of the trials (minimum number of responses per participant was 90%), and these responses deviated from the target less than 22.5° on 91% of the trials (minimum per participant was 74%). We then looked for brain regions showing an effect of Cue 5 versus Cue 3 during the stimulus presentation and delay period, according to the computational model. To this end, we investigate whether there were significant consistent effects for the contrasts (Cue 5/Load 3 vs. Cue 3/Load 3) and (Cue 5/Load 5 vs.

Table 1. Clusters Found for the Contrasts (Cue 5/Load 3 vs. Cue 3/Load 3) and (Cue 5/Load 5 vs. Cue 3/Load 5) Using a Random Effects Conjunction Analysis with Global Null Hypothesis

MNI Coordinates			*Region*	*Voxels*	*p Level of Local Maximum*	*F*
−18	−62	45	Left superior parietal	299	2.22e−06	5.45
30	−80	34	Right superior parietal	480	3.20e−05	4.39
49	3	18	Right inferior frontal	287	2.39e−05	4.50

Cue 3/Load 5) using a random effects conjunction analysis with global null hypothesis (Nichols, Brett, Andersson, Wager, & Poline, 2004; Price & Friston, 1997). Because the analysis was used to define ROIs for subsequent analyses, we used a lenient uncorrected threshold of $p = .001$, but we restricted the analyses to clusters with more than 250 voxels. Three clusters survived these criteria (see Table 1 and Figure 5). Two clusters were located in the left and right superior parietal cortices. The other cluster was located in the frontal cortex (right inferior frontal gyrus), suggesting that differences in cue might cause differential changes in a boost signal.

To investigate the possible role of these clusters on the behavioral effect found, we correlated the level of activation in each participant with the magnitude of the behavioral effects estimated from the behavioral experiment outside the scanner. We extracted the contrast values for both (Cue 5/Load 3 vs. Cue 3/Load 3) and (Cue 5/Load 5 vs. Cue 3/Load 5) from these clusters for each participant. We then tested for significant correlations between the contrast values for each load and the individual precision and capacity effects from the behavioral experiment. The precision and capacity effects were calculated from the difference of precision ($1/s$) and capacity ($1 - u$) estimates between Cue 5 and Cue 3 trials.

We found that the contrast values for Load 5 in the right inferior frontal gyrus cluster correlated positively with the differences in capacity ($p = .035$; Figure 5AI,

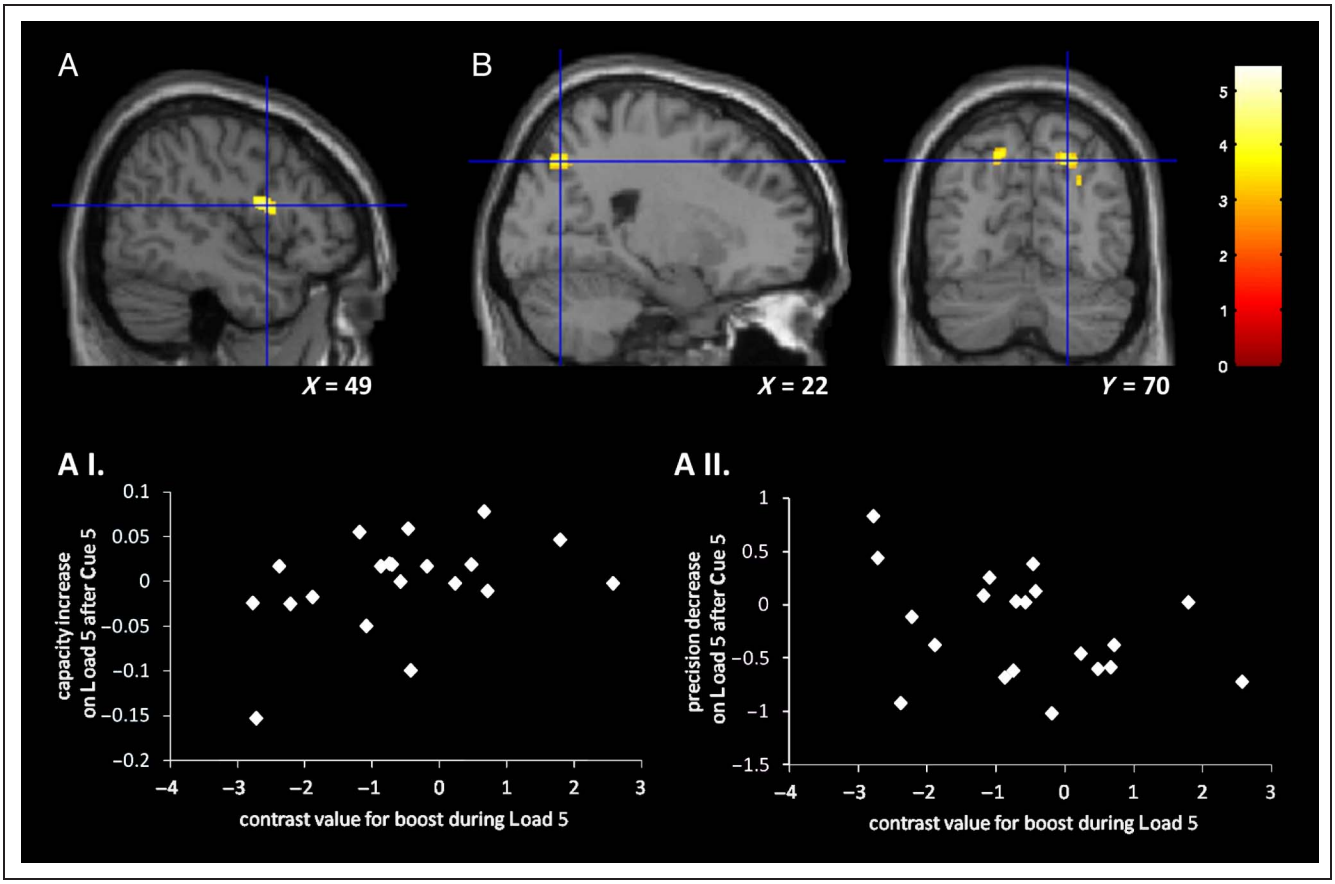

Figure 5. Whole-brain results for the contrast (Cue 5/Load 3 vs. Cue 3/Load 3) and (Cue 5/Load 5 vs. Cue 3/Load 5) using a random effects conjunction analysis with global null hypothesis. The analysis was restricted to clusters with a cluster extent threshold of 250 voxels. (A) Cluster in the right inferior frontal sulcus. (B) Clusters in left and right superior parietal lobe. (AI and AII) Correlations between the contrast values extracted from the right inferior frontal sulcus shown in A and the individual capacity (AI) and precision (AII) effects estimated from the behavioral experiment. The effects are the differences for u and s between the estimates for Cue 5 and Cue 3.

$p = .085$ removing two mild outliers) and negatively with the differences in precision ($p = .048$; Figure 5AII). That is, a larger contrast value was associated with a larger increase in capacity and with a larger decrease in precision across participants. So, for Load 5, participants showing larger effects on the trade-off between capacity and precision also showed larger activation in a region in the right inferior frontal cortex for trials assumed to elicit boosting. No other correlations were found.

DISCUSSION

In this work, we used a computational model consisting of two modules to make the prediction of a trade-off between capacity and precision of vsWM. In the model, the trade-off results from a boost signal to the storage module of WM. This effect would not have been easily predicted departing only from previous behavior experiments. We further hypothesized that the boost signal could be adjusted on a trial-by-trial way by cueing the upcoming memory load. We showed in a behavioral experiment that cueing indeed induces a trade-off between precision and capacity of vsWM at fixed load. Presumably the cue is eliciting the boost signal. Using fMRI, we found that the cue effect at fixed load is associated with activation of a region in the inferior frontal gyrus, and activity in this region across participants correlates with the amount of trade-off showed by each individual. This region is then a candidate region for the origin of the hypothesized boost signal. Taken together, the results support a model of vsWM consisting of two modules that contribute independently to limitations in WM performance. One source of limitations originates in the storage circuit and the other on the trial-by-trial recruitment of a boosting signal.

Our present work is closely related to recent work on the ability to adjust the capacity and the precision of WM to task requirements (Machizawa et al., 2012; Zhang & Luck, 2011). Zhang and Luck (2011) tested whether participants could retain in WM information about a larger number of items at the cost of lower precision of the information retained, when given incentives to do so. In most cases, they found no evidence of such a trade-off adjustment. This is at odds with our findings, although there are several plausible causes of such differences. One is that Zhang and Luck (2011) studied WM of colors while we used spatial locations. We used vsWM because we departed from a biologically inspired model, which relies on a neuronal topographic organization of spatial representations. To our knowledge, an equivalent topographical organization along a color dimension for color-sensitive neurons has not been described experimentally, and hence, the biological plausibility of our model in that case is less certain. This implies that our findings will not necessarily extend to other modalities of vWM that could underlie the differences found. Note, however, that a functional topographic organization of color representation has been proposed, in particular in the context of models of WM (Wei, Wang, & Wang, 2012; Johnson, Spencer, Luck, & Schöner, 2009; Johnson, Spencer, & Schöner, 2009). The success of these models in explaining experimental data suggests that our model might also be applicable in the case of WM of colors and eventually of other parametric dimensions. More work will be needed to determine whether our predictions for vsWM extend to WM of other dimensions and, in general, to characterize possible modality differences in WM limitations.

Another possible origin of the differences in results with the work of Zhang and Luck (2011) is the way in which the two studies induce the trade-off. In our work, we mix trials of different loads and we use cues to elicit capacity boosting selectively, whereas Zhang and Luck (2011) used a fixed load but different precision requirements in the response. Our paradigm might be recruiting a boost signal to increase capacity, which is not recruited in the previously used paradigm (Zhang & Luck, 2011). Finally, interindividual variability might explain the differences found. Zhang and Luck (2011) argue that a trade-off can be achieved only when the number of items is below the individual capacity limit. Our results could be driven by participants having high vsWM capacity (higher or equal to 3), so that they could trade-off precision and capacity at least in some of our experimental conditions. This would be consistent with the small effect sizes reported here (see below).

Another recent work showed that precision of vsWM can be adjusted on a trial-by-trial basis for low memory load (Machizawa et al., 2012). However, this study did not test how manipulations of precision affected vsWM capacity or whether capacity could be similarly adjusted on a trial-by-trial basis resulting in changes of vsWM precision. One other study has suggested that a boosting signal can increase vsWM capacity (Edin et al., 2009). In that work, the eventual adjustment of vsWM capacity resulted from presentation of stimuli with different loads and was in this sense reactive. In the current study, we show that vsWM capacity can be adjusted on a trial-by-trial basis, matching the expected demand of the trial when the actual demand (the load) is the same. More importantly, we demonstrated for the first time a trade-off effect for vsWM performance, where capacity can be increased at the cost of precision. The specific relation between capacity and precision as they are individually manipulated may constrain strongly the nature of plausible models of vsWM.

The differences found in capacity and precision induced by different cues are small. One likely reason for this is the presence of considerable interindividual variability, which results in a small effect size when data is pooled across individuals. Some individuals were likely not trading capacity for precision in this task, presumably by not engaging a boost signal differentially. There are several reasons for this. For some individuals, Loads 3 and 5 might be almost equally difficult (or easy), and therefore, the paradigm used should not imply a change in boost input. Some individuals were maybe not using

the information provided by the cue. Different participants might also use different strategies. Because the items are presented simultaneously, the participants might memorize a pattern of dots instead of the location of each item separately. In this case, the increase in number of dots might not correspond to an increase in load requiring a boost of capacity. We used simultaneous instead of sequential item presentation to reduce the experimental time. Note that we do not expect the possible encoding of groups of positions together to bias the results for a particular load and cue, and hence, this should not affect the conclusions. Instead, this could introduce variability, making the effects smaller. In summary, our experimental results demonstrate that a trade-off between capacity and precision can occur, which poses constraints on models for vsWM. However, our experimental results do not allow us to fully characterize the conditions and the extent of such a trade-off. For example, it would be of interest to study if the trade-off occurs only for some participants or only below a participant-specific number of items (as the model would suggest) or if no such limitation exists.

Using fMRI, we found three clusters showing significant activation for high versus low cue, consistent for both memory loads. The clusters were located bilaterally in the superior parietal cortex and right inferior frontal gyrus. These regions are differentially involved in trials expected to have high versus low demands on vsWM and hence presumably recruiting a boosting signal. An alternative explanation is that these regions are involved in the detection of conflicting information, in this case between the cue and the actual load. However, we found that the strength of the activation in the frontal region correlated across individuals with the amount of decrease in precision and increase in capacity for Load 5. These correlations would not be expected if the region was involved in processing the presence of conflicting information. Hence, our results suggest that this region might be the origin of a boost signal. Moreover, the correlations suggest that individuals might differ on the engagement of such a boost signal. There are several possible explanations for individual differences, as described above. A further possible difference in the context of fMRI is that some individuals might sustain the boost signal during a trial, independent of the stimuli presented, whereas others might adjust it upon stimulus presentation. The latest group of participants would maybe not show a behavior effect or maybe the effect would still be present but would not be associated with frontal activation detectable in our fMRI analysis. Finally, the behavioral results obtained for single participants are estimated using all the clicks in each trial on a behavioral experiment performed outside the scanner. This estimation is made so that there are enough trials per participant at the cost of using data acquired on a different occasion and using less well-controlled data. In fact, using all the responses (mouse clicks) in one trial is problematic for a number of reasons. One problem is that as the number of clicks increases so does the delay to recall. This means that the last two items clicked for Load 5 trials will always correspond to larger delays than for Load 3 items. Another problem is that after each click it is likely that the participant will have a memory of the clicked position, making the responses dependent within each trial. In fact, after the first click, the responses will be guided by previous clicks in the same trial, and hence in the case of memory loss, guesses will be made on just a fraction of the circle and not on 360°. To partly address this last caveat, we used a modified model to estimate individual capacity and precision (see Methods). Despite this, the individual estimates of capacity and precision are potentially contaminated by noise, making it harder to interpret interindividual differences. These factors might explain why we found a correlation between behavior and activation of the frontal area for Load 5 trials but not Load 3 trials.

The prediction of a trade-off between capacity and precision in vsWM was derived from a biological constrained model (Edin et al., 2009; Compte et al., 2000). This model was originally developed to account for single neuron activity measured in monkeys performing a vsWM task (Funahashi, Bruce, & Goldman-Rakic, 1989). The model was shown to reproduce the observed sustained elevated activity during a delay period after stimulus offset and selective for the stimulus visual location. The model was later shown to also account for human behavior (Wei et al., 2012; Macoveanu, Klingberg, & Tegnér, 2007) and fMRI data (Edin et al., 2009; Macoveanu, Klingberg, & Tegnér, 2006). Edin et al. (2009) used a modified version of the original computational model consisting of two coupled networks. This model was used to study WM capacity limitations and the increase of capacity by a boost input to the storage network for the case of stimulus consisting of locations evenly distributed across a circle. The present work departed from the observation that in this computational model the boost signal increasing storage capacity comes with a cost to the precision with which the memories of randomly positioned locations are stored. The decrease in precision is a result of at least two factors. One is the increased noise from increased nonspecific background input to all neurons of the storage network. This decrease in precision is also observed in the model for Load 1. The other factor is that the boost input increases the number of memory traces actually maintained by the network (capacity boosting). When more memory traces are active, the probability of interference between traces increases, resulting in decreased precision (Wei et al., 2012). In this work we used the same parameters as in Edin et al. (2009) without further tuning. The reason for this was that we searched for a qualitative robust prediction to test experimentally and we were not concerned with fitting quantitatively a given result (e.g., the average capacity or precision for each load). For simplicity, we have also kept the simulation protocol where the boost signal and the stimulus appear at the same time. Instead, we could have modeled the onset of boost input occurring before the

stimulus onset, mimicking recruitment of boost input at the time of cue presentation. Qualitatively, the results would remain the same as long as the extra boost input would not destabilize the spontaneous state of the storage network. Experimental results from visual WM tasks have been used to propose and support several influential theories accounting for WM limitations, namely the "slots model" (Anderson et al., 2011; Zhang & Luck, 2008; Luck & Vogel, 1997) and the "resources model" (van den Berg et al., 2012; Huang, 2010; Bays et al., 2009; Bays & Husain, 2008; Wilken & Ma, 2004). The computational model used was not developed to match any of these competing theories of WM limitations. Instead, the computational model was developed to match electrophysiological data recorded in monkeys. In the current and previous work (Wei et al., 2012), this more biological constrained model shows features that are consistent with results that have been used to support both the slots and the resources models of WM. Concretely, the computational model can retain information about a limited number of items. Such a limitation is in line with the slots account of WM. However, the model proposes a plausible mechanism for a boost in WM capacity, which comes at a cost of a decrease in precision. This trade-off is in line with predictions from a resources account of WM. Hence, the computational model can be seen as a possible biological implementation of a hybrid model as proposed at a more abstract level previously (Buschman et al., 2011; Xu & Chun, 2006; Alvarez & Cavanagh, 2004). The approach of using more biological-based computational models is promising. First, we can contribute by proposing concrete mechanisms underlying the observed behaviors. Second, we can help to reconcile more abstract models. Finally, we can formulate testable predictions that would be hard to come across without the model. Here, we have predicted and confirmed a trade-off effect between capacity and precision in vsWM performance, mediated by the activation of an area in frontal cortex.

APPENDIX

Mathematical Description of the Computational Model

The computational model was described by Edin et al. (2009) and was used here without further tuning of parameters. The following paragraphs outline the equations and parameters used in this study. The storage network consisted of 1024 excitatory (E cells) and 256 inhibitory (I cells) interconnected leaky integrate-and-fire neurons (Tuckwell, 1988). A more detailed description of this network can be found in Compte et al. (2000). Integrate-and-fire neurons are described in terms of their subthreshold membrane potential (V_m) according to the equation $C_m\, dV_m/dt = -g_L\,(V_m - E_L) - I_{syn} - I_{ext}$, where C_m is the membrane capacitance, g_L is the leak conductance, E_L is the leak reversal potential, and I_{syn} and I_{ext} are inputs from other neurons inside or outside the network, respectively. When the membrane potential reaches a given threshold V_{th}, it fires an action potential (not explicitly modeled) and is set to a reset value V_{res} during a refractory period τ_{ref}. For E cells, we used $C_m = 0.5$ nF, $g_L = 25$ nS, $E_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 2$ msec; for I cells, we used $C_m = 0.2$ nF, $g_L = 20$ nS, $E_L = -70$ mV, $V_{th} = -50$ mV, $V_{res} = -60$ mV, and $\tau_{ref} = 1$ msec.

The network had a ring structure so that E and I cells were spatially distributed on a ring, where nearby neurons encoded similar spatial locations θ (Figure 1A). Connections were tuned according to this organization so that the connection strength $g_{syn,ij}$ between cells i and j depended on the difference in preferred angle between the cells as $g_{syn,ij} = W(\theta_i - \theta_j)\, G_{syn}$, where $W(\theta_i - \theta_j) = J^- + (J^+ + J^-)\exp(-(\theta_i - \theta_j)^2/\sigma^2)$, with J^- set to satisfy a normalization condition $S_j W(\theta_i - \theta_j) = 1$. We used $\sigma_{E\,to\,E} = 9.4°$ and $\sigma_{E\,to\,I} = \sigma_{I\,to\,E} = 32.4$, $J^+_{E\,to\,E} = 5.7$, $J^+_{E\,to\,I} = J^+_{I\,to\,E} = 1.4$, $J^+_{I\,to\,I} = 1.5$. So, connectivity between E and I cells was wider and flatter than between E cells (Figure 1A). The connectivity between I cells was not spatially tuned. The strengths of the connections were $G_{E\,to\,E} = 0.7$ nS, $G_{E\,to\,I} = 0.49$ nS, $G_{I\,to\,E} = 0.935$ nS, $G_{I\,to\,I} = 0.7413$ nS.

All neurons received uncorrelated random background excitatory input modeled as spike trains described by a Poisson process with a rate of 1800 Hz, through a conductance $g_{ext\,to\,E} = 6.5$ nS, $g_{ext\,to\,I} = 5.8$ nS. Upon presentation of a visual stimulus at θ_{stim}, E cells received an input modeled as current injection: $I_{stim}(\theta,\theta_{stim}) = a \exp\{m\,[\cos(2\,\pi/360(\theta - \theta_{stim})) - 1]\}$, where $a = 0.025$ nA and $m = 39$.

Postsynaptic currents were described by the equation $I_{syn} = g_{syn}\, s\, (V_m - V_{syn})$, where g_{syn} is a synaptic conductance, s is a synaptic gating variable, and V_{syn} is the synaptic reversal potential ($V_{syn} = 0$ for excitatory synapses, $V_{syn} = -70$ mV for inhibitory synapses). Recurrent excitatory connections were modeled to follow the dynamics of NMDAR-mediated transmission, external excitatory inputs to follow AMPAR-mediated transmission, and inhibitory inputs to follow $GABA_AR$ transmission. AMPAR and $GABA_AR$ gating variables were modeled as an instantaneous jump of magnitude 1 when a presynaptic action potential occurs, followed by an exponential decay with time constant 2 msec for AMPA and 10 msec for $GABA_A$. The NMDAR conductance is voltage dependent, and this was modeled by multiplying g_{syn} by $1/(1 + [Mg^{2+}]\exp(-0.062\,V_m)/3.57)$ with $[Mg^{2+}] = 1.0$ mM. The dynamics of NMDA gatings were modeled by $ds/dt = -s/\tau_s + \alpha_s\, x(1 - s)$ and $dx/dt = -x/\tau_x + \Sigma_i\, \delta(t - t_i)$, where x is a synaptic variable representing neurotransmitter concentration in the synapse, t_i are the times of presynaptic action potentials, $\tau_s = 100$ msec is the decay time, $\tau_x = 2$ msec the rise time, and $\alpha_s = 0.5$ kHz controls saturation of NMDAR channels.

The boost network only differs from the storage network in that its connectivity was not tuned (all $J^+ = 1$). The values

used for the conductances were $g_{E\ to\ E} = 0.968$ nS, $g_{E\ to\ I} = 0.723$ n, $g_{I\ to\ E} = 3.66$ nS, $g_{I\ to\ I} = 2.832$ nS, $g_{ext\ to\ E} = 3.0$ nS, $g_{ext\ to\ I} = 2.3803$ nS. The boost input was modeled to impact neurons in the storage network through AMPARs, with conductance of 9 nS.

The integration of the model equations was done using a second-order Runge-Kutta algorithm in a custom code implemented in C^{++}.

Acknowledgments

We thank M. Spencer-Smith and F. Edin for valuable comments on the manuscript. This work was supported by the Karolinska Institutet Strategic Neuroscience Program, the Swedish Research Council (VR 524-2011-1068 to C. R.), the Ministry of Economy and Competitiveness of Spain, the European Regional Development Fund (BFU2009-09537 to A. C.), and the Generalitat de Catalunya (Beatriu de Pinós program 2007BP-B100135 to R. A.).

Reprint requests should be sent to Rita Almeida, Department of Neuroscience, Karolinska Institute, Retzius Väg 8, 171 77 Stockholm, Sweden, or via e-mail: Rita.Almeida@ki.se.

REFERENCES

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15,* 106–111.

Anderson, D. E., Vogel, E. K., & Awh, E. (2011). Precision in visual working memory reaches a stable plateau when individual item limits are exceeded. *Journal of Neuroscience, 31,* 1128–1138.

Baddeley, A. (1986). *Working memory*. New York: Oxford University Press.

Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision, 9,* 1–11.

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science, 321,* 851–854.

Brett, M., Anton, J. L., Valabregue, R., & Poline, J. B. (2002). Region of interest analysis using an SPM toolbox. *Neuroimage, 16,* 497.

Buschman, T. J., Siegel, M., Roy, J. E., & Miller, E. K. (2011). Neural substrates of cognitive capacity limitations. *Proceedings of the National Academy of Sciences, U.S.A., 108,* 11252–11255.

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X. J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex, 10,* 910–923.

Conway, A. R., Kane, M. J., & Engle, R. W. (2003). Working memory capacity and its relation to general intelligence. *Trends in Cognitive Sciences, 7,* 547–552.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–185.

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

Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegnér, J., & Compte, A. (2009). Mechanism for top–down control of working memory capacity. *Proceedings of the National Academy of Sciences, U.S.A., 106,* 6802–6807.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology, 61,* 331–349.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science, 233,* 1416–1419.

Huang, L. (2010). Visual working memory is better characterized as a distributed resource rather than discrete slots. *Journal of Vision, 10,* 8.

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science, 20,* 568–577.

Johnson, J. S., Spencer, J. P., & Schöner, G. (2009). A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain Research, 1299,* 17–32.

Kyllonen, P. C., & Christal, R. E. (1990). Reasoning ability is (little more than) working memory capacity. *Intelligence, 14,* 389–433.

Lee, D. D., Reis, B. Y., Seung, H. S., & Tank, D. W. (1997). Nonlinear network models of the oculomotor integrator. In J. Bower (Ed.), *Computational neuroscience* (pp. 317–377). New York: Plenum Press.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281.

Machizawa, M. G., Goh, C. C., & Driver, J. (2012). Human visual short-term memory precision can be varied at will when the number of retained items is low. *Psychological Science, 23,* 554–559.

Macoveanu, J., Klingberg, T., & Tegnér, J. (2006). A biophysical model of multiple-item working memory: A computational and neuroimaging study. *Neuroscience, 141,* 1611–1618.

Macoveanu, J., Klingberg, T., & Tegnér, J. (2007). Neuronal firing rates account for distractor effects on mnemonic accuracy in a visuospatial working memory task. *Biological Cybernetics, 96,* 407–419.

Nichols, T., Brett, M., Andersson, J., Wager, T., & Poline, J. B. (2004). Valid conjunction inference with the minimum statistic. *Neuroimage, 25,* 653–660.

Price, C. J., & Friston, K. J. (1997). Cognitive conjunction: A new approach to brain activation experiments. *Neuroimage, 5,* 261–270.

Tuckwell, H. C. (1988). *Introduction to theoretical neurobiology*. Cambridge: Cambridge University Press.

van den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences, U.S.A., 109,* 8780–8785.

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. *Journal of Neuroscience, 32,* 11228–11240.

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision, 4,* 1120–1135.

Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature, 440,* 91–95.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453,* 233–235.

Zhang, W., & Luck, S. J. (2011). The number and quality of representations in working memory. *Psychological Science, 22,* 1434–1441.