

Features and Conjunctions in Visual Working Memory

Weiwei Zhang¹, Jeffrey S. Johnson², Geoffrey F. Woodman³, and Steven J. Luck¹

¹Center for Mind & Brain, University of California, Davis

²Department of Psychology, University of Wisconsin, Madison

³Department of Psychology, Vanderbilt University

As discussed throughout this volume, Treisman and her colleagues have proposed that focused attention integrates massively parallel sensory inputs together to form bound object-based representations. As a result, the features of the environment, which are processed in separate sensory channels at early stages of the visual system, are ultimately perceived as integrated objects. When focused attention is absent or distracted away, features cannot be precisely localized and combine haphazardly. This leads to perceptual errors in healthy individuals when their attention is distracted (see Chapter XXX, this volume), and it leads to a chaotic visual experience for individuals with bilateral damage to the attentional control systems of the parietal lobes (see Chapters XXX and YYY, this volume).

Although the research described in this volume provides overwhelming evidence that attention binds features together during perception, this does not solve the entire binding problem faced by the brain, because perceptual representations fade after the sensory input terminates, and they are easily overwritten by the visual transients that occur every time we blink or move our eyes (Irwin & Thomas, 2008; Simons & Rensink, 2005). More stable representations of bound objects are needed to support reasoning, decision making, and other aspects of behavior that are extended over time (Cowan, 2005). These stabilized representations are called visual short-term memory (VSTM) or visual working memory (VWM) representations (see review by Luck & Hollingworth, 2008).

In this chapter, we will consider how features and conjunctions are represented in VWM, an issue that applies the conceptual framework of Feature Integration Theory to the storage of information in memory. In particular, we will focus on the role of attention in maintaining bindings of features in VWM, which was brought to the forefront of VWM research by Wheeler and Treisman (2002, accompanying article).

Storage of Features and Conjunctions in Visual Working Memory

The storage of features and conjunctions in VWM has typically been studied in change detection tasks. Consider, for example, the task and results shown in Figure 1 (from the study of Luck & Vogel, 1997; see also Vogel, Woodman, & Luck, 2001). On each trial, observers were shown a 100-ms *sample array* consisting of colored, oriented rectangles, followed after a 900-ms retention interval by a 2000-ms *test array*. The test array was either identical to the sample array or differed in terms of a single feature (e.g., the red vertical bar in the sample array changed to a green vertical bar in the test array). In the *color-only* condition, observers were instructed to remember the colors of the bars and make an unspeeded button-press response to indicate whether or not a color change was detected in the test array. Only color could change in this condition. In the *orientation-only* condition, observers were instructed to remember orientation because only orientation could change. If observers are able to store 3 of the items in VWM, then the changes should be accurately detected on virtually every trial when the displays contain 3 or fewer items (assuming that the changes are sufficiently discriminable). If the displays contain 4 items, however, then the item that changed will be one of the items that was stored in VWM on only $\frac{3}{4}$ of trials, and the observers should be unable to detect the change on $\frac{1}{4}$ of trials. Indeed, change detection accuracy for color or orientation is typically near ceiling for arrays of 1-3 items and then falls systematically as the set size increases. Moreover, recent evidence indicates that observers maintain precise representations of the items that are remembered, but they store no information at all about the other items, as if VWM storage capacity is determined by a limited number of slots (Zhang & Luck, 2008).

Is storage capacity limited by the number of features that must be retained (i.e., one feature per memory slot), or is it instead limited by the number of objects (i.e., one object per slot)? To answer this question, an *either* condition was tested in which either color or orientation might change on a given trial. To perform this task successfully, observers must store both the color and the orientation of each item in VWM, because they do not know which feature will change on a given trial. If VWM capacity is determined by the number of features that must be retained, then change detection accuracy should be worse for a given number of objects in the either condition than in the single-feature conditions. If, however, VWM capacity is determined by the number of objects that must be retained, then performance should be the same across all three conditions. As shown in Figure 1, change detection accuracy was indeed the same in the either condition as in the color-only and orientation-only conditions, consistent with the hypothesis the VWM capacity is limited by the number of objects (the *integrated object hypothesis*)ⁱ.

An alternative explanation is that color and orientation are stored completely independently so that features from different dimensions simply do not compete for the limited capacity of VWM. This could account for the comparable change detection accuracy observed by Luck and Vogel (1997) in the single-feature and either conditions. To address this possibility, Luck and Vogel tested a condition in which each object was composed of two colored squares. With these two-color objects, performance was again limited by the number of objects rather than by the

number of features. However, this pattern has been difficult to replicate (Delvenne & Bruyer, 2004; Olson & Jiang, 2002; Wheeler & Treisman, 2002; Xu, 2002).

Wheeler and Treisman (2002) provided much stronger evidence for the binding of features in VWM. This study included the single-feature and either conditions of the Luck and Vogel (1997) study, but using shape and color (or location and color) as the two features. Wheeler and Treisman also added a *binding* condition, illustrated in Figure 2. In the *single-probe test* version of the task, a single item was presented at the center of the display at the time of test (see upper half of Figure 2). In the binding condition, this item either consisted of an item from the sample array or an item that combined two features that had been present in different objects in the sample array. For example, a red circle and green cross might be presented in the sample array, and a changed item would be a red cross. Under these conditions, subjects detected binding changes just as well as they detected feature changes (see Figure 8D in Wheeler & Treisman, 2002). These results are consistent with the integrated object hypothesis.

However, different results were obtained in the *whole-array test* version of the task, in which the test array contained all the items from the sample array (see bottom half of Figure 2). In this condition, change detection performance was impaired for the binding condition relative to the single-feature conditions (see Figure 8B in Wheeler & Treisman, 2002). Note, however, that accuracy was only a few percent worse in the binding condition than in the either condition. Thus, this study demonstrated that observers can retain bound objects just as well as individual features (when a single item is present in the test array) or almost as well as individual features (when multiple items are present in the test array).

The Role of Attention in Maintaining Feature Bindings

Although the results of Wheeler and Treisman (2002) supported the conclusion that bound objects can be stored in VWM, it is important to consider why accuracy was reduced for the binding condition in the whole-display version but not in the single-probe version. Wheeler & Treisman (2002) proposed that attention is required to maintain bindings in VWM and that attention is disrupted by the appearance of multiple objects in the test array in the whole-display version. That is, attention presumably needs to be divided across different items in the test array in the whole-display version so that each item in the test array can be compared to the corresponding item in the sample array. Consequently, attention is distracted away from maintaining feature bindings in VWM, causing the bindings to fall apart. Less division of attention is necessary when the test array contains a single item, according to this account, and so there was little or no disruption of the bindings in the single-probe condition.

Wheeler & Treisman proposed a dual-storage mechanism in VWM. That is, features are stored in independent systems in VWM and can be maintained without focused attention, whereas attention is required to maintain bindings of features. The parallel storage of features could have two complementary manifestations. On the one hand, different features within a given dimension compete for the limited capacity for that dimension; on the other hand, features from different dimensions do not compete for limited capacity. Feature bindings are maintained separately from features themselves, and the bindings require deliberate maintenance through focused attention. As a consequence, when attention is degraded or distracted, feature bindings are lost but the individual features are not. Just as attention is responsible for initially gluing features together in perceptual encoding, this account proposes that attention keeps features glued together during VWM maintenance. In contrast, the integrated object hypothesis of Luck and Vogel (1997) proposes that attention may be necessary for creating integrated object representations during perception, but attention plays no special role in maintaining bindings, compared to features, in VWM.

Testing the Role of Attention in Maintaining VWM Bindings

As they acknowledged in their paper, Wheeler & Treisman (2002) did not explicitly manipulate attention, and their results served as an inspiration for, rather than a test of, their hypothesis about the role of attention in VWM. This interpretation influenced subsequent studies by other investigators who directly tested this hypothesis by using a secondary task to explicitly disrupt attention during the maintenance interval of a VWM task. If attention plays a special role in maintaining bindings in VWM, then this disruption of attention should impair VWM performance more when the VWM task requires maintenance of bindings than when it requires the maintenance of simple features.

In the first of these studies, Yeh, Yang, and Chiu (2005) presented a small letter either at fixation or at a peripheral location during the retention interval of a change detection task. Subjects made an immediate response to letter, indicating whether it was a vowel, and they made an unspeeded change/no-change response at the end of the trial for the change detection task. They also included trials on which the change detection task was performed without the letter classification task. The change detection task required either the detection of two new color values in the test display or a change in color-location binding (i.e., swapping the colors of two items, leading to a test display that contains the same colors and the same locations but different color-location bindings). Yeh et al. assumed that discriminating the peripheral letter would require a shift of spatial attention, which should specifically impair change detection performance in the color-location binding condition if attention plays a specific role in maintaining bindings in VWM. They found that change detection performance was indeed worse for the peripheral-letter condition compared to the central-letter and no-letter conditions. However, this decline in performance for the peripheral-letter condition was modest and was equivalent

across the color-only and color-location binding tasks, indicating that attention plays no special role in maintaining bindings in VWM.

In a similar study, Gajewski and Brockmole (2006) presented an exogenous attention cue (the abrupt onset of a dot) at the location of one of the items in the memory array during the retention interval of a color-shape change detection task. Attention was indeed allocated to the cued item since participants' recalled the cued item much better than the uncued items. When memory at an uncued location was tested, participants either remembered both the color and the shape of the tested item or neither feature, suggesting that capture of attention by the abrupt onset during the retention interval did not make the feature binding collapse into unbound features.

To provide a stronger manipulation of attention, Johnson, Hollingworth, and Luck (2008) used a highly demanding visual search task during the retention interval (see Figure 3A). Specifically, participants searched for a target square with a gap on either the left or right side among distractor squares with a gap on either the top or bottom. Either the left-gap or right-gap target was present on every trial, and subjects made a speeded two-alternative forced-choice response to report the target's gap position. This is a demanding search task that typically yields search slopes of approximately 60 ms/item (Woodman, Vogel, & Luck, 2001), and electrophysiological studies have verified that attention shifts serially among the distractors until the target is found (Woodman & Luck, 1999, 2003). Just as in the Yeh et al. (2005) study, Johnson et al found that distracting attention during the maintenance interval of the VWM task led to a modest reduction in change detection performance, and that this reduction was equivalent for feature and binding tasks (see Figure 3B).

Allen, Baddeley, & Hitch (2006) addressed the possible role of a different variety of attention in maintaining feature bindings in VWM. Specifically, they tested whether the maintenance of bindings relies on the central executive. To accomplish this, they combined a change detection task with a simultaneous backward counting task or a near-capacity verbal working memory task. As in the studies that manipulated visuospatial attention, Allen et al. found that these executive load tasks caused a modest and equivalent decline in performance in feature and binding VWM tasks.

These four studies used a broad range of methods for distracting attention while subjects performed change detection tasks, and none of them found evidence that distraction of attention leads to a selective deficit in the representation of bindings in VWM. One study, however, has found evidence of a specific impairment in VWM for bindings when attention is distracted. Specifically, Fougne and Marois (2009) used a multiple-object tracking (MOT) task during the retention interval of a change detection task and found that the MOT task produced a larger decrement in VWM for bindings than in VWM for features. Why did they find a differential effect on bindings when no such effect was found by the other four studies reviewed above? One possibility is that the MOT task is simply more attention-demanding than the

tasks used in the other studies. This seems unlikely, because the other studies found significant effects of distraction on VWM performance—indicating that the distraction was potent—but simply failed to observe a differential effect of distraction on bindings. Moreover, the serial visual search task used by Johnson et al. (2008) should be maximally attention-demanding from the viewpoint of Feature Integration Theory, and this specific task is known to involve the serial focusing of attention (Woodman & Luck, 1999, 2003). Another possibility is that MOT involves continuous focusing of attention, whereas the other tasks might permit rapid switching between the task and the concurrent VWM maintenance process. This is also unlikely in the case of the Johnson et al. study, because any such switching of attention during visual search would lead to a large elevation in reaction time. However, reaction time for the visual search task was not elevated when subjects were maintaining bindings compared to when they were maintaining individual features. Moreover, previous research demonstrates that observers can successfully switch between simultaneous MOT and visual search tasks (Alvarez et al., 2005).

The most likely explanation for the discrepancy is that the retention interval and distraction task used by Fougne and Marois (2009) were substantially longer than those used in the other studies, with a total retention interval of 7.2 seconds and a 3-second period during which subjects performed the MOT distraction task. The other tasks involved brief stimulus presentations and shorter retention intervals. This raises two possibilities. One is that bindings fade away slowly when attention is unavailable, and a long period of distraction is therefore necessary to see a differential effect of distraction on bindings. Another possibility is that different memory mechanisms are recruited when the retention interval becomes long. Zhang and Luck (2009) have shown that VWM representations begin to fail sometime after 4 seconds. Moreover, Olson et al. (2006) found that patients with long-term memory deficits due to medial temporal lobe damage are specifically impaired at retaining bindings at an 8-s retention interval but not at a 1-s retention interval. Thus, long-term memory mechanisms may have contributed to performance in the Fougne and Marois (2009) study, and attention may play a role in the use of these mechanisms. Additional research is necessary to settle this issue. However, the existing data provide clear evidence that the withdrawal of attention does not cause an immediate collapse of bindings in VWM.

If attention is not necessary to maintain bindings in VWM, then why did Wheeler and Treisman (2002) find impaired performance for bindings in the whole-display condition but not in the single-probe condition? One possibility is that observers have a memory of the statistical properties of the sample array in addition to a memory for the individual objects. Research by Chong and Treisman (2003, 2005) indicates that people can extract statistical properties from arrays of stimuli (e.g., mean size; but see Myczek & Simons, 2008, for a different account). When faced with a simple feature task, subjects may be able to detect changes by noting differences in the statistics of the two arrays, but only if the entire array is presented at the time of test (i.e. in whole-display versions but not in single-probe versions). A memory for statistical properties would not be helpful in detecting

binding changes. Thus, participants in the study of Wheeler and Treisman (2002) may have used memory for statistical information to boost accuracy in the feature tasks in the whole-display version but not in the single-probe version. Indeed, performance was better overall in the whole-display version than in the single-probe version. Thus, the findings of Wheeler and Treisman (2002) may reflect the benefit of being able to use statistical information to detect changes in simple features rather than the cost of maintaining feature bindings in VWM. Other explanations have been offered by other investigators (see Allen, Hitch, & Baddeley, 2009; Gajewski & Brockmole, 2006; Yeh, et al., 2005).

Shared Capacity Across Feature Dimensions

Even if attention plays no special role in the maintenance of bindings in VWM, it is still possible—as proposed by Wheeler and Treisman (2002)—that each feature dimension is stored independently in VWM. This would be analogous to the separate feature maps for different feature dimensions proposed by Feature Integration Theory. Indeed, other investigators have made similar proposals regarding VWM storage. For example, there is fairly broad agreement that location information is stored at least partly independently of object identity information in VWM, especially within posterior areas of cortex (Courtney, Ungerleider, Keil, & Haxby, 1996; Logie & Marchetti, 1991). A few investigators have also provided evidence that different ventral-stream feature dimensions, such as orientation and spatial frequency, are stored independently in VWM (see, e.g., Kyllingsbaek & Bundesen, 2007; Magnussen, Greenlee, & Thomas, 1996). However, other investigators have proposed that all visual dimensions are stored together in VWM (e.g., Rainer, Asaad, & Miller, 1998), and some have even argued that working memory capacity is shared across sensory modalities (Saults & Cowan, 2007; Scarborough, 1972).

Perhaps the most straightforward procedure for assessing whether different working memory representations compete for storage capacity was developed in a classic study by Scarborough (1972). Scarborough was interested in whether visual information and verbal information were stored in the same working memory system. To test this, he asked whether the amount of information that could be retained was greater when half of the to-be-remembered information was verbal and half was visual compared to when all of it was verbal or all of it was visual. He found that performance was much better when the information was divided between visual and verbal, consistent with separate storage for these two types of information.

Woodman, Vogel, and Luck (unpublished data) used this same approach to ask whether color and orientation are stored in a common VWM system. Examples of stimuli and the task are shown in Figure 4A. On each trial, observers were shown a sequence of two arrays of three items, one above and one below the center of the screen. The first array consisted of either three gray oriented bars or three colored

squares, and the second array consisted of a second set of three gray oriented bars or three colored squares. The four possible sequences (orientation-orientation, color-color, orientation-color, color-orientation) were tested in separate trial blocks, so the subjects were fully prepared for the type of stimuli presented in each of the two sample arrays. At the end of the trial, memory was tested either for the first array or the second array; this varied randomly from trial to trial, so subjects were motivated to store both sets of items in VWM.

If color and orientation are stored independently in VWM, with no sharing of storage capacity, then observers should be able to store three colors and three orientations in VWM at the same time. If, however, VWM storage capacity is limited by the number of objects independent of the stimulus dimensions of those objects, then performance should be the same when subjects are asked to remember six colors, six orientations, or three colors plus three orientations. As shown in Figure 4B, subjects exhibited a nearly perfect sharing of capacity across color and orientation. That is, change detection accuracy was nearly identical when subjects were asked to remember six items from a single dimension (six colors or six orientations) and when they were asked to remember three items from each of two dimensions (three colors and three orientations).

This experiment provides direct evidence that different ventral-stream dimensions are not stored independently, but instead share a common storage capacity limit. Interestingly, a second experiment examining memory for color and location yielded substantially better performance when subjects were asked to remember three colors and three locations compared to when they were asked to remember six colors or six locations. This is consistent with prior research indicating that location information is stored at least partly independently of object identity information, and it demonstrates the ability of this approach to detect independence of VWM representations.

It should be noted, however, that the sharing of storage capacity by color and orientation does not indicate that color and orientation are maintained together in the same memory storage units (i.e., units that explicitly store specific color-orientation conjunctions). It is quite possible that different dimensions are stored in different units, and yet the limit on the number of items that can be simultaneously active in VWM is independent of which units are storing the information. It is therefore also possible that an observer could encode one feature of an object but fail to successfully encode the other (e.g., Woodman & Vogel, 2008), or that one feature could be lost from VWM while the other feature remains. The key is that the limiting factor is the number of items being remembered, not the features that compose these items. In the next section, we provide an explicit theory of VWM storage that explains how different dimensions can be stored separately and yet be subject to a shared capacity limit.

What limits storage capacity for features and conjunctions in VWM?

The evidence reviewed above indicates that the storage capacity of VWM is determined by the number of objects that must be remembered, not by the number of features. In this section, we will refine this hypothesis and provide an explanation of why VWM exhibits this sort of storage capacity limit. Understanding the factors that limit VWM storage capacity is of great theoretical importance, because capacity limits can inform us about the overall architecture of the cognitive system. It is also of great practical importance, because VWM storage capacity is predictive of overall cognitive functioning among normal individuals (Cowan, et al., 2005) and is reduced in individuals suffering from brain disorders (Gold, et al., 2006; J. Lee & Park, 2005).

At the neural level, visual working memory representations appear to be implemented by means of sustained neural firing during the retention interval, and this sustained firing occurs in the same neurons that are active when the stimulus is initially presented (Chelazzi, Miller, Duncan, & Desimone, 1993; Harrison & Tong, 2009). The sustained firing is assumed to be achieved by means of recurrent connections that feed a neuron's activity back onto itself (either directly or via a loop with other neurons; see Deco & Rolls, 2008). For example, a red stimulus might cause a group of red-selective neurons to fire, and a VWM representation could be created by allowing the outputs of these neurons to project to another set of neurons that then feed back to the original set of red-selective neurons. This *reverberatory* activity will then persist indefinitely after the stimulus terminates, creating an oscillatory pattern as the activity bounces back and forth like a ping-pong ball between the two sets of neurons. A potential problem is that this reverberatory activity will tend to spread to neurons coding similar feature values (e.g., orange), and this may eventually recruit all of the neurons in the system into a state of persistent oscillation. This is essentially a seizure, and it is obviously not a good thing. To prevent this kind of unchecked spreading of activation, it is possible to add lateral inhibition to the system. The right amount and pattern of inhibition can lead to very stable and precise VWM representation. Spencer and his colleagues have developed a detailed dynamic systems model of VWM that implements this form of storage, and this model has generated several interesting predictions (see, e.g., Johnson, Spencer, Luck, & Schoner, 2009; Schutte, Spencer, & Schoner, 2003; Simmering, Spencer, & Schoner, 2006).

Luck and Vogel (1997, 1998) proposed a neural synchrony hypothesis to explain how different features could be bound together into object representations, and Raffone and Wolters (2001) developed a neural network simulation of this general hypothesis. According to this proposal, the neurons representing a given object fire synchronously, forming what Hebb (1949) called a *cell assembly*. As shown in Figure 5, a display containing a red object and a blue object would be stored by allowing the red-selective neurons to fire in synchrony with each other and the blue-selective neurons to fire in synchrony with each other. Synchrony tends to

potentiate neural transmission, and synchronizing a set of neurons may therefore create a sufficiently strong recurrent signal to allow the representation to be self-sustaining. That is, the establishment of synchrony among a set of neurons may transform a transient sensory response into a persisting VWM representation. Evidence for this has been obtained both in human EEG recordings, which exhibit gamma-frequency oscillations during the retention interval of a VWM task (Tallon-Baudry, Bertrand, & Fischer, 2001; Tallon-Baudry, Bertrand, Peronnet, & Pernier, 1998; Tallon-Baudry, Kreiter, & Bertrand, 1999) and in macaque microelectrode recordings, which exhibit phase synchrony between oscillating local field potentials and spiking activity during the retention interval (H. Lee, Simpson, Logothetis, & Rainer, 2005).

This conceptualization of VWM storage also provides an explanation of storage capacity limitations. To maintain separate representations of red and blue, as illustrated in Figure 5A, it is necessary to ensure that the synchronized set of neurons coding the red item never fires at the same time as the synchronized set of neurons coding the blue item. If these two cell assemblies fire at the same time, they will either merge into a single synchronous cell assembly (representing a single intermediate color) or collide in a manner that disrupts the synchrony of each cell assembly, causing the representations to terminate. To avoid this, Raffone and Wolters (2001) proposed that each cell assembly mutually inhibits the other cell assemblies.

This combination of excitatory recurrent activity and mutual inhibition will tend to cause the two cell assemblies to alternate in time (see Figure 5A). More than two cell assemblies can be active in this manner, but as the number of active cell assemblies increases, the amount of time between each firing of a given cell assembly will increase. The activation level of a cell assembly fades over time, and if too much time passes between one firing of a cell assembly and the next firing of that same cell assembly, then the activation level will fall too low and the cell assembly will be unable to fire again. This will cause the representation to be suddenly lost (as has been observed psychophysically by Zhang & Luck, 2009). This puts a limit on the number of items the network can maintain. Using realistic biophysical parameters, Raffone and Wolter (2001) found that 3-4 items could be retained by their model. Thus, this conceptualization of the maintenance of information in VWM can explain why storage capacity is limited (for an alternative neural account of capacity limits in VWM, see, e.g., Johnson, Spencer, & Schöner, in press).

It can also explain why the limit is based on the number of objects rather than the number of features. As illustrated in Figure 5B, the features of an object can be linked by means of synchronous firing of the pools of neurons coding those features, just as the neurons that code each individual feature are linked via synchrony. Although linking pools of neurons from different dimensions may require additional neural circuitry, it does not change the fundamental limit on capacity. That is, even

if each cell assembly contains neurons that code different stimulus dimensions, storage capacity is limited by the need to separate the firing of the different cell assemblies over time and by the fact that a given cell assembly will be lost if too much time elapses since the last firing of that assembly.

This mechanism of binding can also explain why most studies have found that objects containing two values along the same feature dimension (e.g., two different colors) cannot be stored together in the same “slot” in VWM (see, e.g., Xu, 2002). If, for example, red-selective and yellow-selective neurons fired simultaneously in the same cell assembly, there would be nothing to indicate whether this reflected separate red and yellow colors or a single orange color. Consequently, each part of a multi-part object will typically need to be represented by a separate cell assembly, which can explain why storage capacity appears to be smaller for complex objects than for simple objects (see Alvarez & Cavanagh, 2004; Awh, Barton, & Vogel, 2007).

The neural synchrony hypothesis can therefore explain several key aspects of the representation of features and conjunctions in VWM. First, because synchrony potentiates neural transmission, it can allow recurrent connections to become sufficiently potent to yield sustained firing in the absence of a stimulus. Second, the need for the cell assemblies representing different objects to fire sequentially, but without too much time between firings of any given cell assembly, can explain the limited storage capacity of VWM. Third, the use of synchrony between as well as within dimensions can explain why capacity is limited by the number of objects rather than by the number of features. Finally, this method of binding explains why two different values along the same dimension cannot be stored together in a single VWM representation.

Conclusions

Treisman’s Feature Integration Theory has had an enormous impact on the explosion of research on visual working memory that has occurred over the past 15 years. Even though it was not intended to be a theory of VWM, Feature Integration Theory played a major role in setting the VWM research agenda, in providing a principled way of manipulating complexity, in defining rigorous and powerful experimental methods, and in linking cognitive theory with the underlying neural systems. In addition, the study of Wheeler and Treisman (2002) provided the field with a better experimental method for assessing the binding of features in VWM, and it forced a reconsideration of how features and their bindings are represented in VWM.

In addition to the enormous influence of her research, Anne Treisman has played a key role in the careers of a many researchers in this area, including the authors of this chapter. She has served as a mentor, a font of wisdom, a source of encouragement, and a model for how to make important and lasting contributions

to science. The importance of these informal contributions should not be overshadowed by her more formal impact on the field.

Footnotes

ⁱ This conclusion is limited to single-part objects. Each part of a multipart object may be stored in a separate slot (see Sakai & Inui, 2002; Xu, 2002)

References

- Allen, R. J., Baddeley, A. D., & Hitch, G. J. (2006). Is the binding of visual features in working memory resource-demanding? *J Exp Psychol Gen*, *135*(2), 298-313.
- Allen, R. J., Hitch, G. J., & Baddeley, A. D. (2009). Cross-modal binding and working memory. *Visual Cognition*, *17*(1&2), 83-102.
- Alvarez, G., & Cavanagh, P. (2004). The Capacity of Visual Short-Term Memory Is Set Both by Visual Information Load and by Number of Objects. *Psychol Sci*, *15*(2), 106-111.
- Alvarez, G. A., Horowitz, T. S., Arsenio, H. C., Dimase, J. S., & Wolfe, J. M. (2005). Do multielement visual tracking and visual search draw continuously on the same visual attention resources? *J Exp Psychol Hum Percept Perform*, *31*(4), 643-667.
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychol Sci*, *18*(7), 622-628.
- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, *363*(6427), 345-347.
- Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Res*, *43*(4), 393-404.
- Chong, S. C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Percept Psychophys*, *67*(1), 1-13.
- Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1996). Object and spatial visual working memory activate separate neural systems in human cortex. *Cerebral Cortex*, *6*, 39-49.
- Cowan, N. (2005). *Working memory capacity*. New York, NY: Psychology Press.
- Cowan, N., Elliott, E. M., Scott Saults, J., Morey, C. C., Mattox, S., Hismjatullina, A., et al. (2005). On the capacity of attention: its estimation and its role in working memory and cognitive aptitudes. *Cognit Psychol*, *51*(1), 42-100.
- Deco, G., & Tolls, E. T. (2008). Neural mechanisms of visual memory: A neurocomputational perspective. In S. J. Luck & A. Hollingworth (Eds.), *Visual Memory* (pp. 247-289). New York: Oxford University Press.
- Delvenne, J.-F., & Bruyer, R. (2004). Does visual short-term memory store bound features? *Visual Cognition*, *11*(1), 1-27.
- Fougnie, D., & Marois, R. (2009). Attentive tracking disrupts feature binding in visual working memory. *Visual Cognition*, *17*(1-2), 48-66.
- Gajewski, D. A., & Brockmole, J. R. (2006). Feature bindings endure without attention: evidence from an explicit recall task. *Psychon Bull Rev*, *13*(4), 581-587.
- Gold, J. M., Fuller, R. L., Robinson, B. M., McMahan, R. P., Braun, E. L., & Luck, S. J. (2006). Intact attentional control of working memory encoding in schizophrenia. *J Abnorm Psychol*, *115*(4), 658-673.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632-635.
- Hebb, D. O. (1949). *Organization of Behavior: A Neuropsychological Theory*. New York: Wiley, Inc.
- Irwin, D. E., & Thomas, L. E. (2008). Visual sensory memory. In S. J. Luck & A. Hollingworth (Eds.), *Visual Memory* (pp. 9-41). New York: Oxford University Press.

- Johnson, J. S., Hollingworth, A., & Luck, S. J. (2008). The role of attention in the maintenance of feature bindings in visual short-term memory. *J Exp Psychol Hum Percept Perform*, 34(1), 41-55.
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schoner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychol Sci*, 20(5), 568-577.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (in press). A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain Research*.
- Kyllingsbaek, S., & Bundesen, C. (2007). Parallel processing in a multifeature whole-report paradigm. *J Exp Psychol Hum Percept Perform*, 33(1), 64-82.
- Lee, H., Simpson, G. V., Logothetis, N. K., & Rainer, G. (2005). Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron*, 45, 147-156.
- Lee, J., & Park, S. (2005). Working memory impairments in schizophrenia: a meta-analysis. *J Abnorm Psychol*, 114(4), 599-611.
- Logie, R. H., & Marchetti, C. (1991). Visuo-spatial working memory: Visual, spatial or central executive? *Advances in Psychology*, 80, 105-115.
- Luck, S. J., & Hollingworth, A. (Eds.). (2008). *Visual Memory*. New York: Oxford University Press.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279-281.
- Luck, S. J., & Vogel, E. K. (1998). Response from Luck and Vogel (Response to Commentary by Nelson Cowan). *Trends in Cognitive Sciences*, 2, 78-80.
- Magnussen, S., Greenlee, M. W., & Thomas, J. P. (1996). Parallel processing in visual short-term memory. *J Exp Psychol Hum Percept Perform*, 22(1), 202-212.
- Myczek, K., & Simons, D. J. (2008). Better than average: alternatives to statistical summary representations for rapid judgments of average size. *Percept Psychophys*, 70(5), 772-788.
- Olson, I. R., & Jiang, Y. (2002). Is visual short-term memory object based? Rejection of the "strong-object" hypothesis. *Percept Psychophys*, 64(7), 1055-1067.
- Olson, I. R., Page, K., Moore, K. S., Chatterjee, A., & Verfaellie, M. (2006). Working memory for conjunctions relies on the medial temporal lobe. *J Neurosci*, 26(17), 4596-4601.
- Raffone, A., & Wolters, G. (2001). A cortical mechanism for binding in visual working memory. *J Cogn Neurosci*, 13(6), 766-785.
- Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, 393(6685), 577-579.
- Sakai, K., & Inui, T. (2002). A feature-segmentation model of short-term visual memory. *Perception*, 31(5), 579-590.
- Saults, J. S., & Cowan, N. (2007). A central capacity limit to the simultaneous storage of visual and auditory arrays in working memory. *Journal of Experimental Psychology: General*, 136, 663-684.
- Scarborough, D. L. (1972). Memory for brief visual displays of symbols. *Cognitive Psychology*, 3, 408-429.

- Schutte, A. R., Spencer, J. P., & Schoner, G. (2003). Testing the dynamic field theory: working memory for locations becomes more spatially precise over development. *Child Dev*, 74(5), 1393-1417.
- Simmering, V. R., Spencer, J. P., & Schoner, G. (2006). Reference-related inhibition produces enhanced position discrimination and fast repulsion near axes of symmetry. *Percept Psychophys*, 68(6), 1027-1046.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and future. *Trends Cogn Sci*, 9(1), 16-20.
- Tallon-Baudry, C., Bertrand, O., & Fischer, C. (2001). Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance. *Journal of Neuroscience*, 21, RC177.
- Tallon-Baudry, C., Bertrand, O., Peronnet, F., & Pernier, J. (1998). Induced gamma-band activity during the delay of a visual short-term memory task in humans. *Journal of Neuroscience*, 18, 4244-4254.
- Tallon-Baudry, C., Kreiter, A., & Bertrand, O. (1999). Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. *Visual Neuroscience*, 16, 449-459.
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 92-114.
- Wheeler, M. E., & Treisman, A. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology-General*, 131(1), 48-64.
- Woodman, G. F., & Luck, S. J. (1999). Electrophysiological measurement of rapid shifts of attention during visual search. *Nature*, 400(6747), 867-869.
- Woodman, G. F., & Luck, S. J. (2003). Serial deployment of attention during visual search. *J Exp Psychol Hum Percept Perform*, 29(1), 121-138.
- Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychon Bull Rev*, 15(1), 223-229.
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychol Sci*, 12(3), 219-224.
- Xu, Y. (2002). Limitations of object-based feature encoding in visual short-term memory. *Journal of Experimental Psychology: Human Perception & Performance*, 28(2), 458-468.
- Yeh, Y.-Y., Yang, C.-T., & Chiu, Y.-C. (2005). Binding or prioritization: The role of selective attention in visual short-term memory. *Visual Cognition*, 12(5), 759-799.
- Zhang, W., & Luck, S. J. (2008). Discrete Fixed-Resolution Representations in Visual Working Memory. *Nature*, 453(7192), 233-235.
- Zhang, W., & Luck, S. J. (2009). Sudden Death and Gradual Decay in Visual Working Memory. *Psychol Sci*.

Figure Captions

Figure 1. Stimuli and results from the study of Luck and Vogel (1997).

Figure 2. Examples of the conditions tested by Wheeler and Treisman (2002). In the single-probe test conditions, a single test item was presented at fixation. On change trials, this item either contained a new feature value that was not present in the sample array (feature change) or contained a combination of features that were present in different objects in the sample array (binding change). In the whole-array test conditions, the entire array was presented at test. On change trials, either two items changed to new feature values that were not present in the sample array (feature change) or two items swapped features (binding change). The positions of the items were scrambled between sample and test to avoid the use of location-based binding.

Figure 3. Stimuli (A) and results (B) from the study of Johnson et al. (2008). In this example, the orientations of two objects were swapped in the test array. Object locations were not scrambled between sample and test in this experiment, and so bindings with location were also disrupted on change trials. Change detection accuracy is shown averaged across color changes and orientation changes.

Figure 4. Stimuli (A) and results (B) from an unpublished study by Woodman, Vogel, and Luck. Change detection accuracy is averaged across color and orientation.

Figure 5. Neural synchrony in the representation of individual features (A) and conjunctions of features (B). Each vertical line represents the firing of an individual neuron (A) or the firing of neurons within a population of similarly selective neurons (B).

Figure 1

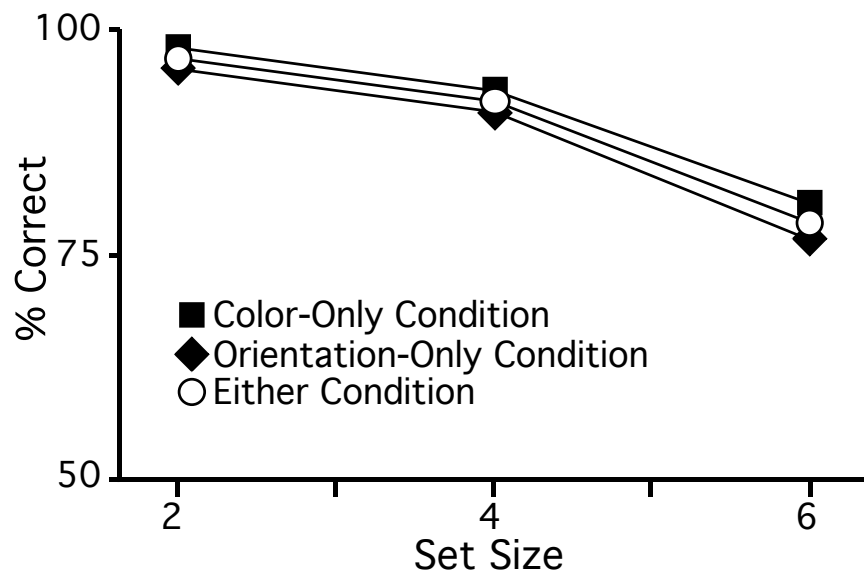

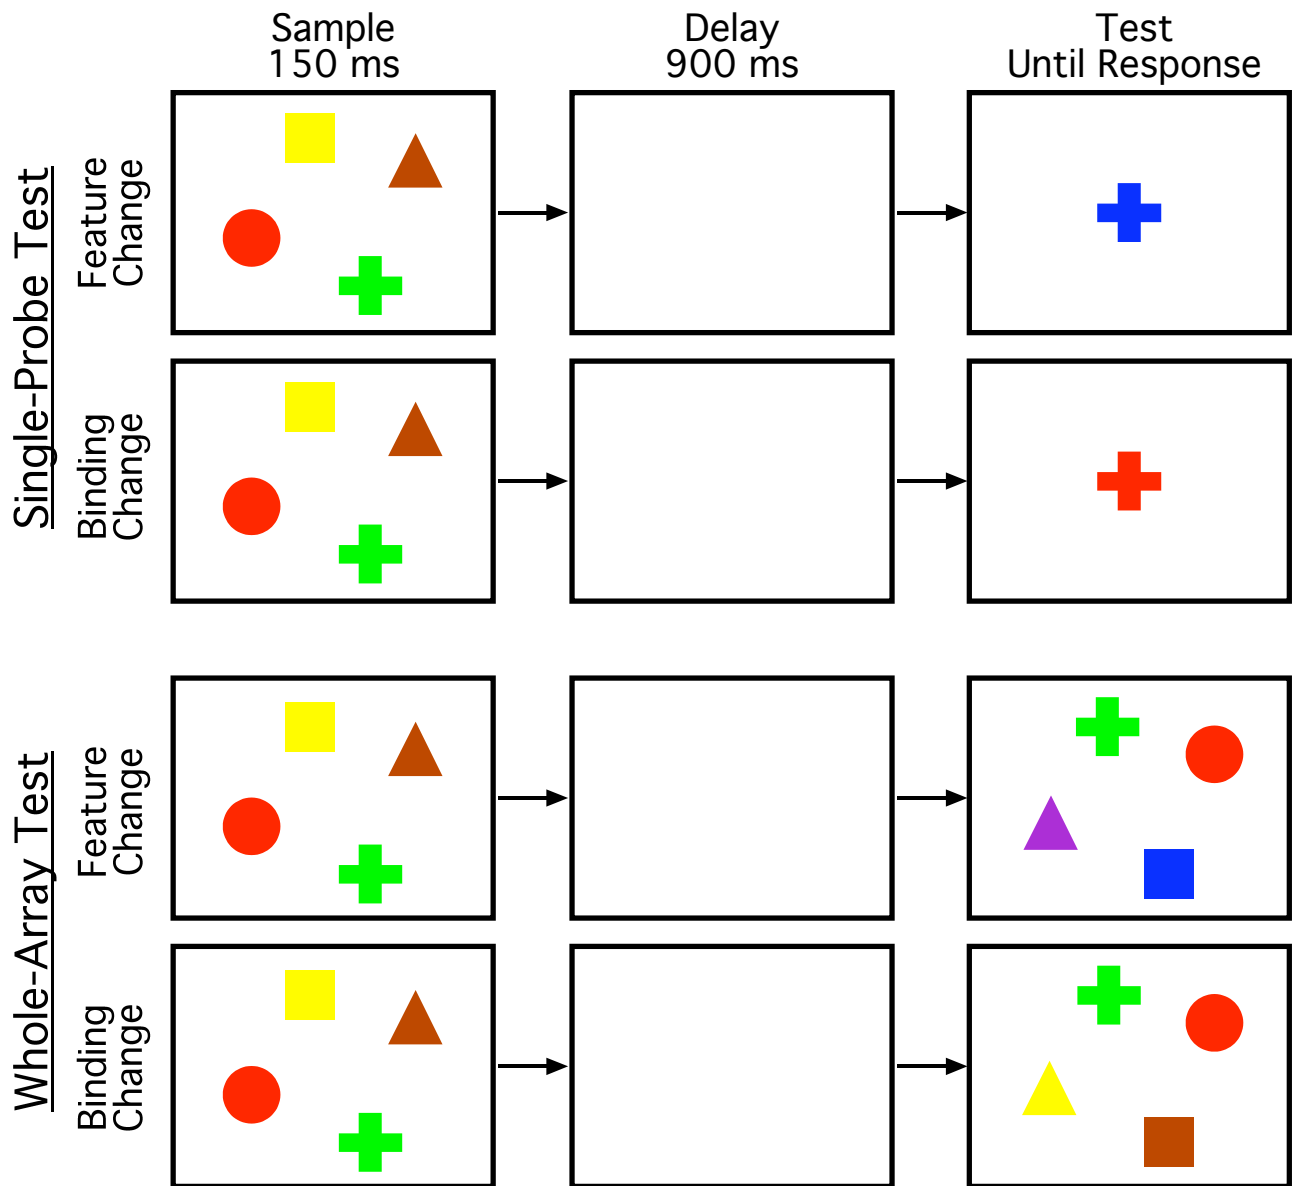
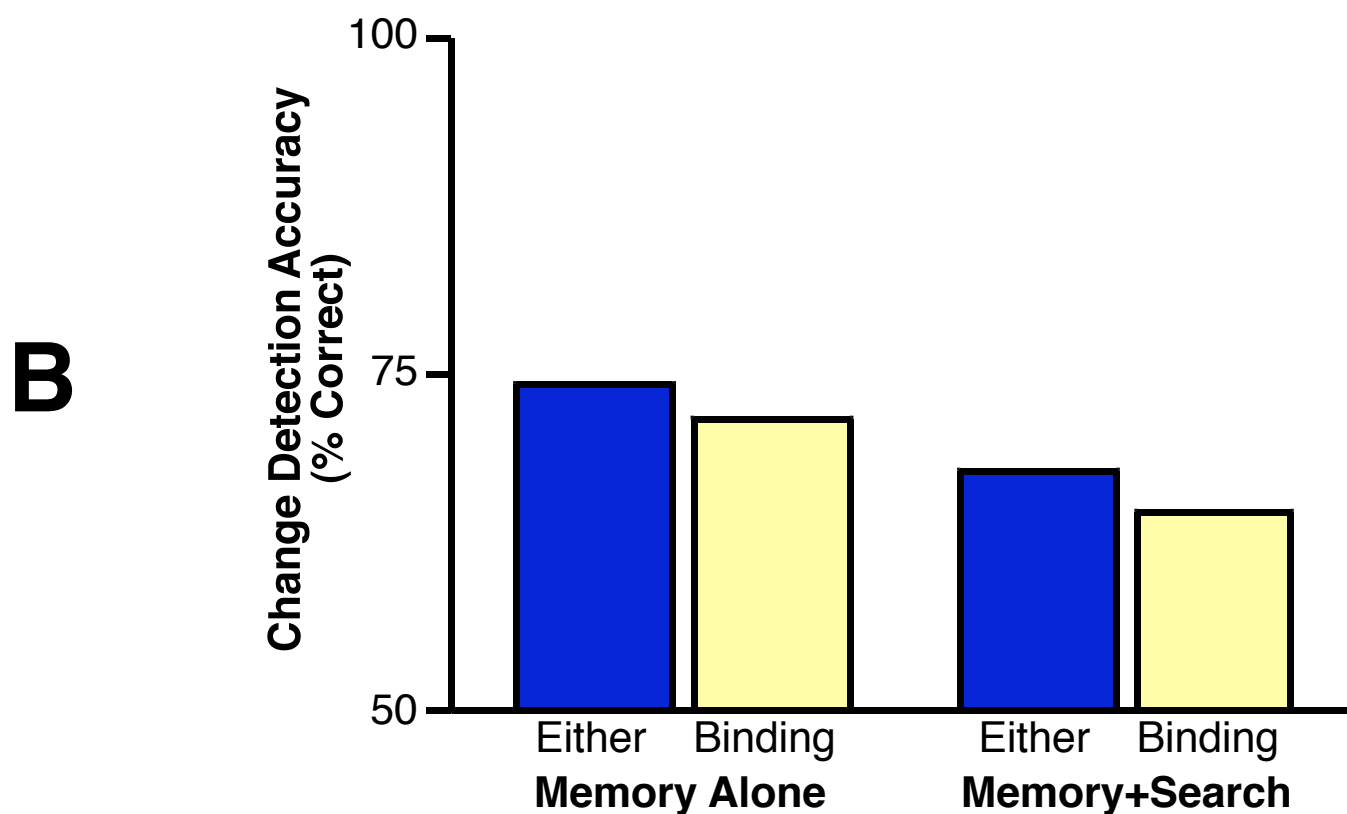




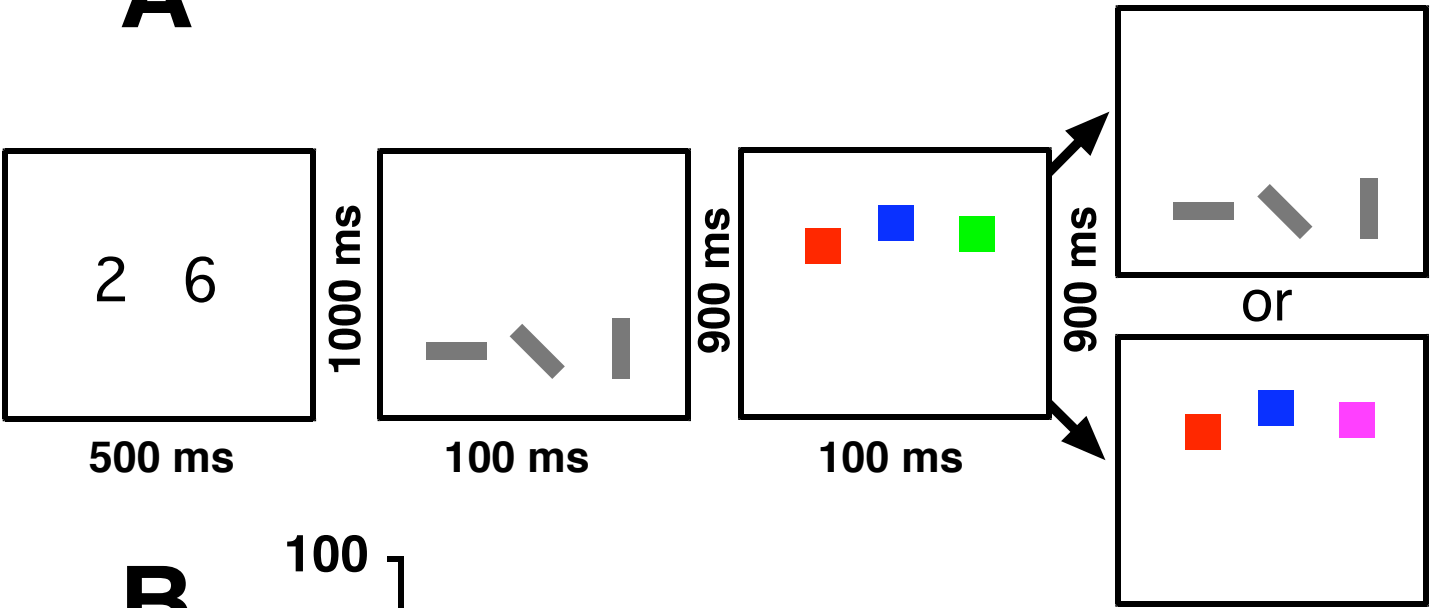
Figure 2







A



B

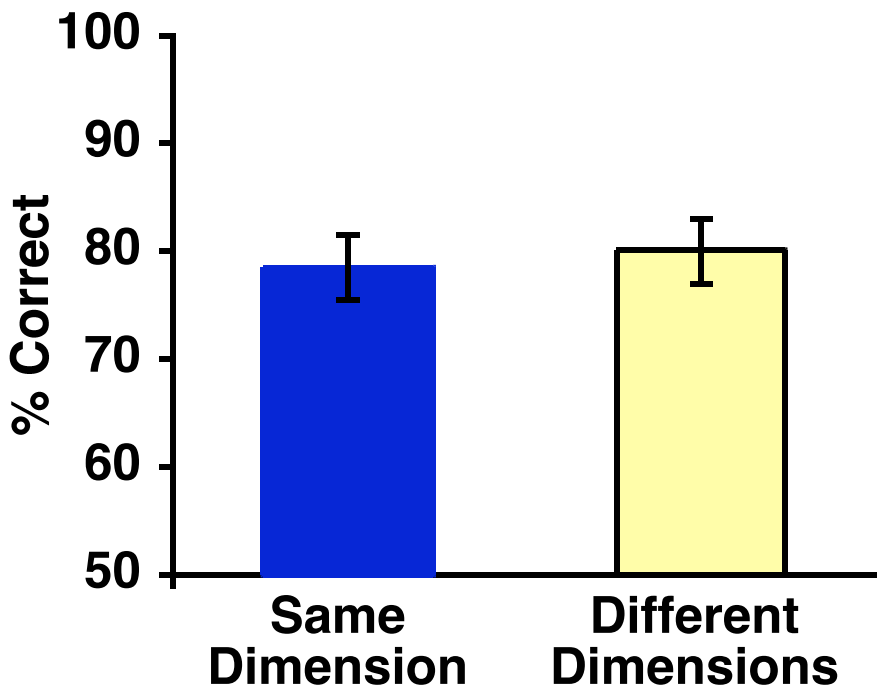




Figure 5


A




Red-Selective
Neurons



Blue-Selective
Neurons




B




Red-Selective
Population




Blue-Selective
Population



Vertical-Selective
Population



Horizontal-Selective
Population



Time →