

Object tokens, binding and visual memory

Anne Treisman
Princeton University

Address correspondence to:

Anne Treisman

Department of Psychology

Green Hall

Princeton University

Princeton, NJ 08544

e-mail: Treisman@princeton.edu

At any moment of time the scene around us is filled with multi-featured objects, which we see from particular angles, distances, and illuminations, and which may themselves move and change. We must recognize their identities in order to retrieve semantic information relevant to our behavior. But we also need to represent their current state in order to interact with them, and store an episodic memory of the particular events in which they play a role. Kahneman, Gibbs, and I developed the object file metaphor to help capture this type-token distinction in the perceptual domain and collected data showing object-specific priming to support the idea (Kahneman, Treisman, & Gibbs, 1992). I also applied it to feature binding (Treisman, 1992), to visual working memory (Wheeler and Treisman 2002), to negative priming (Treisman and DeSchepper 1996), and to long term learning in visual search (Treisman 1992).

Object files in perception

I will start by outlining the perceptual framework. We drew a distinction between the activation of stored types, which mediates identification, and the creation of temporary episodic tokens, which mediates seeing (Kahneman and Treisman 1984). We called the episodic representations “object files” and suggested that they hold the current description, location, orientation and distance of the objects around us. Object files are temporary structures, addressed through their spatial locations, in which information accrues over time, just as information about a particular crime might be collected in a specific police file. If the object moves or changes, the object file is updated, creating a perceptually continuous entity

Object files are assumed to play a number of different roles in perception and in memory.

1. They are the perceptual units into which a scene is parsed, becoming the potential objects of attention. It is much easier to attend to a whole object than to single out one of its properties. It is also much easier to divide attention between two properties of the same object than between the same two properties when they are seen as belonging to separate objects (Treisman, Kahneman, & Burkell, 1983; Duncan 1984).

2. Object files allow us to represent novel objects for which we have no prior representations, no object type to reactivate. If an object file matches a known type in semantic memory, its identity is accessed. But object files can also represent unknown objects.

3. They allow us to represent multiple identical objects, - to differentiate a flock of sheep from a single very large or sheepish sheep. This might be difficult if all we could vary was the degree of activation of the type node for sheep.

4. Object files serve as vehicles to bind features. Feature integration theory (Treisman and Gelade, 1980) proposed that the correct selection of features is achieved through spatial attention. Features that share the same location usually belong to the same object, so the correct binding can be determined by serial attention to different locations, temporarily excluding features from unattended locations. Evidence for this view includes the facts that search for conjunctions of features is usually serial, and that when focused attention is diverted or overloaded, illusory conjunctions are often formed (Treisman and Schmidt 1982). For example a red S and a blue O might generate an illusory blue S or red O.

5. Object files also bind successive states of an object over time, updating their representations as the objects move and change. The visual system assumes the continuity of objects in the real world. Kahneman, Treisman, & Gibbs (1992) suggested that each object that appears in the scene "looks for" its immediate past history. When a spatiotemporally compatible match is found, the contents of the earlier object file are compared to the currently present features. If there is a mismatch, we update the file, or, if the change is too great, we open a new one and see a new object. Evidence came from the object-specific priming that we found when participants were asked to name letters presented successively, either in the same or in a different moving frame.

6. Object files provide an interface between early vision, top-down knowledge, and conscious experience. The hypothesis is that without an object file, the object is not consciously seen. However, the reverse may not be the case: not all object files become conscious (see the section below on negative priming).

The main theme of this chapter concerns what happens to object files once the objects are no longer visible. How do we store bindings in memory and recognize them over short intervals of time. Wolfe, Oliva, Butcher & Arsenio (2002) recently gave a short and simple answer: "We don't". The conclusion stemmed mostly from the marked "change blindness" shown when we try to detect a difference between two alternating versions of the same scene (Rensink, O'Regan, & Clark, 1997; Simons and Levin 1997). Large changes like the disappearance of an engine on an airplane may take many seconds to detect. Instead of the teeming, colorful world of detail we think we have in memory, it seems instead to be a sparse, abstracted, schematic world, inhabited at most by 3 or 4 more detailed objects. We retain the illusion of a rich and

detailed memory because in everyday life we are seldom tested. As O'Regan (1992) points out, the real world normally provides an external memory to supply details as we need them.

Within the framework of feature integration theory, there are several forms in which information is encoded and a number of levels at which it could be stored: (1) Features and locations are encoded in separate parallel maps before attention binds them to form integrated objects. They may also leave temporary unbound traces in those separate maps. (2) The object tokens that are formed in object files through focused attention can be transferred to visual working memory for early explicit recall and recognition. (3) They can also be stored in episodic long-term memory as consciously retrievable traces of particular incidents. (4) The object files may remain as implicit traces of perceptual experiences that prime or interfere with later re-perceptions. (5) Finally they can contribute to the formation of learned associations in semantic memory between features that have been repeatedly perceptually bound. I will outline some research on each of these kinds of memory and some tentative suggestions about how they may be related over short and longer intervals, in both explicit and implicit memory.

Visual working memory (VWM): features and objects

Luck and Vogel (1997) presented displays of colored squares and tested recognition 900 ms later. The second display either matched the first exactly or had one feature changed. The maximum number of objects that participants could accurately remember was only 3 to 4. However, the number was the same whether the objects had one relevant feature (color), or two, or even four (color, length, orientation

and a break in the line). Luck & Vogel suggested that what is encoded and stored in visual working memory is a small set of bound objects, to which any number of features can be added for free if they characterize one of the original objects. There is, however, an alternative explanation to be considered: Performance in Luck & Vogel's experiment might also reflect a set of separate feature stores, holding a set of four colors in one, a set of four shapes in another, and so on. Luck and Vogel claimed to have ruled out this alternative by showing that objects that were each characterized by two features within a single dimension (color,) rather than, for example, a color and a shape, also functioned as single memory units. In this case, memory could not recruit new storage space for the additional features, since only one dimension (color) was varied.

Figure 1 about here

Wheeler and Treisman (2002) were unable to replicate this result. We found that three bicolored objects were remembered no better than six unicolored objects. There was no benefit at all from combining colors within fewer objects (see Figure 1). This was true whatever patterns we used to integrate the colors, including checkerboards, center squares with surrounds and so on. This result is consistent with a similar finding testing parts of shapes. Lee & Chun (2001) compared the capacity of visual memory in change detection for what they called locations and objects. They used either superimposed or separated presentation of three pairs of boxes and lines (cf. Duncan, 1984) and found no difference in memory. If we think of the component shapes as equivalent to the colors in our experiment, the inference is the same – that integrating two shapes into a single compound unit produces no increase at all in the capacity of visual working memory. These failures revive the possibility of parallel independent feature stores, each with its own limited capacity, as an explanation of memory for multidimensional stimuli.

A similar conclusion was reached by Magnussen (2000) from very different data – the psychophysics of delayed discrimination with sine wave gratings varying in spatial frequency, orientation, color, or motion. He describes a “low-level perceptual memory mechanism located early in the visual processing stream..... composed of a series of parallel special-purpose perceptual mechanisms with independent but limited processing resources. Each mechanism is devoted to the analysis of a single dimension and is coupled to a memory store.” Magnussen suggests that even higher-level memory may reflect the same feature traces. ‘It remains a distinct possibility that our visual memories, like on-line perceptions, are based on the resurrection of a pattern of precisely-stored single attributes.’ The data I describe in the remainder of this chapter suggests instead that higher-level or longer-term recall depends on a separate memory for object tokens, which incorporates the particular binding in which the features are presented. Working memory can draw both on these object tokens and on the traces of separate features lingering in separate sensory feature maps.

My students and I have run a number of further experiments using the change detection paradigm to explore the parameters of explicit voluntary memory over brief intervals of time. We used two different ways of testing the idea that visual working memory holds bound objects. In the first, we made it impossible to rely on feature stores by specifically asking for recognition of the binding. In the second, we tested how far participants could *avoid* remembering the binding when they were asked to ignore it and to remember only the features.

Figure 2 about here

In the first two experiments, Wheeler and Treisman (2002) compared memory for features (color and location in one experiment, color and shape in the other) with

memory for their binding. In the first experiment, we presented participants with a 150 ms. array of 3 or 6 colored squares, followed by a delay of 900 ms, and then a second array which was identical on 50% of trials or could contain either a new feature (color, or location) or unchanged features but a change in their binding – the colors switched between two locations (see Figure 3). The conditions were blocked so that participants knew whether they would be tested on color, or location, or either of the two features, or the bindings. In another version of the experiment, different shapes were used and changes of locations were replaced by changes of shapes. Figures 3(a) and 4(a) show the results. In both experiments, the changes of binding were less likely to be detected than new features. Thus the features seem to be stored at least partly independently of the bound object tokens.

Figures 3 and 4 about here

However, the story changed when we used a single probe rather than a whole display. In this condition, the changes of binding were detected as accurately as changes in the least good single feature (Figures 3(b) and 4(b)). This was true both for color-location and for color-shape binding. One possibility is that there was simply a greater decision load with the whole display. There were several items to be checked in a whole display probe and only one in the single probe condition, giving more chance for error. This did not seem to be the reason. When we presented the whole display at test, then cued one relevant item to be checked, performance was equally bad. The binding must be stored, since it was available with the single probe. But it may be especially vulnerable to the visual interference created by the whole display at test.

We tried to distinguish two forms this interference might take. (1) Focused attention may be needed, not only to bind the features during the initial encoding of the

display, but also to maintain the bindings in VWM. The new visual stimuli in the whole display probe may draw attention away from the stored items in working memory and destroy their binding. The whole display competes for attention more than the single item. (2) Feature bindings may be especially vulnerable to new visual stimuli overwriting previously registered visual information. This would be consistent with the object file model: When Kahneman, Gibbs and I were exploring the object-specific priming effects that prompted the theory, we asked how far back the priming could originate – just one change back or more? We tried showing two prime letters successively in the same object frame and checked whether there was any object-specific priming from the letter before last as well as from the last one. The answer was no. Each new state of an object apparently erases the one that precedes it in the same object file. In the present paradigm, the whole display probes appeared in the same locations as the original memory display. The visual system may simply have updated the preceding traces to represent the current state of the world. In the color-shape experiment, the single probe was presented in the center of the display rather than in one of the original locations. The color-location probe was necessarily presented in one of the original locations, but since it was the only item present, it could draw focused attention, which might allow participants to detect the change when it occurred.

Wheeler and I leaned towards the first account, that it takes attention to maintain bindings in VWM. Now I am not so sure, as I will explain below. Whichever is correct, the fact that the features survive even when the binding is lost suggests that they may be stored at the feature level as well as forming part of the stored object files.

So far I have described experiments looking at explicit memory for binding. In the next experiment, Zhang and I turned the question around to ask “when items are attended, is binding automatic, so that object files are the only memory traces available? Or can we choose to store features separately and unbound when this is what the task requires?” If WM holds only object files, and the bindings are changed in the probe, this should disrupt memory for the features as well.

Figure 5 about here

We presented three colored shapes for 150 ms, followed after 900 ms by a probe display (Figure 5). Participants looked for a change in one feature – shape or color. The bindings could be the same as in the original display or they could be switched around. Participants were told that they should respond “New” only if a new shape or color was present, ignoring possible changes of binding. Feature recognition was significantly worse when the binding changed – 62% with changed binding versus 69% when the binding matched. It seems that when participants attend to the display, they automatically store the binding and hold it in working memory, even when it is irrelevant and may interfere with recognition. We had claimed that attention is necessary for binding. We now find that it may also be sufficient.

Figure 6 about here

Can we say any more about how the features are stored? Are they bound to each other, independently of location, or is each of the features bound to its location and only indirectly to the other features in the same location? In half the blocks we tested recognition in the same locations and in half the blocks we presented the recognition probes in new locations, previously empty (Figure 6). Surprisingly we found that changing the locations had no overall effect on feature recognition. It is not that the locations were forgotten, because we found large interactions of location with binding (Figure 7). People were better with the same location when the binding was

the same and worse when it was different, but there was no overall advantage to using the original locations.

Figure 7 about here

A related experiment was reported by Prabhakaran, Narayanan, Zhao, & Gabrieli (2000). Memory for letters was tested with probes that were presented either in the same location or in the location previously occupied by a different letter. Unlike our results, recognition in their experiment was more accurate (88.5% compared to 85.7%) and faster (1139ms compared to 1261 ms) when the location was the same, suggesting that “verbal and spatial information in the bound displays were maintained in an integrated fashion in working memory.” However the advantage of keeping the same locations was not large. It seems that features are stored in object files, but that memory can also be supplemented by information from separate feature stores when binding is not directly tested. Prabhakaran et al. also looked at the brain areas activated selectively in the bound condition and found that right prefrontal areas were involved. They suggested that there might be a specialized memory in this brain area for integrated information, perhaps analogous to Baddeley’s Episodic Buffer (Baddeley, 2001).

While feature information can be retrieved either from integrated representations or from independent feature stores, feature integration theory suggests that locations are critical for the encoding of conjunctions. Are they equally critical for their storage and retrieval, or is it the case that, once the features are bound, the location information can be dropped without losing the conjunctions? Zhang and I used the same stimuli to test VWM for binding. In this experiment, the same features always appeared in the initial display and in the recognition probe (as shown in Figure 6). The task was to detect a change of binding rather than the presence of a new

feature. On half the trials the binding between two of the objects was changed and on half it remained the same. Participants now did considerably worse when the objects were presented in new locations, an average drop of 9% ($F(1,23)=57.5$, $p<.001$). Unlike the previous tests of memory for features, memory for the bindings was sharply impaired when the locations were new, as though the binding is maintained, as well as initially established, through attention to the locations. Again this fits well with the object file model. A crucial claim was that object files are addressed by their spatio-temporal coordinates. When these allow a plausible continuity in the object (as when the locations are the same, or when the object is seen in real or apparent motion between the original and new location), the contents of the object file are compared automatically and updated if necessary to match the features currently present. When continuity is not established, as when the objects are presented in arbitrary new locations, new object files are likely to be opened and the comparison to the old ones is less efficient.

We ran some further variants of this experiment, each with 12 participants. In one, the delay was reduced to 100 ms – in the range of iconic memory - to see if the role of location in memory for binding is different immediately after presentation. In fact the pattern was similar but at 100 ms, the effect of location was considerably stronger, averaging 18% instead of 9% (see Figure 8). The interaction of delay and location was significant ($F(1, 23) = 13.68$, $p<.01$) with errors on 40% of trials when the locations were changed compared to 14% in the same location. The result suggests that the object files are initially closely dependent on location and that the direct link between the features is gradually consolidated over the first second of delay, reducing without eliminating the dependence on location.

Figure 8 about here

Articulatory suppression, which is presumed to reduce or eliminate verbal coding, had little effect. Performance was slightly worse after a 900 ms delay of repeating “coca-cola” relative to an unfilled delay in which naming was possible, but the difference was not significant ($F(1, 23) = 2.37, p = .137$; see Figure 9(a) and (b)), suggesting that verbal coding is little used, at least at these brief delays. What about visual interference? Presenting a visual word in the center of the display during the interval between display and test had no effect at all (Figure 9(c)), although reading is automatic and participants almost certainly read the word. The word may be sufficiently different from the colored shapes not to be taken as an update on the same object files.

Figure 9 about here

What conclusions on VWM can we draw so far? They seem to fit well with the ideas that came out of the earlier work and also add to them. We can set up at most 2 or 3 object files which contain bound features and which can survive for at least 900 ms, and probably longer. Adding features on more different dimensions within the same number of object files comes free, or at low cost, if any. The binding is vulnerable to overwriting when a new display of objects is presented, although the features may survive intact. Recognizing features is harder when their binding is changed, suggesting that when we attend to objects, we bind their features automatically, whether the binding is relevant or not. Moving the objects to new locations leaves feature recognition unimpaired but damages recognition of the bindings quite severely, especially immediately after presentation. Features are stored in two different forms: as entries in an object file, and also as separate traces, independent of their locations, as predicted by the hypothesis that features and locations are registered initially in separate maps. The object files integrate information across the two sets of representations, binding the features to their locations. Initially the locations are strongly integrated with the object file, but this

dependence diminishes as the memory delay increases and direct bindings between features within object files are consolidated.

Relation between VWM and long-term memory (LTM)

Another question concerns the relation between WM and long-term memory. There have been a number of hypotheses (see Figure 10). In an early account by Atkinson and Shiffrin (1971), short-term memory was a gateway or way-station to LTM. The representations were presumably transformed into a different format to gain admission, since they were vulnerable to different forms of interference. A second account, proposed by Anderson (1983) and also adopted by Cowan (1995, 1999), makes WM a temporarily activated subset of LTM, presumably implying that it shares both the format and the information encoded in that subset. In a third account, WM is seen as a separate system with its own properties, accessed in parallel with LTM rather than necessarily preceding it, although information can be transferred in either direction between the two. Baddeley's multi-store model of WM is the best-known elaboration of this view (Baddeley, 1986). It is also supported by neuropsychological evidence of dissociations resulting from brain lesions in different brain systems (e.g. Della Sala & Logie, 1993). There could be an evolutionary advantage for a separate WM, since the properties it needs are quite different from those needed in LTM. They include instant or rapid access, literal recall, quick erasure, and the ability to carry out various online operations or transformations of the material. Baddeley (2000) has more recently added a bridging store that he calls the episodic buffer. It has many of the characteristics of long term memory, including semantic and syntactic chunking and multimodal binding, but it retains other characteristics of WM, such as limited capacity, only temporary availability, and immunity to the forms of amnesia that devastate long term memory.

A number of recent papers (including Crowder 1993, Cowan, 1999, Ruchkin & Grafman, 2003, and Ranganath, Johnson, & D'Esposito, 2003) have argued against any separation between working memory and long-term memory on the grounds that similar brain areas are activated in imaging studies of both. However differences in brain activation between short and long term memory tasks can also be found; (see for example Braver, Barch, Kelley, Buckner, Cohen, Miezin, Snyder, Ollinger, Akbudak, Conturo, & Petersen, 2001; Cabeza, Dolcos, Graham, & Nyberg, 2002). The assumption that items can enter either or both WM and LTM in parallel, directly from the perceptual encoding, makes it difficult to test whether they are distinct systems or a single unitary store activated in different ways. The more similar the information in WM is to that stored in LTM, the more plausible the single store view becomes.

Conditions that increase the likelihood of testing WM rather than early stages of LTM (assuming that they are separate systems) include repeatedly reusing the same items in different contexts, so that the information from each trial must be deleted and replaced to minimize inter-trial interference. On the other hand, using trial-unique items increases the probability of automatic long-term storage and retrieval, since novelty gives an advantage in retention, thus blurring the distinction between WM and LTM when recall is tested at short intervals.

Miriam Berkowitz and I ran a preliminary experiment to test possible interactions between visual working memory and long-term memory. If WM is an activated subset of LTM, one might expect any relevant learning to show up in both. We explored the possibility that WM would benefit from long-term learning of contingencies between properties, so that consistent property bindings would provide facilitation also when tested in WM paradigms. We presented a circular array of 5

colored shapes in a change detection paradigm similar to those described earlier (Figure 11). This time the task was to attend to one dimension of the stimuli (either their color or their shape) and say whether the probe matched on the attended dimension. Participants always knew which dimension was relevant.

Figure 11 about here

The new factor in this experiment was the introduction of long-term consistencies between the two dimensions – color and shape. Some combinations of color and shape were much more frequent than others. For example, for one subject in 80% of the trials red might be paired with triangle, blue with circle, green with square and so on. In a surprise test after the experiment, we asked whether these contingencies had been learned, measuring explicit recall of the high frequency pairings. That was the long term learning component, recording the trial-to-trial accumulation of frequencies.

We tested a WM expression of this long-term binding by looking for dependencies between color and shape in the within-trial responses. If the predictable pairings were represented in WM as well as LTM, performance should be better for a group exposed to the frequent color-shape conjunctions than it was for a control group in which there were no contingencies – shape and color varied orthogonally.

We also looked at the effects of short-term, within-trial color-shape binding. The irrelevant dimension could either agree with the relevant one in being old or new or it could disagree - (old when the relevant one was new or new when it was old). Thus each test item could either match or mismatch, in terms of the response the two features

would evoke on any given trial. If the shape and color are automatically bound in WM, a mismatch should interfere with the response and a match should facilitate it.

Thus we used three different measures of memory for color-shape bindings: Recall of the high frequency color-shape pairings in the surprise test at the end of the experiment. (LTM); better WM performance with associated pairs than with random pairings (LTM effect on WM); and within-trial binding in WM, shown in the effects of consistency in the responses evoked by the relevant and irrelevant dimensions.

Figure 12 about here

Finally, for each of these three measures of associative memory, we also asked if it matters whether the colors and shapes are presented in an integrated form or spatially separated, either as parts within an object or as different objects (Figure 12). Xu (2002) found a large benefit in change detection in either of two different features (color and orientation) when the features characterized different parts of the same object (mushroom stem and cap) relative to separate stems and caps. Walker & Hinkley (2003) in contrast found that recognition memory for colors cued by their associated shapes was at chance when naming was prevented and the colors were presented as a non-integral background to the white shapes to which they were associated. On the other hand, we found no significant effect overall of integrality (68.6% Integral and 68.2% Separable). Surprisingly, accuracy was closely matched for recognition of a grey shape with a colored central dot, a colored shape with a grey central dot, a separate colored shape adjacent to a grey dot, and a separate grey shape adjacent to a colored dot. The difference from Walker & Hinkley's results may be due to the fact that both our integral and our separable stimuli were presented as figures, whereas their separable colors were presented as backgrounds to the figural shapes and may therefore have been subject to some attentional suppression. Xu's results may differ from ours because memory was tested explicitly for both colors and shapes, whereas our participants

responded explicitly to only one of the two features. Implicit processing, shown in indirect measures of consistency, may be less constrained by perceptual organization than explicit processing. Another difference is that the components of our separate pairs were still spatially grouped, whereas those in Xu's experiment were independently located. Since we found no effects of integrality, the remaining discussion will refer to results pooled across the integrality-separability factor.

To see if there was any long term learning of the color-shape contingencies, we looked at the post-experiment questionnaire. Participants did far better than chance when asked to report the pairings at the end of the experiment. In the Shape-relevant condition, participants had a mean accuracy of 65%, and in the Color-relevant condition they had a mean accuracy of 36%, where the chance probability of getting a correct pairing was about 15%. So participants did learn at least some of the consistencies across trials. However, we found absolutely no benefit from these long-term contingencies in WM recognition. Compared to control groups for whom the colors and shapes were orthogonally combined, there was a small and non-significant benefit of correlation in the experiment with integral or separable parts and an almost significant cost of correlation in the experiment with integral or separable objects ($p=0.059$). We have no explanation for the apparent reversal, but we can conclude that there is no evidence for any consistent use of long-term contingencies in the working memory task. There was also no within-trial difference between the correlated items and the one randomly paired item within the Correlated condition. Finally, there was no correlation between long-term learning and WM performance. The Pearson correlation between the number of correct pairings reported by each participant and his or her overall mean accuracy in the WM task was $-.19$. Thus, although participants did learn some of the

pairings between the relevant and irrelevant dimensions in the correlated conditions, that "knowledge" did nothing to improve their recognition performance in WM.

What about binding *within* WM? Perhaps we found no benefit from the correlated colors and shapes simply because the task did not require any binding. Recognition was tested only for one of the two features (shape or color). Our earlier experiment suggests that binding is automatic when the objects are attended, and we have evidence here too that the colors and shapes were bound, even though recall of the bindings was not required. We compared performance when the irrelevant feature evoked the same response as the relevant feature or the opposite response (match versus mismatch to the initial display). We did find an effect of consistency, suggesting again that binding occurs automatically in WM. However it was present only when shape was relevant, giving a significant interaction between Match/Mismatch and Color versus Shape. The findings suggest an asymmetry in the processing of color and shape in WM: Shapes retrieve colors more readily than the reverse. In sum, we seem to have some binding in WM, some binding in LTM, but no interaction between the two sets of associations. The results fit better with separate, differently specialized stores than with the activated subset view. Of course it is possible that both forms of short-term storage are used – persisting passive activation in LTM and voluntary active rehearsal in WM. The question would then become which form is responsible for performance in particular specific conditions. In the next section we discuss a passive form of memory revealed through priming that can affect immediate as well as long-delayed responses.

Implicit memory for bound object tokens

One of the surprising findings in this research has been how little we store of what we have just seen. Three or four bound objects is not much. The change-

blindness phenomenon is another dramatic demonstration of this limit (Rensink, O'Regan, & Clark, 1997; Simons and Levin 1997). Does the brain really not have the capacity to lay down more information about the world around us?

Perhaps the limit is more on conscious access than on storage. Some implicit measures of memory offer quite dramatic evidence of long-term storage of a large number of novel objects to which participants had not even paid attention. Tipper (1985) found that responses are slower to a currently attended object when it has previously been ignored – a phenomenon he called 'negative priming'. He used a small set of familiar objects, so his results could perhaps be explained with the notion of object types – the representation of a trumpet (for example) in the mental dictionary of objects could be made temporarily less accessible, so that it is harder to respond with its name. DeSchepper and I ruled out this possibility by using novel nonsense shapes that participants had never seen before, similar to those used earlier by Rock and Gutman (1981). Participants had to decide whether the green shape in overlapped red and green pairs matched a white shape to the right of it (Figure 13). We found clear negative priming from the previously unattended red shapes after a single exposure to a novel shape (DeSchepper and Treisman 1996).

Figure 13 about here

Our results cannot be explained in terms of activation or inhibition of stored representations of familiar types. We used 270 nonsense shapes, each seen only once before being tested. In our original experiment, the representation of the previously ignored red shape must have been held in memory at least for 1 or 2 seconds to affect the next trial. We tested whether it lasted longer than that and found that the negative priming effect was undiminished after 3, 10, 100 and 200 intervening trials. Growing more and more surprised, we asked new groups of participants to

come back after a day, a week, and a month and still saw effects of the initial exposure. There were substantial individual differences, so we divided the participants according to whether they showed negative priming after one trial, then tested them on a *different* set of shapes after the various delays shown in Figure 14. Both groups showed significant effects, but they took the form of inhibition for the group who showed negative priming and facilitation for the others. What seems to happen with the passage of time is that the inhibition decreases and the positive facilitation increases, as if the response tags saying “IGNORE ME” and creating the inhibitory effect are lost earlier than the traces of the shapes. These survive independently of the response tags and prime perception of the shapes when they are presented again.

Figure 14 about here

Are these memory traces available to the participants when we ask for explicit recognition of the unattended shapes? It seems not. Participants were at chance both in immediate and in delayed tests in deciding which shapes they had seen but ignored and which had not been presented. Apparently attention is needed to make memories explicitly available. This may help to account for the apparent discrepancy between the negative priming results, implying persistent detailed traces of a large set of novel shapes after a single unattended exposure, and the very limited storage suggested by the change blindness results. The memory traces that mediate priming are usually thought to be changes in the perceptual system that initially registered them – a kind of “greasing of pathways”. This may be plausible in some cases of positive priming. But in our negative priming experiments the implicit traces seem to be individuated tokens. They involve some binding of color to shape and they must be labeled as relevant or irrelevant to the action, to be attended or to be ignored. Whether these are the same tokens that, if attended, would become available to conscious voluntary retrieval, or different tokens with different properties is not yet clear. However, the fact that the

traces in negative priming last a long time and that they are not simply increases in the fluency of processing, as could be the case in positive or repetition priming, does suggest distinct object representations which can be in some way tagged as relevant or irrelevant.

Search, priming and long term contingencies

Another form of conjunction priming can be seen across trials in visual search tasks. Treisman, Vieira & Hayes (1992) introduced long-term contingencies here too. We showed that conjunction targets are detected faster when they frequently occur in a particular location than when they can appear anywhere in the display (see Figure 15). For example if the pink Q target usually appears at 3 o'clock in a circular array whereas the green R target can appear anywhere, the pink Q in its habitual location is detected up to 300 ms faster – a large effect. This was true, even when participants were quite unaware of the contingencies. It seems to reflect implicit learning.

Figure 15 about here

In another experiment we introduced contingencies in other irrelevant properties besides location (see Figure 16). We found that detection was also faster for the frequent pairings, when, for example, a solid blue bar target was usually vertical, or when a tilted broken target was usually pink, even though the targets were defined only by the first two features listed. This reflects some long term but implicit learning of object tokens with all their features bound together, even those irrelevant to the task.

Figure 16 about here

In an experiment that is so far unpublished we tried setting up contingencies in the distractor items rather than the targets. Chun and Jiang (1998) found implicit

priming from a consistent spatial layout of distractor items in a visual search task. We found the same. Participants apparently learned associations between particular distractors and the locations in which they typically appeared. They showed faster responses when the display conformed to those associations. But we observed a further surprising effect, suggesting very specific memorial matching. We analyzed the results according to the degree of contrast between a current target and the distractor that had in the past most often occupied the same location, and found a substantial effect on detection times. So if a particular location normally contained a green P, participants were quicker to detect the very different pink Q target in that location than a more similar green R. It was as if they were matching the new display against an expected spatial layout of object tokens and detecting mismatches faster when they were more salient. I think this fits well with the object file framework. It lends some reality to the notion that new objects are matched to the traces of past object tokens that occupied the same locations earlier in time. All these conjunction priming effects suggest that implicit traces of many bound object tokens can survive in memory for long periods, despite interference from multiple intervening displays, even when explicit access to them has apparently been lost or was never established.

Conclusions

(1) For explicit visual working memory, object tokens are stored in a limited capacity, vulnerable store that maintains the bindings of features for just 2 to 4 objects. Attention is required to sustain the memories. They are matched to new stimuli on the basis of their locations, and the object files are updated, erasing memory for their earlier states. In addition we retain some traces of unbound features, which allow better recognition of feature changes than of binding changes.

(2) Although the data do not yet rule out a unitary store, our evidence is consistent with the idea that VWM is a separate system from long term visual memory, specialized for different tasks and storing information in different forms. In the task we used, in which the bindings were not explicitly relevant, we found little effect of feature integrality. However, it is possible that there would be more effect of the presentation format if the features had not been perceptually grouped in both cases, or if participants were asked to recall the bindings explicitly, as in the experiment by Xu (2000).

(3) Perception also leaves implicit traces that can speed or slow performance when the tokens are presented again. Capacity here is much larger and the duration of the traces much longer. These implicit traces do meet some of the criteria for tokens. They bind features and they can be tagged for their relevance to action. It seems unlikely that simply changing the ease of activating previously used perceptual pathways could mediate these effects.

The story that begins, tentatively, to emerge can be summarized as follows: In perceiving the world we form representations first of the separate features that are present. We then bind the features through their locations to form integrated object tokens. Memory traces remain both of the separate features and of the object tokens. The former are independent of their locations, while the tokens are at least initially accessed through the locations in which the bindings were formed. Only a few bound tokens (2 to 4) can be explicitly maintained in working memory and subsequently accessed in recognition memory tasks. This requires continued attention. Without attention, or with new intervening presentations, the bindings are quickly lost., at least to explicit conscious awareness. However, many additional object tokens may remain implicitly available, to prime their later re-perception or to inhibit it if they were

previously ignored. Thus there seems to be a dissociation between, on the one hand, forming multiple bound object tokens that persist over time and may prime future perception, and, on the other hand, storing the tokens explicitly in a voluntarily accessible form. The main costs in human memory appear to relate to voluntary retrieval. Why the latter should pose such a severe memory bottleneck if the former is possible is an intriguing topic for future research.

Acknowledgements

The research was supported by grants from the NIH Conte Center # P50 MH62196, from the Israeli Binational Science Foundation, # 1000274, and from NIH, # RO1 MH 58383 on Visual Coding and the Deployment of Attention.

References

Anderson, J.A., (1983). The Architecture of Cognition. Cambridge, MA: Harvard University Press.

Atkinson, R. C. and Shiffrin, R. M. (1971). The control of short term memory. Scientific American, 225, 82-90.

Baddeley, A. (1986). Working memory. Oxford, Oxford University Press.

Baddeley, A. D. (2000). The episodic buffer: a new component of working memory? Trends in Cognitive Sciences, 4, 417-423

Braver, T., Barch, D., Kelley, WM, Buckner, RL, Cohen, NJ, Miezin, FM, Snyder, AZ, Ollinger, JM, Akbudak, E, Conturo, TE, & Petersen, SE (2001). Direct comparison of prefrontal cortex regions engaged by working and long-term memory tasks. Neuroimage, 14, 48-59.

Cabeza, R., Dolcos, F., Graham, R., & Nyberg, L (2002). Similarities and differences in the neural correlates of episodic memory retrieval and working memory. Neuroimage, 16, 317-330.

Cowan, N. (1995). Attention and memory: An integrated framework. Oxford, Oxford University Press.

Cowan, N. (1999). An embedded-processes model of working memory. Models of working memory. Miyake, A. and Shah, P. New York, Cambridge University Press: 62-101.

Della Sala, S. and Logie, R. H. (1993). When working memory does not work: The role of working memory in neuropsychology. Handbook of Neuropsychology, Volume 8. H. Spinnler and F. Boller. Amsterdam, Elsevier: pp. 1-62.

DeSchepper, B. and Treisman, A. (1996). Visual memory for novel shapes: Implicit coding without attention. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 27-47.

Duncan, J. (1984). Selective attention and the organization of visual information. Journal of Experimental Psychology: General, 113, 501-517.

Kahneman, D. and Treisman, A. (1984). Changing views of attention and automaticity. Varieties of Attention. R. Parasuraman and R. Davies. New York, Academic Press: 29-61.

Kahneman, D., Treisman, A. & Gibbs, B. (1992). The reviewing of object files: Object-specific integration of information. Cognitive Psychology, 24, 175-219.

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

Magnussen, S. (2000). Low level memory processes in vision. Trends in Neuroscience, 23: 247-251.

O'Regan, K. (1992). Solving the 'real' mysteries of visual perception. The world as an outside memory. Canadian Journal of Psychology, 46, 461-488.

Prabhakaran, V., Narayanan, K., Zhao, Z., & Gabrieli, J.D. (2000). Integration of diverse information in working memory within the frontal lobes. Nature Neuroscience, 3, 85-90.

Rensink, R. A., O'Regan, J. K. & Clark, J. (1997). To see or not to see: The need for attention to perceive changes in scenes. Psychological Science, 8, 368-373.

Rock, I. and Gutman, D. (1981). The effect of inattention and form perception. Journal of Experimental Psychology: Human Perception and Performance, 7, 275-285.

Simons, D. J. and Levin, D. T. (1997). Change blindness. Trends in Cognitive Sciences, 1, 261-267.

Tipper, S. P. (1985). The negative priming effect: Inhibitory effects of ignored primes. Quarterly Journal of Experimental Psychology, 37A, 571-590.

Treisman, A. (1992). Perceiving and re-perceiving objects. American Psychologist, 47, 862-875.

Treisman, A. and DeSchepper, B. (1996). Object tokens, attention, and visual memory. In Attention and Performance XVI: Information Integration in Perception and Communication. T. Inui and J. McClelland (Eds.). Cambridge, MA, MIT Press: pp. 15-46.

Treisman, A. and Gelade, G. (1980). A feature integration theory of attention. Cognitive Psychology, 12, 97-136.

Treisman, A., Kahneman, D., & Burkell, J. (1983). Perceptual objects and the cost of filtering. Perception and Psychophysics, 33, 527-532.

Treisman, A. and Schmidt, H. (1982). Illusory conjunctions in the perception of objects. Cognitive Psychology, 14, 107-141.

Treisman, A., Vieira, A. & Hayes, A. (1992). Automaticity and preattentive processing. American Journal of Psychology, 105, 341-362.

Wheeler, M. E. and Treisman, A. M. (2002). Binding in short-term visual memory. Journal of Experimental Psychology: General, 131, 48-64.

Wolfe, J. M., Oliva, A., Butcher, S. J. & Arsenio, H. C. (2002). An unbinding problem? The disintegration of visible, previously attended objects does not attract attention. Journal of Vision, 2, 256-271.

Xu, Y. (2002). Limitations in object-based feature encoding in visual short-term memory. Journal of Experimental Psychology: Human Perception and Performance, 28, 458-468.

Figure captions

Figure 1. Results of Wheeler & Treisman (2002) experiment on VWM for combinations of colors. There are no significant differences between these various patterns and the sets of six separate squares, suggesting that integrality has no effect on WM for color combinations.

Figure 2. Schematic description of experiment by Wheeler and Treisman on memory for colors and locations.

Figure 3. Results of Wheeler and Treisman experiment on VWM for color and location. With the Whole Display probe, accuracy of detection of changes of binding is significantly lower than detection of changes of either feature separately. The difference disappears with the Single Item probe.

Figure 4. Results of Wheeler and Treisman experiment on VWM for color and shape. Again with the Whole Display probe, memory for binding is worse than memory for the separate features, whereas with the Single Item probe there is no difference.

Figure 5. Design of experiment by Treisman & Zhang on memory for features with the same or different binding and location.

Figure 6. Different conditions in the feature recognition experiment by Treisman & Zhang.

Figure 7. Accuracy of detection of feature changes in Treisman & Zhang experiment as a function of whether the binding and locations were changed or not. Changing the

binding is disruptive. Changing the location has no overall effect, but there is an interaction with changed binding, suggesting that locations are stored only with the bound form of feature memory.

Figure 8. Accuracy of detection of changes of binding in experiment 2 by Treisman and Zhang, as a function of location changes and of probe delay. Changes of location are much more disruptive than they were for detection of feature changes, and the disruption is greater at the shorter delay.

Figure 9. Accuracy of detection of changes of binding as a function of articulatory suppression and of visual word interference.

Figure 10. Three possible relationships between WM and LTM. (1) The traditional sequential model. (2) WM as a temporarily activated subset of LTM. (3) Separate and independent stores with the possibility of parallel access.

Figure 11. Design of experiment by Treisman and Dishon-Berkowitz. Participants attended either to color or to shape, and the experiment explored the effect of consistent pairings of the two features on both LTM and WM.

Figure 12. Different forms of integrality or separability of shape and color in two versions of the experiment by Treisman and Dishon-Berkowitz.

Figure 13. Example of displays in negative priming experiment by DeSchepper and Treisman (1996). Participants decided whether the green shape in the overlapped pair matched the white shape to the right. The ignored red shape on one trial could become

the attended green shape on a subsequent trial, giving slightly longer mean response latencies.

Figure 14. Negative and positive priming results at different delays for participants divided according to whether they had shown negative priming or not in an initial test with a different set of shapes.

Figure 15. Example of display in search for conjunction targets (e.g. pink Q and green R) with two targets appearing mostly in consistent locations and two others in randomly varying locations.

Figure 16. Example of display in search for conjunction targets (e.g. solid vertical bar) with irrelevant feature (e.g. blue) consistently paired on most trials..

Figure 1



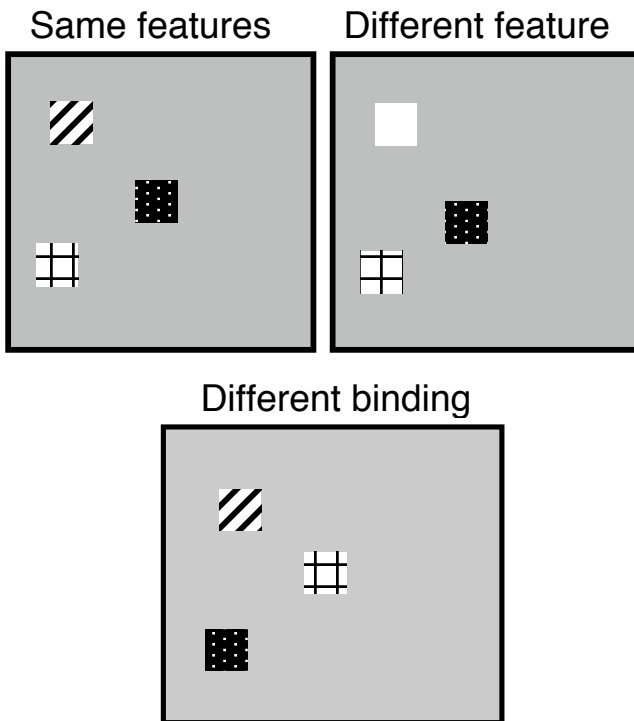


Figure 2



(a) Test: Whole display



(b) Test Single Probe

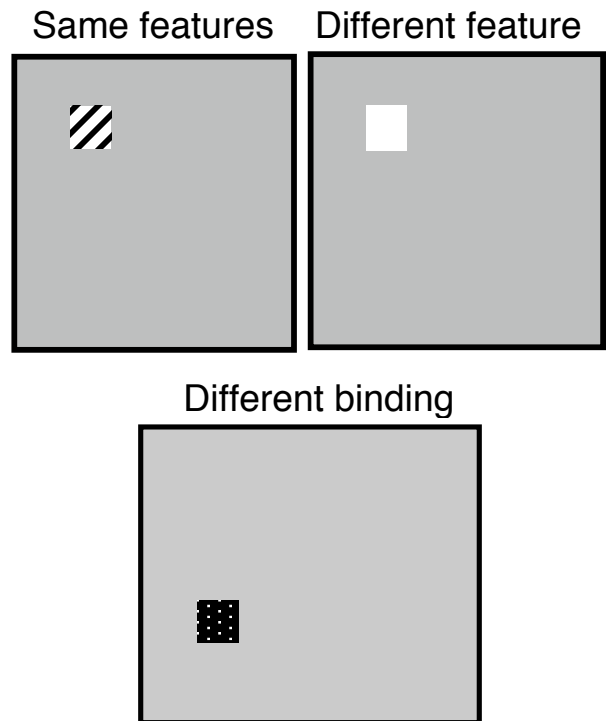


Figure 3

Location and color

(a) Whole display as probe

(b) Single Probe




Figure 4

Shape and color

(a) Whole display as probe

(b) Single Probe

% Correct





Figure 5

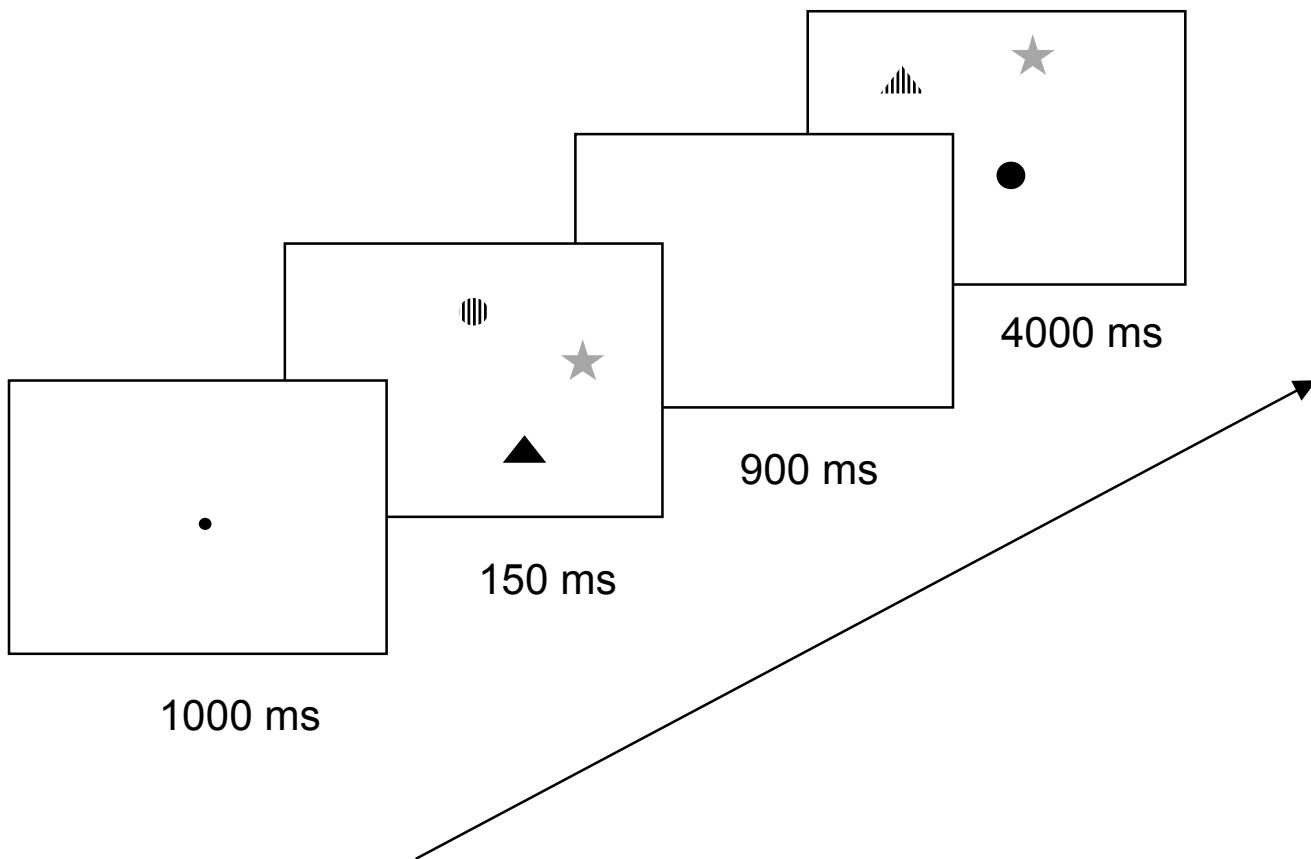


Figure 6

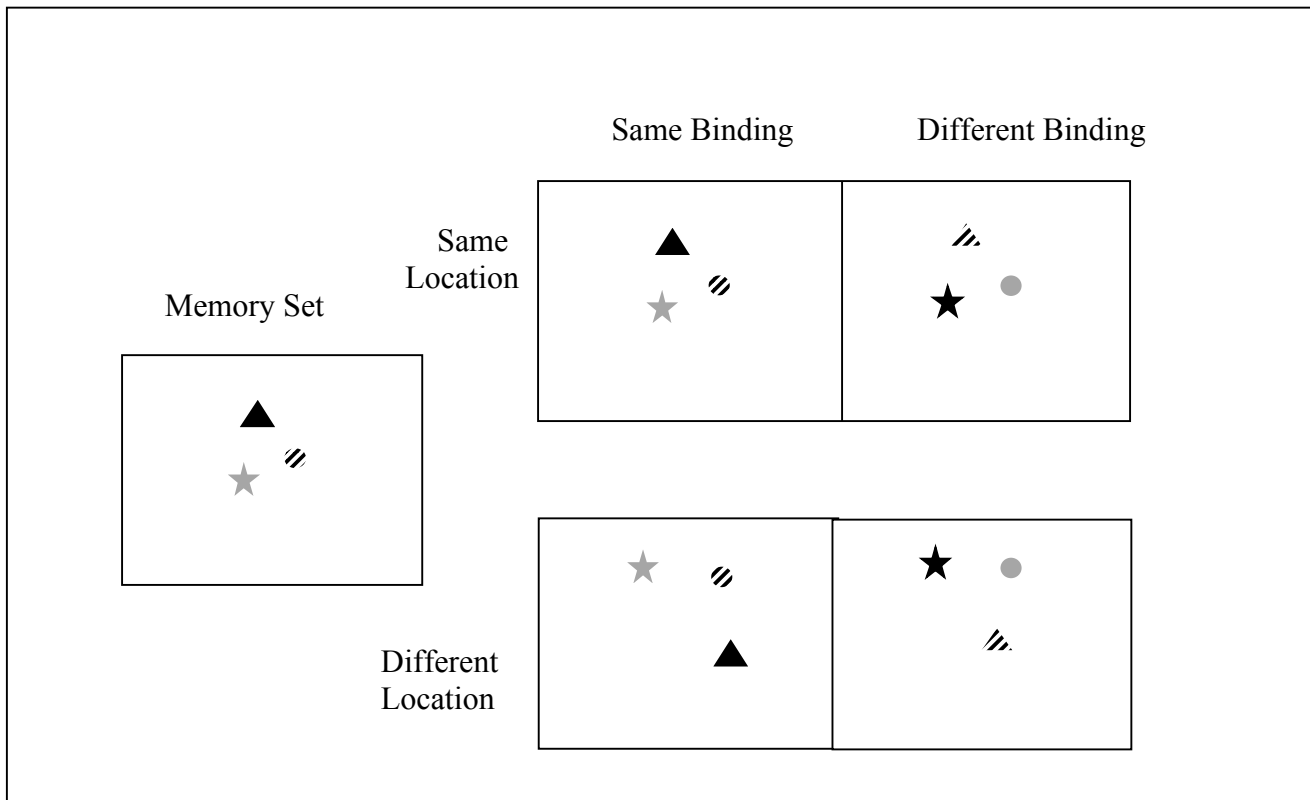


Figure 7

VWM for Features

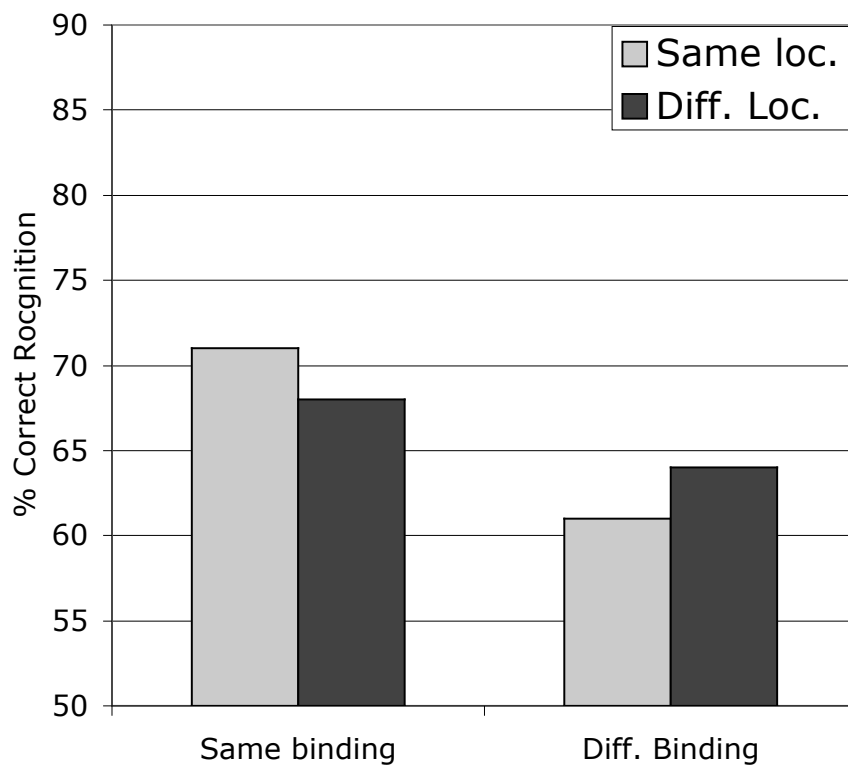


Figure 8

VWM for Binding: Effect of Exposure Duration

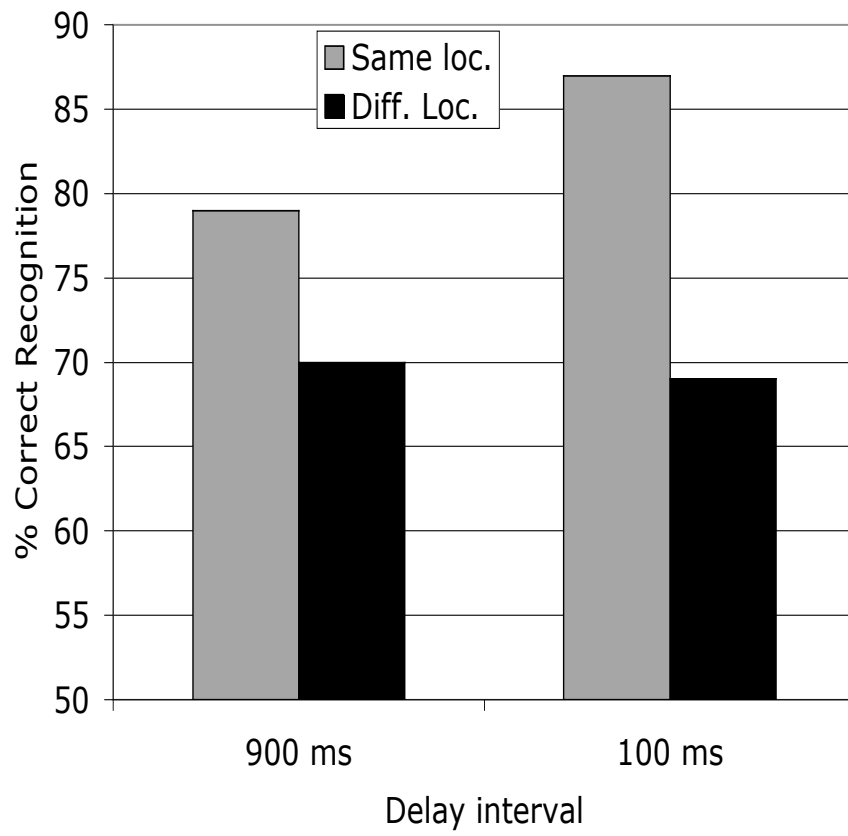


Figure 9

Interference effects in VWM

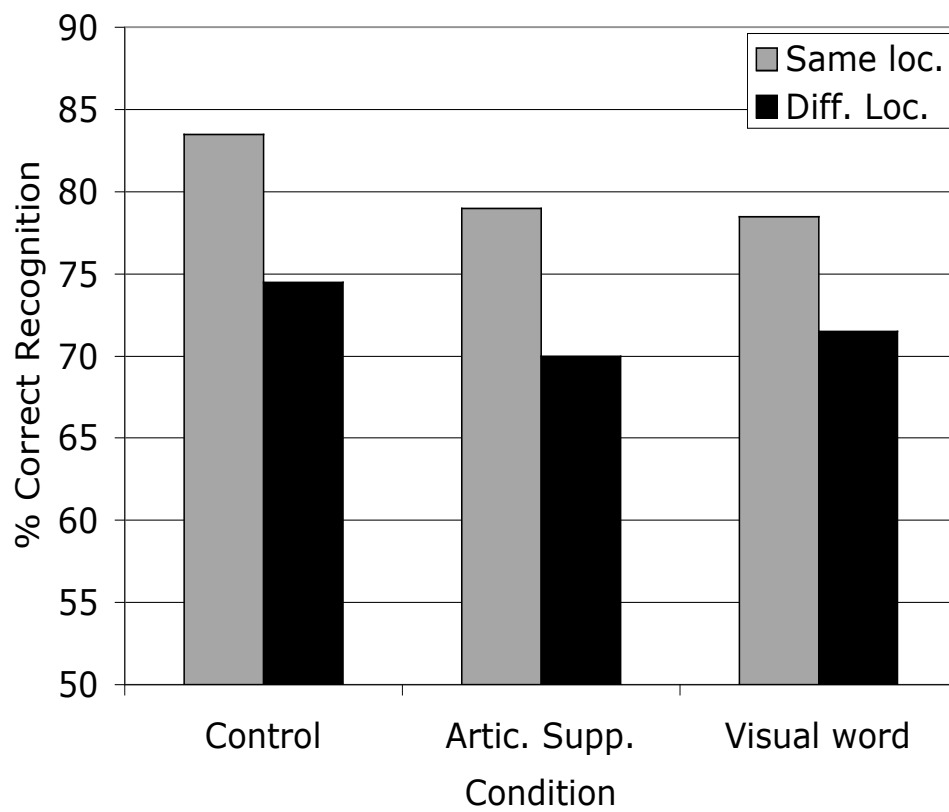


Figure 10

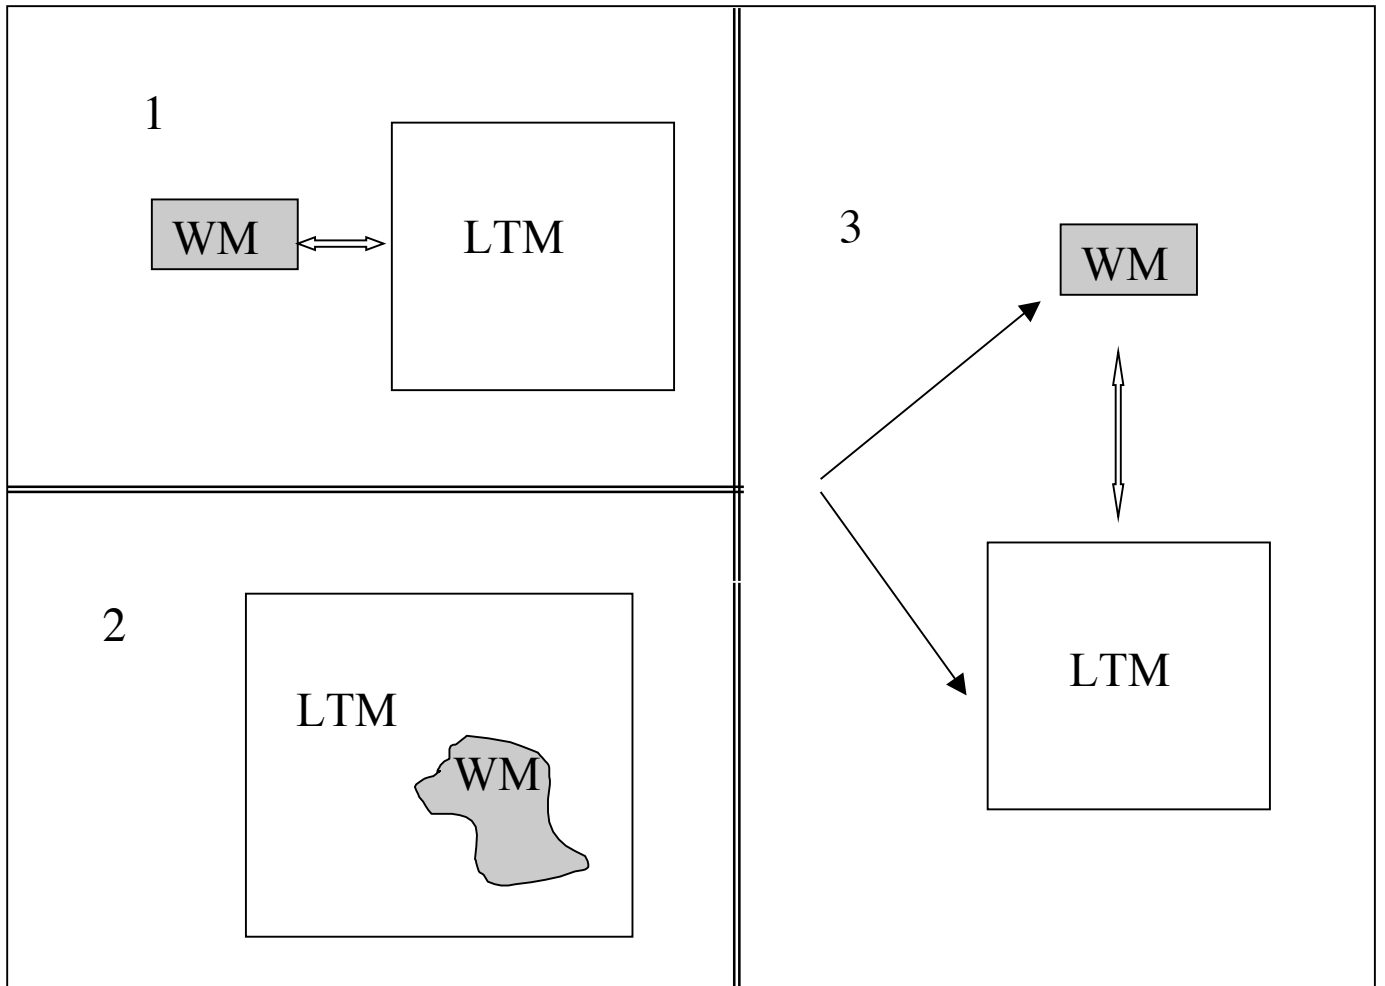


Figure 11




Figure 12

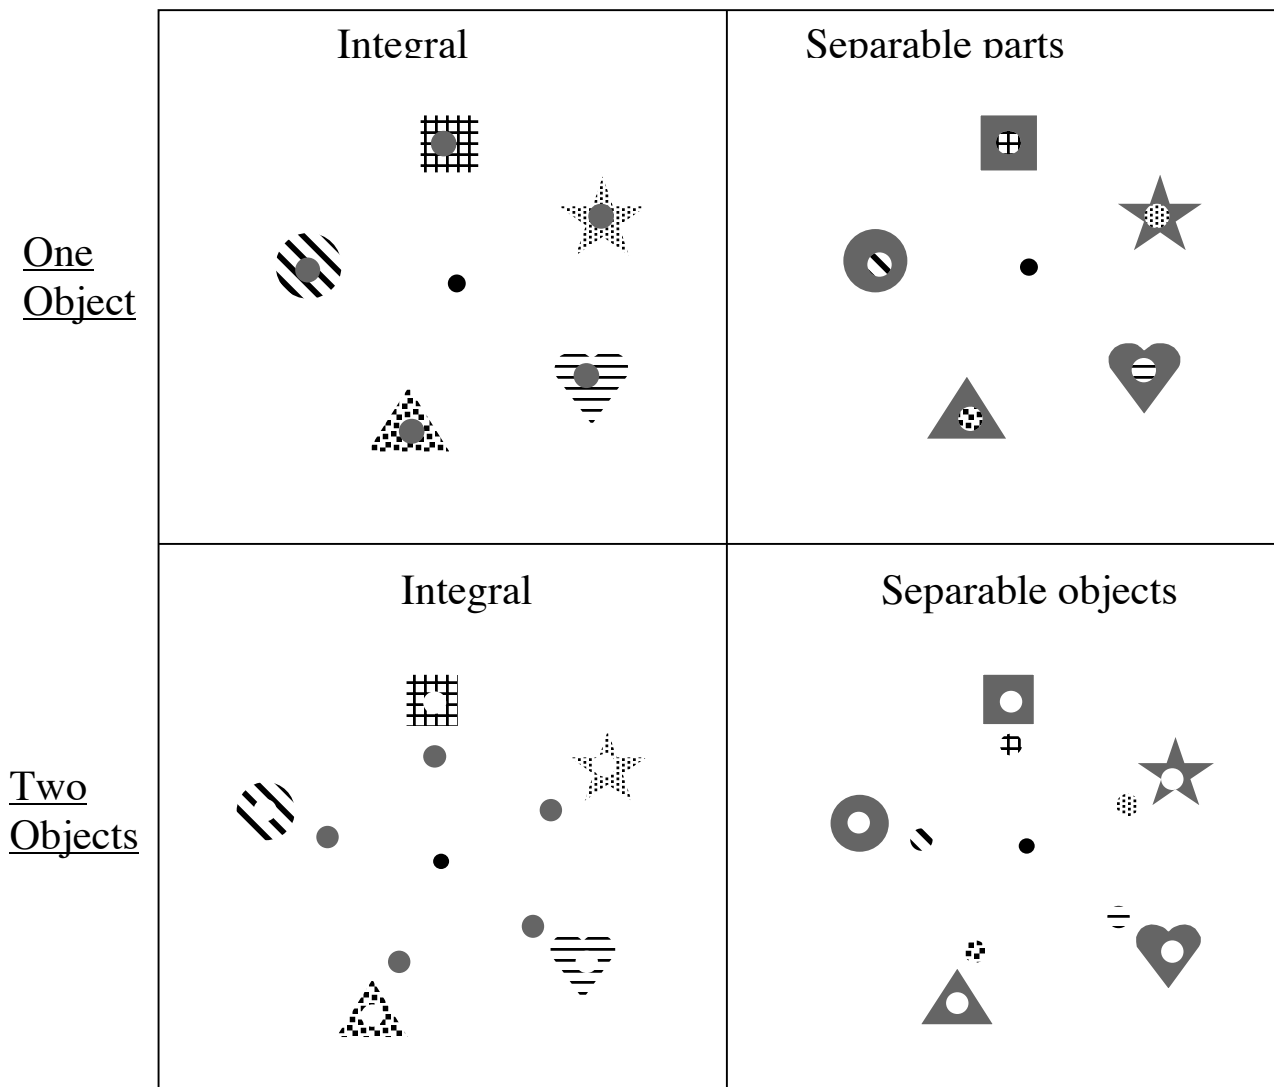


Figure 13




Figure 14




Figure 15

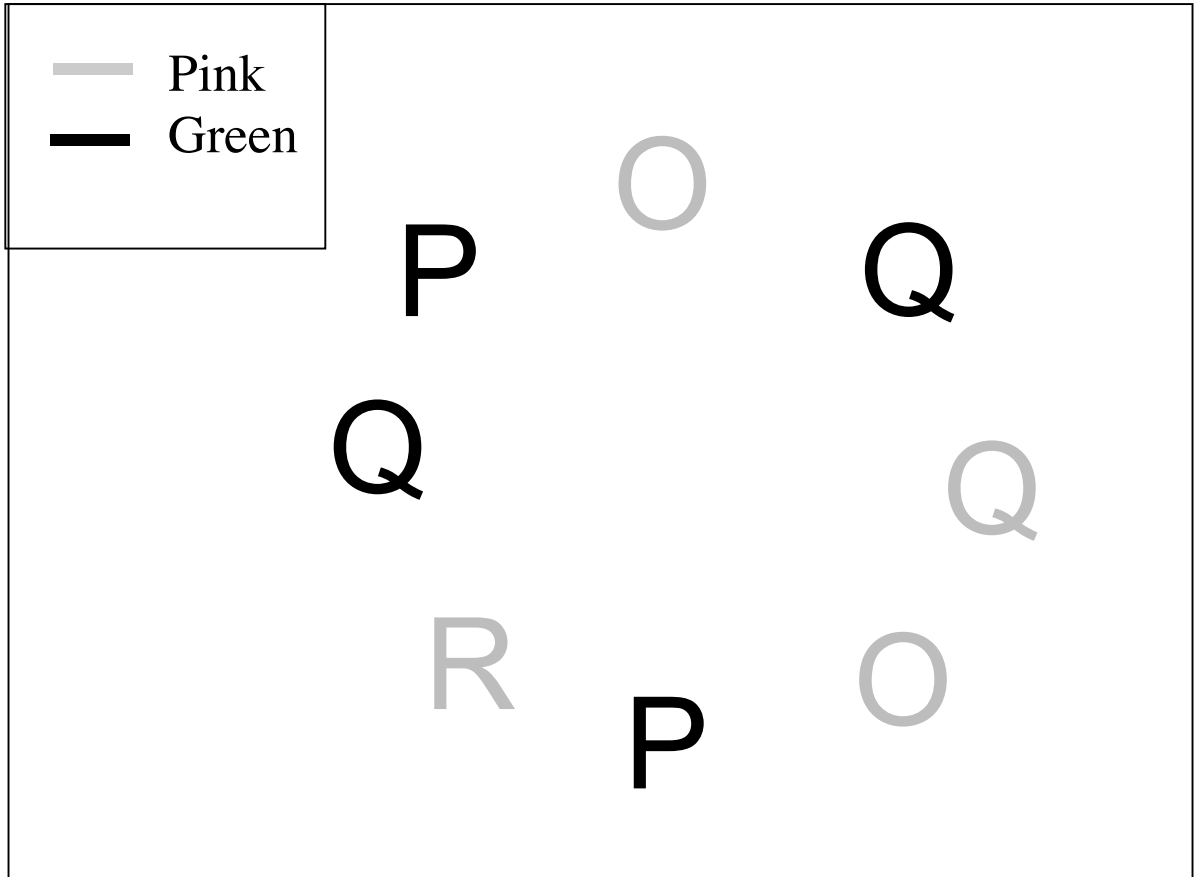


Figure 16

