Is short-term storage correlated with fluid intelligence? Strategy use
explains the apparent relationship between 'number of remembered items'
and fluid intelligence

Zeljana Babic¹, Mark W. Schurgin¹, Timothy F. Brady¹

¹ Department of Psychology
University of California, San Diego

Word count: 3947

* Correspondence to:

Timothy Brady

9500 Gilman Drive #0109

La Jolla, CA 92093

Email: timbrady@ucsd.edu

Abstract

Working memory is a core cognitive system that actively maintains information in an accessible state to support a variety of everyday tasks. Crucially, working memory performance has frequently been shown to strongly correlate with fluid intelligence. Traditionally when these correlations have been observed, the working memory tasks involved required a high degree of manipulation and executive function, as opposed to solely utilizing short-term storage capacity. However, recent work has claimed that simple storage capacity is also correlated with fluid intelligence, and that this is driven by a particularly special and dissociable component of capacity, the ‘number of items represented’ (rather than the precision of those representations). These results have been used to argue that investigating the underlying mechanisms of capacity limitations may be critical to understanding aspects of fluid intelligence. Here we demonstrate that such correlations do not arise solely or primarily from simple storage capacity (nor a single dissociable component of capacity), but are driven by the availability of strategic encoding of different kinds of visual representations. Specifically, a working memory task that decreased the utility of storing and making use of spatial ensemble information, while holding constant the number of items to be remembered and the exact changes participants needed to detect, significantly reduced the correlation between working memory performance and fluid intelligence. Thus, despite being probed on the same items, with the same foils, at the same set size, only working memory displays that allowed for the strategic use of both item and ensemble representations correlated with fluid intelligence. These results provide evidence against the hypothesis that simple storage alone is related to fluid intelligence. They also demonstrate that participants make use of more complex and structured representations rather than solely individual item representation, and that strategic utilization of these representations is what correlates strongly with fluid intelligence.

Working memory is a core cognitive system that actively maintains information in an accessible state to support a wide range of behaviors (Baddeley, 2000; Cowan, 2005; Schurgin, 2018). One of the key features of working memory is that the storage capacity is extremely limited, with performance dropping quickly when people are asked to hold in mind more than just a few items (Cowan, 2005; Ma, Husain & Bays, 2014; Brady et al., 2011). Interestingly, these performance limitations of working memory have been shown to strongly correlate with fluid intelligence (Gf), the ability to solve novel abstract reasoning problems (e.g., Conway et al., 2002). However, many of the working memory tasks used in this research have focused on complex span measures -- tasks that involve a high degree of manipulation and processing in addition to the short-term storage of items (e.g., Daneman & Carpenter, 1980; Turner & Engle, 1989; Süß, Oberauer, Wittmann, Wilhelm, & Schulze, 2002). As a result, these strong correlations have traditionally been thought to be driven largely by the executive demands of the task or of the memory access strategy (e.g., Unsworth & Engle, 2007) rather than short-term storage capacity per se (Conway et al., 2002; Engle, 2002, 2018; Engle & Kane, 2004).

However, a separate line of work has argued that such correlations do not arise solely from tasks which require complex manipulation tasks, but are also present in simple working memory tasks with few strategic components (e.g., Cowan et al., 2005; Fukuda et al., 2010; Unsworth et al., 2014; Colom et al., 2008). This has led many researchers to argue that “pure” measures of working memory capacity (e.g., change detection with simple stimuli) also correlate with intelligence.

Some of this work has gone even further, claiming that a particular aspect of short-term storage capacity -- ‘number of items remembered’ -- is related to fluid intelligence. For example, in a study done by Fukuda and colleagues (2010) based on a similar task designed by Awh et al. (2007), participants performed a visual working memory task where participants had to hold in mind over a brief delay a display of simple shapes, a mix of ovals and rectangles. After the delay, they had to detect whether a change occurred. If a change did occur, this change could be a “small,” within-category change (i.e., an oval changed to another oval) or a large, cross-category change (i.e., an oval would change to a rectangle). Fukuda et al. (2010) found that only the ability to detect large changes -- which they argued only measured “how many” items were

remembered -- correlated with fluid intelligence, whereas the ability to detect small changes, which they argued as depending on another factor (the “precision” of the memories) did not relate to intelligence.

Structured memory representations even in simple displays

What drives the connection between such ‘simple’ storage tasks and fluid intelligence? If simple storage capacity alone is related to fluid intelligence, this provides a fundamentally different interpretation on the relationship between working memory and fluid intelligence than is suggested by a previous large body of work using complex span tasks (e.g., Engle & Kane, 2004). Specifically, it suggests working memory capacity (as opposed to other factors of working memory, such as executive function or attention) may be a primary constraint on measures of fluid intelligence. However, this argument depends strongly on the assumption that simple storage tasks involve minimal processing, chunking, and other strategic elements (Cowan et al., 2005). Intuitively, many of these tasks appear to be straightforward enough that this assumption is often made (e.g., remember a few colored squares and detect changes to them after a 1s delay).

Critically, however, a separate large body of literature has shown participants make use of complex strategies and utilize structured memory representations, even in such relatively simple displays (Brady & Alvarez, 2015a, 2015b; Brady & Tenenbaum, 2013; Nassar, Helmers & Frank, 2018; Chunharas et al., 2019; Orhan & Jacobs, 2013; Morey et al., 2015; Sims, Jacobs, & Knill, 2012; Schurgin & Brady, 2019; Son, Oh, Kang & Chong, 2019). For example, Brady & Alvarez (2015a) showed that even with just 3 colored circles to remember, participants rely on chunking and summary statistic representations; and that these representations are used extensively to support performance at higher set sizes (e.g., Brady & Tenenbaum, 2013). Even memory tasks where participants only have to remember two items (far below typical claims about capacity limits), they still may make use of relational structure between the items under certain circumstances, such as when items are similar (Chunharas et al., 2019; Bae & Luck, 2017). Thus, it remains unresolved how pure a measure of storage even the simplest working memory tasks are. Participant performance may vary not only due to some difference in raw

“capacity” but also because they differentially make use of chunking, ensemble representations, summary statistics, relational encoding or other strategies that help improve performance.

More complex tasks designed to tap different working memory storage constructs (‘number of items’ and ‘precision,’ e.g., Fukuda et al., 2010; Awh et al., 2007; Barton et al., 2009) are subject to even greater difficulties of this kind. The majority of the best fitting working memory models disagree with even the basic distinction these papers make between ‘represented items’ varying in precision and ‘unrepresented’ items. Most current models argue participants represent information about all objects even when many are presented, though less accurately than when only a few are present (e.g., van den Berg et al., 2012; Fougner et al., 2012; Bays, 2014, 2015; Ma, Husain & Bays, 2014; Schurgin, Wixted & Brady, 2019). Even allowing for such distinctions between ‘how many items are represented’ and ‘precision’, recent evidence has shown that the techniques used to study this -- asking participants to detect either ‘small’ or ‘large’ changes, to differentially tap ‘precision’ vs. ‘number of represented items’ -- both introduces difficulties in quantifying these concepts (e.g., Morey et al., 2012) and almost invariably allows participants to use more structured memory representations, opening the door to a wide variety of strategies whose usage may differ between participants (e.g., Brady & Alvarez, 2015b; Schurgin & Brady, 2019).

For example, Brady and Alvarez (2015b) showed that when asked to detect large changes (cross-category; a 3D cube to a Chinese character, for example), participants relied quite heavily on the global “texture” of the display (Fig. 1, 2; i.e., the idea that the “dark” part of the display was on the left, for example), which they argued was similar to how participants store spatial ensemble information in memory (e.g., Alvarez & Oliva, 2009). Alvarez and Oliva (2009) found that people naturally store -- with little to no attentional cost -- information about the texture structure of displays (i.e., the orientation information present in approximate spatial locations). Brady and Alvarez (2015b) suggested that people do this naturally, and that since such information can be remembered independently of any specific information about the objects participants just saw, participants use this separate, non-item based information to detect ‘large’ changes in many cases. Brady and Alvarez (2015b) also demonstrated that limiting the use of such texture representations dramatically decreases performance with ‘large’ changes while having little

impact on ‘small’ changes. This finding suggests that, contrary to ideas that such large, cross-category changes measure ‘how many’ objects participants can remember (Awh et al., 2007; Fukuda et al. 2010), in fact, they appear to measure in large part how adaptive participants are at making use of multiple distinct kinds of visual representations to help detect changes (see also Schurgin & Brady, 2019).

The current study

In the current study, we sought to examine whether the strategic use of ensemble or texture-based structure suggested by Brady and Alvarez (2015b) and Schurgin and Brady (2019) explains why large cross-category changes appear to have a special relationship to fluid intelligence (e.g., Fukuda et al., 2010; Unsworth et al., 2014). We hypothesized that the relationship between fluid intelligence and detection of cross-category changes is primarily a product of the strategic use of the texture of the display -- consistent with executive accounts between the connection of working memory performance and intelligence (e.g., Engle, 2002) -- rather than because of some ‘maximum number of items’ a person can hold in mind (as suggested by Fukuda et al., 2010).

To examine this, we conducted a study suggested in the discussion sections of Brady and Alvarez (2015) and Schurgin and Brady (2019), where we compared how well participants could detect large, cross-category changes in relatively homogenous displays, like those used in previous work (Awh et al., 2007), in relation to how well they could detect the exact same changes in heterogeneous displays that attempt to significantly reduce the influence of grouping and ensemble structure (see Fig 1. and Fig. 2). We hypothesized that performance for cross-category changes in the homogenous displays would be correlated with fluid intelligence, as in Fukuda et al. (2010), since this is where strategic use of ensemble information is a large part of performance. However, we hypothesized performance for the exact same changes in heterogeneous conditions (where such strategies are less useful) would be not only reduced, but also significantly less related to fluid intelligence. This is contrary to the predictions of Awh et al. (2007) and Fukuda et al. (2010), which argue such changes should continue to measure the exact same quantity (“number of represented items”).

Figure 1. The conditions of the working memory task were designed to vary the utility of spatial ensemble representations, with the critical hypothesis being that strategic use of such representations, rather than core “working memory capacity,” might underlie the correlation between fluid intelligence and large changes previously observed by Fukuda et al. (2010). In the homogenous displays, participants were shown only cube and characters, where allows for the strategic use of ensemble information to be particularly informative. In the heterogeneous displays, participants were shown cubes, characters, polygons and line drawings, which greatly reduced the utility of ensemble information relative to the homogenous displays. Critically, the size of the change at test (small, within-category or large, across-category) was the same across both conditions.

Figure 2. Schematic of a spatial ensemble account of working memory performance. Here, participants encode a display of cubes and characters and, at test, are presented a display with an across-category change (i.e., character to cube). An observer could rely on an item representation (i.e., memory of that particular character) to correctly indicate a change has occurred. However, in the absence of a strong item memory or in addition to this memory, the observer could utilize an ensemble representation, such as they may know in a particular area of the display (outlined in red) there were no cube-like features present, and subsequently use this information to correctly indicate a change has occurred. This is a particularly useful strategy in homogenous displays. Thus, we hypothesized these ensemble strategies could be driving the correlations between working memory capacity estimates and fluid intelligence, suggesting that executive function is still critical to this relationship (adapted from Schurgin & Brady, 2019).

Method

Participants

Based on a power calculation using the degree of correlation between working memory and fluid intelligence observed by Fukuda et al. (2010), we aimed for a final sample of 100 participants. Four participants were excluded and replaced (see below), so a total of 104 University of California, San Diego undergraduates participated in Experiment 1. All participants reported normal or corrected-to-normal visual acuity and no previous experience reading or writing

Chinese/ Japanese characters. Participation was voluntary and in exchange for course credit. The experimental protocol was approved by the University of California, San Diego institutional review board (IRB).

Overall procedure

Our experiment consisted of 3 tasks (1 fluid intelligence task; 2 different working memory tasks). Each task took approximately 45 minutes. Pilot data indicated that when the two working memory tasks were done on the same day, participants may use strategies developed in one task in the other. Because our central hypothesis concerned the use of such strategies, and in order to reduce fatigue in our participants, participants performed one of these 3 tasks on 3 separate days, with all three occurring less than a week apart. Thus, all subjects completed a fluid intelligence task on one day, and one of two visual working memory tasks on the other two days, in an order counterbalanced across participants.

Fluid intelligence measure

We used the Raven's Advanced Progressive Matrices (RAPM) test as a measure of fluid intelligence. The RAPM is a standardized paper and pencil test, in which participants must pick which of eight figures is the best fit as the missing part of a larger pattern. Participants began with a practice task (Set 1), which lasted 5 minutes and consisted of 12 questions. This practice test was simple and introduced subjects to the format of questions they would encounter in the main task. The main task (Set 2) consisted of 36 questions, and consistent with Fukuda et al. (2010), we gave participants 30 minutes to complete it. In this set, the patterns become increasingly complex as the questions go on. We used the number correct on this main task as our dependent measure for fluid intelligence.

Working memory measures

To measure working memory capacity, we adapted the tasks of Brady and Alvarez (2015) Experiment 1 (homogenous display task) and Experiment 2 (heterogeneous display task), each of which contained both across-category and within-category changes. In particular, in both tasks participants had to remember 8 'complex' objects, and were then probed on their memory for a single item using a change detection task. In the homogenous display task, the items were four

3D cubes and four Chinese characters, and the foils consisted of either the same category item that had been present at that location or an item from the other category. In the heterogeneous display task, the items consisted of two 3D cubes, two Chinese characters, two Snodgrass objects, and two gray polygons. Again, a given foil during the change detection task could be the same category as the item that had been presented at that location or a different item. Participants completed these two tasks -- homogeneous and heterogeneous displays -- in different sessions on different days to reduce the chance of strategies from one task influencing their performance on the other task.

In both working memory tasks, participants first saw a white rectangle on an otherwise grey screen with a central fixation point for 1200ms. Then they were briefly presented with objects of eight different objects inside this rectangle for 1000ms. These objects were randomly located subject to the constraint that they did not overlap and that there were 2 objects in each quadrant of the display. The eight objects disappeared, and after a 1s delay, a single probe object appeared, and participants needed to determine whether that exact image was the image that had been presented at that location. Participants used the keyboard to indicate whether the object was the same ('s') or different ('d') than what they had seen in that location (see Fig. 1). This probe could be either the same object that had been seen at that location (50% of the time), or, if not the same object, either from the same category (cube/character etc) as the encoded stimulus at that location (25% of total trials) or a different category (25%). In the heterogeneous display task, if the object changed category, cubes always changed to characters and vice versa; and Snodgrass objects always changed to polygons and vice versa. This allowed us to analyze the exact same set of changes in the homogeneous and heterogeneous display conditions. To analyze within-category changes, we also focused on only the comparable trials in the two groups: those where cubes or characters were tested.

In each overall condition (homogeneous/heterogeneous display), there were 200 trials, 100 'same' trials, 50 within-category foils and 50 across-category foils. Figure 1 illustrates a single trial.

Exclusion criteria

Based on a priori criteria we set, we excluded and replaced participants whose d' was less than 0 in either across-category change condition. This resulted in the exclusion of 3 of the 104 participants. We also excluded and replaced 1 participant for having a RAPM score (8) that was a significant outlier ($M=23.7$, $SD=3.96$; range excluding that participant: 14-33), leaving a final sample of 100 participants.

Figure 3. Results of fluid intelligence (violin plot; shows full distribution of scores) and working memory conditions (error bars are within-subject S.E.M.). Overall, participants performed better with across-category changes and were more accurate at across-category changes in the homogenous display condition, consistent with previous work on how participants use spatial ensemble structure in such displays (Brady & Alvarez, 2015; Schurgin & Brady, 2019).

Results

Participants performed well on the RAPM ($M=23.6$, $SD=3.9$; range: 14-33 out of 36), and, as expected, participants performed much better on across-category comparisons than within-category comparisons in both working memory display types ($ps<0.001$). Replicating Brady and Alvarez (2015), participants performed significantly better on identical cross-category changes

(cubes->characters and vice versa) when the display was overall more homogenous. In particular, a 2x2 repeated measures ANOVA revealed main effects of within vs. cross-category changes ($F(1,99)=726.4$; $p<0.0001$), a main effect of homogenous vs. heterogeneous ($F(1,99)=14.6$; $p<0.001$) and as expected, an interaction ($F(1,99)=34.3$; $p<0.0001$). Participants in the homogenous condition were better at cross-category changes ($t(99)=5.48$, $p<0.001$, $d_z=0.55$), but this was not true for within-category changes ($t(99)=0.72$, $p=0.47$, $d_z=0.07$). This is consistent with previous work showing that participants make use of additional spatial ensemble structure in homogenous displays that is particularly informative for the cross-category changes (Fig. 2; Brady & Alvarez, 2015; Schurgin & Brady, 2019).

	RAPM	Homogenous Cross-category	Homogenous Within-category	Heterogeneous Cross-category	Heterogeneous Within-category
RAPM	--				
Homogenous Cross-category	0.30**	--			
Homogenous Within-category	0.29**	0.33**	--		
Heterogeneous Cross-category	0.13	0.54***	0.19	--	
Heterogeneous Within-category	0.07	0.18	0.16	0.40***	--

Table 1. Correlations between tasks. * $p < .05$, ** $p < .01$, *** $p < .001$.

We next turned to the correlation structure between the tasks. Consistent with our predictions, we found that the homogenous cross-category task was significantly correlated with RAPM ($r=0.30$, $p=0.002$), whereas the heterogeneous cross-category task was not significantly correlated with RAPM ($r=0.13$, $p=0.18$), and the difference between these two correlations was significant (Steiger, 1980; Meng et al., 1992; $z=1.76$, $p=0.039$). That is, while the two cross-category conditions were strongly correlated ($r=0.54$; $p<0.001$), it appears to be the non-shared component -- the unique variance present only in the homogenous cross-category task -- that drives this correlation with RAPM. In particular, the homogenous cross-category task still correlated equally well with RAPM after partialing out variance from the heterogeneous cross-

category task ($r=0.27$, $p=0.008$), whereas the heterogeneous cross-category task still did not correlate with RAPM after partialing out variance from the homogenous cross-category task ($r=-0.04$, $p=0.69$). This is consistent with our hypothesis that the homogenous task contains an additive factor relative to the heterogeneous cross-category task -- strategic use of spatial ensemble structure -- that is the dominant reason this task correlates with fluid intelligence.

Contrary to our initial predictions, we found nearly the same results for the within-category changes. That is, contrary to the claims of Fukuda et al. (2010), who claimed a selective correlation between large changes and fluid intelligence with nearly no relationship between small changes and fluid intelligence, we found nearly identical correlations for within- and across-category changes (Table 1). Despite the fact that participants performed about equally well in the homogenous within-category and heterogeneous within-category tasks (consistent with ensemble structure not being particularly informative in these tasks; Brady & Alvarez, 2015; Schurgin & Brady, 2019), performing better in the within-category changes in the homogenous condition was predictive of RAPM score ($r=0.29$, $p=0.004$), but not in the heterogeneous condition ($r=0.07$, $p=0.50$), and the difference between these correlations was again significant ($p=0.04$, $z=1.70$).

Thus, performance in both aspects of the homogenous display type -- where more strategic encoding decisions were available overall -- was predictive of RAPM, but not performance in the heterogeneous displays. This was true despite the same overall working memory load in the heterogeneous and homogeneous displays and despite participants being probed on identical tests of identical objects in both cases.

General Discussion

We found strong evidence that a measure of fluid intelligence (RAPM) correlates differentially with performance in two extremely similar working memory tasks. Despite being probed on the same items, with the same foils, at the same set size, only the working memory task where multiple relevant representations could be used (item-based and ensemble-based) correlated with fluid intelligence. This is consistent with our hypothesis that strategic encoding decisions are

responsible for a significant portion of the correlation between working memory performance and fluid intelligence. Our results are contrary to claims that working memory storage capacity alone underlies this correlation, or the more particular claim that only the ‘quantity’ measure of working memory storage and not the ‘precision’ measure of storage is correlated. Our results are consistent with the broader literature arguing this relationship is driven by the executive attention aspects of working memory, and in addition demonstrates a novel way executive processes might be utilized in this relationship (i.e., utilizing global ensemble representations) (Engle, 2002; Conway et al., 2002).

Broadly, these results reinforce the importance of considering non-item based information in working memory studies. While displays of segmented objects appear to be distinct objects, and thus it is natural to appeal to memory representations of single items, participants’ memory representations may not treat them this way. Here, we provide a new kind of evidence that working memory researchers must be careful to dissociate the contributions of item-level and ensemble-level representations to working memory performance (Brady & Alvarez, 2011; Orhan & Jacobs, 2013; Brady & Tenenbaum, 2013; Brady & Alvarez, 2015b), and highlight the importance of taking into account the potential use of global ensemble information or explicit chunking strategies (Nassar, Helmers & Frank, 2018) in both visual working memory tasks and the models used to explain those tasks.

A surprising aspect of the current data is that despite equivalent performance on small within-category changes in the homogeneous and heterogeneous conditions, performance for the small, within-category changes in the homogeneous condition correlated with fluid intelligence. While this is not a necessary prediction of our account, this does appear to be consistent with differential use of ensemble-based strategies by participants with different RAPM scores. In particular, if all participants switched to a completely texture/group-based encoding in the homogeneous condition, and thus did not represent individual item information at all, this should have *impaired* performance in the within-category changes in the homogenous relative to the heterogeneous condition. Thus, finding the same overall level of performance in this condition may hide heterogeneity: participants with low RAPM scores may on average abandon item representations completely, causing lower within-category performance and higher across-

category performance; by contrast, participants with higher RAPM scores may be more likely to choose a different and appropriate encoding strategy for each display, continuing to represent individual items as appropriate. Thus, if the items are generally spatially clustered (cubes all on one side, for example), these participants may default to a more ensemble or group-based encoding. By contrast, on displays where the items are scattered and the overall texture of the display is of little use, these participants may encode more individual item information. Thus, overall performance on the within-category condition could reflect a blend of participants who largely switched to an ensemble/group-based strategy and those who more fluidly remembered item information, causing greater performance on the within-category changes to be indicative of higher RAPM scores while resulting in no overall increase in performance in this condition.

Critically, this result also fails to replicate the distinction between small and large changes observed by Fukuda et al. (2010) in predicting fluid intelligence, since in our data detection of both kinds of changes was related to RAPM scores. What explains the discrepancy between our results and those of Fukuda et al. (2010) with respect to the correlation between small change and fluid intelligence? It is possible that Fukuda et al. (2010) simply had less power to detect such a relationship. However, one possibility noted by Fukuda et al. is that their experiment had twice as many large change trials as small change trials, so strategic use of this fact may have been more common among participants with higher fluid intelligence. This is broadly consistent with our account that strategic awareness may explain the overall correlations with intelligence in these tasks. In retrospect, it seems surprising even under the “slot-based” hypothesis of Fukuda et al. that there would be no correlation between detecting small changes and detecting large changes, as in that framework performance detecting both kinds of changes is necessarily limited by having at least some item representation, and their experiment had half the trials at set size 8. Even according to their own theory, most of the variance in the small change condition at set size 8 should arise from “quantity” rather than “quality.” A specific correlation with large changes also appears to be somewhat implausible in frameworks that reject the “number” vs. “precision” distinction, since lower accuracy when discriminating more difficult changes appears to be directly predictable from a fixed perceptual similarity measure of how much more difficult these changes are (Schurgin, Wixted & Brady, 2018) -- that is, small and large changes appear to

fundamentally measure the same underlying item-based capacity if ensemble/grouping strategies are limited.

Overall, we find no evidence for a distinction between “quantity” and “quality” in the relationship between fluid intelligence and working memory capacity. More broadly, we believe that while the current result is a case study in the difficulty of isolating simple storage, similar effects may be found in much of the other work that appears to demonstrate that fluid intelligence is correlated with short-term storage capacity above-and-beyond executive attention and cognitive control (e.g., Cowan et al., 2005), as such previous work has likely significantly underestimated the complexity of participants memory representations and strategies in ‘simple’ working memory displays (Brady & Alvarez, 2015a, 2015b; Brady & Tenenbaum, 2013; Nassar, Helmers & Frank, 2019; Chunharas et al., 2019; Orhan & Jacobs, 2013; Morey et al., 2015; Sims, Jacobs, & Knill, 2012; Schurgin & Brady, 2019).

References

- Alvarez, G. A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences*, 106(18), 7345-7350.
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, 18, 622– 628.
- Baddeley, A. (2000). The episodic buffer: a new component of working memory?. *Trends in Cognitive Sciences*, 4(11), 417-423.
- Bae, G. Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception, & Psychophysics*, 79(8), 2376-2395.
- Barton, B., Ester, E. F., & Awh, E. (2009). Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1359.

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

Bays, P. M. (2015). Spikes not slots: noise in neural populations limits working memory. *Trends in cognitive sciences*, 19(8), 431-438.

Brady, T. F. and Alvarez, G.A. (2015a). Contextual effects in visual working memory reveal hierarchically structured memory representations. *Journal of Vision*, 15(15):6.

Brady, T. F. and Alvarez, G.A. (2015b). No evidence for a fixed object limit in working memory: Ensemble representations inflate estimates of working memory capacity for complex objects. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 41(3), 921-9.

Brady, T. F., & Tenenbaum, J. B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, 120(1), 85.

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5), 4-4.

Chunharas, C., Rademaker, R. L., Brady, T.F., and Serences, J. (2019). Adaptive distortions in visual working memory. *PsyArXiv Preprint*: <https://doi.org/10.31234/osf.io/e3m5a>

Colom, R., Abad, F. J., Quiroga, M. Á., Shih, P. C., & Flores-Mendoza, C. (2008). Working memory and intelligence are highly related constructs, but why?. *Intelligence*, 36(6), 584-606.

Conway, A. R., Cowan, N., Bunting, M. F., Theriault, D. J., & Minkoff, S. R. (2002). A latent variable analysis of working memory capacity, short-term memory capacity, processing speed, and general fluid intelligence. *Intelligence*, 30(2), 163-183.

Cowan, N. (2005). *Working Memory Capacity*. Psychology Press.

Cowan, N., Elliott, E. M., Sauls, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, 51(1), 42-100.

- Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of verbal learning and verbal behavior*, 19(4), 450-466.
- Engle, R. W. (2002). Working memory capacity as executive attention. *Current Directions in Psychological Science*, 11(1), 19-23.
- Engle, R. W. (2018). Working memory and executive attention: A revisit. *Perspectives on Psychological Science*, 13(2), 190-193.
- Engle, R. W., & Kane, M. J. (2004). Executive attention, working memory capacity, and a two-factor theory of cognitive control. *Psychology of Learning and Motivation*, 44, 145-200
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, 3, 1229.
- Fukuda, K., Vogel, E., Mayr, U., & Awh, E. (2010). Quantity not quality: The relationship between fluid intelligence and working memory capacity. *Psychonomic Bulletin & Review*, 17(5), 673–679. <http://doi.org/10.3758/17.5.673>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17(3), 347.
- Meng, X. L., Rosenthal, R., & Rubin, D. B. (1992). Comparing correlated correlation coefficients. *Psychological Bulletin*, 111(1), 172.
- Morey, C. C., Cong, Y., Zheng, Y., Price, M., & Morey, R. D. (2015). The color-sharing bonus: Roles of perceptual organization and attentive processes in visual working memory. *Archives of Scientific Psychology*, 3(1), 18.
- Morey, R. D., Morey, C. C., Brisson, B., & Tremblay, S. (2012). A critical evaluation of c as a measure of mnemonic resolution. *Journal of Experimental Psychology: Human Perception and Performance*, 38(4), 1069.
- Nassar, M.R., Helmers, J. & Frank, M.J. (2018). Chunking as a rational strategy for lossy data compression in visual working memory. *Psychological Review*, 125, 486-511.

Orhan, A. E., & Jacobs, R. A. (2013). A probabilistic clustering theory of the organization of visual short-term memory. *Psychological Review*, 120, 297–328

Schurgin, M. W. (2018). Visual Memory, The Long and the Short of it: A Review of Visual Working Memory and Long-Term Memory. *Attention, Perception, & Psychophysics*, 80(5), 1035-1056.

Schurgin, M. W., and Brady, T.F. (2019). When "capacity" changes with set size: Ensemble representations support the detection of across-category changes in visual working memory. *Journal of Vision*, 19(5):3.

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2018). Psychophysical scaling reveals a unified theory of visual memory strength. *BioRxiv*, 325472.

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2012). An ideal observer analysis of visual working memory. *Psychological Review*, 119, 807– 830

Son, G., Oh, B. I., Kang, M. S., & Chong, S. C. (2019). Similarity-based clusters are representational units of visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

Steiger, J. H. (1980). Tests for comparing elements of a correlation matrix. *Psychological Bulletin*, 87(2), 245.

Süß, H. M., Oberauer, K., Wittmann, W. W., Wilhelm, O., & Schulze, R. (2002). Working-memory capacity explains reasoning ability—and a little bit more. *Intelligence*, 30(3), 261-288.

Turner, M. L., & Engle, R. W. (1989). Is working memory capacity task dependent? *Journal of Memory and Language*, 28(2), 127-154.

Unsworth, N., & Engle, R. W. (2007). On the division of short-term and working memory: an examination of simple and complex span and their relation to higher order abilities. *Psychological Bulletin*, 133(6), 1038.

Unsworth, N., Fukuda, K., Awh, E., & Vogel, E. K. (2014). Working Memory and Fluid Intelligence: Capacity, Attention Control, and Secondary Memory Retrieval. *Cognitive Psychology*, 71, 1–26. <http://doi.org/10.1016/j.cogpsych.2014.01.003>

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*, 109(22), 8780-8785.