4 Hybrid Search

Picking up a Thread from Schneider and Shiffrin (1977)

Jeremy M. Wolfe, Trafton Drew, and Sage E. P. Boettcher

"Rabbi Ben Bag-Bag used to say of the Torah (the first five books of the Bible): Turn it and turn it again, for everything is in it. Pore over it, and wax gray and old over it." (Pirkei Avot, 5:25). Now, one would not want to mistake the works of Richard Shiffrin for Holy Scripture; however, we have returned to those works many times as at least one of the authors of this chapter has waxed somewhat older and grayer. In particular, those of us with interests in visual search have frequent reasons to return to the pair of papers he wrote with Walter Schneider in 1977 (Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). For a surprising number of topics, there turns out to be something relevant in Schneider and Shiffrin (1977; henceforth, "SS77" will refer to the pair of papers). This is especially true when one becomes interested in tasks that combine visual search and memory search. We decided to name such tasks "hybrid search" (Wolfe, 2012). We subsequently noticed that SS77 was there first with that name, talking about "hybrid, visual, memory-search tasks" (p8). Indeed, the basic experiments of those papers are hybrid search tasks where observers look through a visual display for any of several targets, held in memory.

Shiffrin and Schneider were particularly interested in the distinction between versions of hybrid tasks that involved "consistent" versus "variable mapping". In consistent mapping tasks, the targets are always drawn from one set of items and the distractor items are always drawn from a different set. In the basic SS77 version, targets might be numbers while distractors were letters, or vice versa. In variable mapping, targets on one trial could be used as distractors on another. In SS77, different, arbitrary sets of letters (or numbers) served as the targets and distractors, changing on each trial. On each trial, observers saw a memory set of 1, 2, or 4 items and then a briefly presented search array with 1, 2, or 4 characters. There are many variations of this basic design. For instance, if the search display is clearly presented, response time (RT) is the measure of interest and errors are rare. If the visual search stimuli are presented briefly and masked, the pattern of errors becomes meaningful. Schneider and Shiffrin found that consistent mapping tasks were easier than variable mapping. In variable mapping tasks, RTs were roughly linear functions of the product of memory and visual set sizes.

The results were consistent with a serial, self-terminating search in which items in the visual and/or memory sets were examined, one after another, until the target was found or the search was terminated. With an overlearned, consistently mapped task like numbers among letters or vice versa, RT was often independent of set size. According to Schneider and Shiffrin, this reflected *automatic* processing of these stimuli.

As noted, SS77 used small visual and memory set sizes. The memory sets, in particular, would fit into visual working memory, with its capacity of about four items (Baddeley & Hitch, 1974; Cowan, 2001). We are interested in hybrid search behavior when the memory set is much larger than what was used in SS77 and related work, and much larger than any reasonable estimate of the capacity of working memory. Consider what we could call the Facebook version of hybrid search. For the sake of the example, you have 1000 Facebook friends. One of those friends posts a picture of a crowd of 100 people on their page and you want to determine if any your 1000 friends is in the photo. That is a visual set size of 100 and a memory set size of 1000. If we assume that RT is a linear function of the product of visual set size and memory set size and if we use the 40–50 msec/item slope, found in SS77 and in many studies of visual and/or memory search, that would yield a search time on the order of 1000 x 100 x 50 = 5,000,000 msec = 83 minutes. That would be the time for an exhaustive search. A serial, self-terminating search might be expected to take a mere 40 minutes or so on average. Clearly, there is something wrong here because we know that we could perform a task of this sort and we know that it would not take the better part of an hour.

The question can be addressed by using photorealistic pictures of objects as stimuli, rather than alphanumeric characters. Since Shepard (1967) and, particularly, Standing (1973; Standing, Conezio, & Haber, 1970), we have known that humans have massive memory for pictures of scenes and objects. More recent work has made it clear that this memory is surprisingly detailed and very easy to produce (Brady, Konkle, Alvarez, & Oliva, 2008; Konkle, Brady, Alvarez, & Oliva, 2010). Two or three seconds of exposure per image is all that is required to encode thousands of images into a form of long-term memory, adequate to perform old/new recognition tasks at high (typically > 85%) accuracy for hours or days, even when the distractor items are quite similar to the target items. This capability allowed us to do a hybrid search experiment in which observers are given a memory set of 1, 2, 4, 8, 16, or 100 items. Once that set is in memory, our observers performed several hundred visual search trials with visual set sizes of 1, 2, 4, 8, or 16 items.

The basic result of this experiment is shown in Figure 4.1, replotted from Wolfe (2012). On the left, the data are plotted with the visual set size on the x-axis. The different functions are different memory set sizes as labeled from 1 to 100. The figure shows correct target-present data, averaged over 10 observers. Lines show the results of linear regression. The pattern of

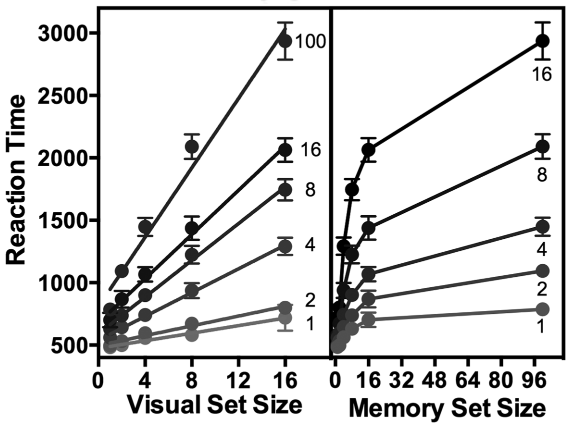

Figure 4.1 Results of the basic hybrid search experiment. Figure 4.1a plots RTs as a function of visual set size with each curve representing a different memory set size as labeled. Figure 4.1b plots the same data as a function of memory set size with each curve representing a different visual set size.

results is the same for target-absent but the slopes of the regression lines are about twice as steep. For present purposes, the important result is that the RT x visual set size functions are essentially linear. This is consistent with other studies of visual search for arbitrary objects in random arrays when the memory set size is one (Vickery, King, & Jiang, 2005; Wolfe, Alvarez, Rosenholtz, Kuzmova, & Sherman, 2011). These data show that the visual search remains linear even when the memory set size gets large.

On the right, the same data are replotted with memory set size on the x-axis and visual set size as the parameter that labels each curve. As should be obvious, these functions are curvilinear, rather than linear.

Figure 4.2 replots the data with memory set size on a log scale. It appears that RT rises linearly with log of the memory set size. Wolfe (2012) found a logarithmic function was a very good fit to the data. One way to show this is to use the data from memory set sizes of 1–16 to predict the RT for a memory set size of 100. In Figure 4.2, the solid lines show the best fitting linear regression for set sizes 1–16 and the dashed lines show the extrapolations of those lines to a memory set size of 100. The log estimate is a remarkably good predictor of the actual data. Returning to Figure 4.1b, we can see that a linear regression on set sizes 1–16 would vastly overestimate the RT at set size 100 on a linear scale. The precision of the logarithmic fit is all the more surprising because the data for memory set size 100 come from a group of observers separate from those who gave the data for the smaller memory set sizes. It should be noted that the curvilinear form of these functions is not a result of higher errors at the larger set sizes. Error rates are quite low and Wolfe (2012) found the same results with a localization version of the task

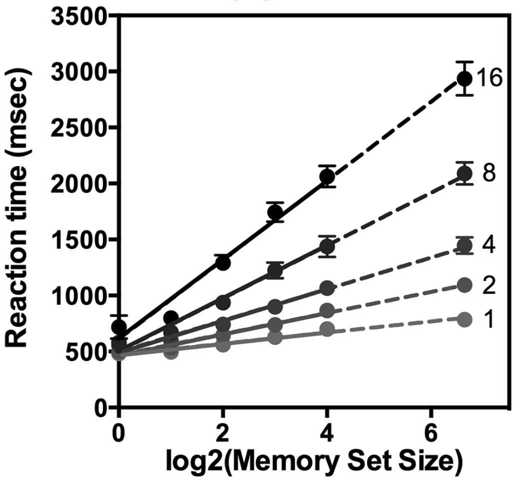

Figure 4.2 Data from Figure 4.1b, replotted on a log scale. Lines are linear regressions of RT on log2 (memory set size) for set sizes 1–16. Dashed lines are the extrapolation to set size 100, showing that the data are well-modeled as a logarithmic function.

(targets present on 100% of trials), which makes the error rate nearly zero in this task.

Very large memory set sizes were possible in this case because observers can easily hold 100 specific objects as some subset of long-term memory and can use that memory to perform the task. Indeed, in another version of the experiment, observers could still perform the task with 500 objects in memory. In that experiment, each trial of the experiment contained a single new object that served as the target for that trial. The target on the current trial became a member of the memory set on the next trial. There was no explicit memorization phase of this experiment. Observers simply increased their memory set size by one on each trial. By the time they were holding 500 items in memory, they were starting to make an appreciable number of errors but the RT x memory set size functions were still strongly curvilinear and roughly logarithmic (Cunningham & Wolfe, 2012).

FAMILIARITY AND LONG-TERM MEMORY

We came to the question of hybrid search from an interest in visual search. In the visual search literature, there is an assumption that some representation or "template" of what you are looking for is held in some version of working memory (e.g. Olivers, Peters, Houtkamp, & Roelfsema, 2011). However, working memory is very limited in its capacity. A set of 16 or 100 or 500 objects is not being maintained in working memory, as that term is usually understood. Our observers must actually determine if an item in the visual display is a member of a set of items in some long-term memory store. Thus, a hybrid search task becomes a sequence of recognition memory tasks of the sort that Shiffrin has spent many years studying. Though this chapter is framed as an homage to SS77, it would be a mistake to ignore Shiffrin's massive body of subsequent work on long-term memory mechanisms. Familiarity plays a very important role in those accounts of recognition memory (Gillund & Shiffrin, 1984; Shiffrin & Steyvers, 1997). Thus, it is worth saying a word about familiarity in these hybrid search experiments.

All of the items in our memory sets were studied for approximately the same amount of time, though this means that learning the longer memory sets took longer in aggregate than learning a smaller set. Given a random selection of targets from the memory set on target-present trials, a specific item in the memory set would be tested more frequently when the memory set was small than when it was large. So, over a block of a few hundred trials, items from small memory set sizes would become more familiar and, thus, would be responded to more quickly if familiarity is the relevant signal (Nelson & Shiffrin, 2013). Even if all items were equally familiar at large and small set sizes, Shiffrin's models, as well as others, would predict an effect of memory set size because the global familiarity signal becomes noisier as the memory set size increases. This mechanism should influence the matching process regardless of how many times a given item is tested.

In Wolfe (2012), we had about 2000 objects to deploy as distractors. These were sampled without replacement until they ran out. Then they were recycled. A block of 500 trials with set sizes 1–16 required about 3000 distractors. Thus, each distractor or foil would be seen just once or twice in a block. However, observers ran several blocks with the same items so the foils would become increasingly familiar. Moreover, targets on one block could appear as distractors on subsequent blocks. In experiments where the task was to find the new item on each trial, the foils were the old, more familiar items. This creates a complex familiarity landscape. None of the manipulations of familiarity in our experiments should cripple performance, but they should have effects on RT and accuracy based on Shiffrin's work. As he and Nelson say, "very precise and clever studies combined with careful modeling will likely be needed" (Nelson & Shiffrin, 2013, p. 382) to characterize those effects. In recent work with Nosofsky, Cox, and Cao (2014), Shiffrin has done some of that careful research. Using the same images that we used, they deliberately manipulated and successfully modeled relative familiarity in the case where the visual set size is 1.

HOW GENERAL IS THIS RESULT?

This result, linear search through the visual display and logarithmic search through the memory set, might have been limited to targets defined by specific views of specific items. That would be interesting but of relatively limited application to real world search. As a general rule, real world targets are more generically defined. You might go to the grocery store with a memory set of "targets" to place in your shopping cart but you would not, typically, go on a search for precisely this T-bone steak. You would be searching for something that could appear in many guises. It could be specified at any of several levels. A T-bone steak target might be considered to be specified at a subordinate level of categorization. "Steak" might be the basic- or entry-level target category, with "meat" or "dinner" as superordinate categories (Jolicoeur, Gluck, & Kosslyn, 1984; Rosch, 1973).

In a different version of the experiment (Cunningham & Wolfe, 2014), 10 observers memorized 1, 2, 4, or 8 categories like Plants, Furniture, Animals, Weapons, and Musical Instruments. On each target-present trial, one of the objects in the visual display would be drawn from one of the target categories in the memory set but, unlike the previous experiments, this target would never have been seen before. That is, you might be looking for "Furniture" but you would not have previously been exposed to the sofa that happened to be the target on this trial. This is a harder task than looking for specific items. To begin, observers cannot trivially memorize 100 target categories the way that they can memorize 100 pictures of items. Even with a smaller number of categories in memory, the RTs for the search for instances of those categories are slower than the RTs shown in Figure 4.1. Deciding that this cat is one of the four specific items in your memory set is faster than deciding if the same cat is a member of any of four categories held in memory. Nevertheless, the pattern of RTs seen above remains. RTs are a linear function of the visual set size and a curvilinear function of the memory set size (the number of categories). Using the lower set sizes to predict the highest memory set size is much more accurate if the memory set sizes are log transformed (as in Figure 4.2) than if they are not.

It is worth mentioning that SS77 discusses the formation of categories extensively. With extensive practice, it was found that sets of characters became "automatized", as if they were being responded to as a single thing (see also Czerwinski, Lightfoot, & Shiffrin, 1992).

Cunningham and Wolfe (2014) obtained similar results with memory sets of an arbitrary 1, 2, 4, or 8 alphanumeric characters. The memory set stayed fixed for 350 trials while observers search through 1–16 characters in each visual display, so, in the language of SS77, this was consistent mapping. In SS77, consistent mapping yields almost no effect of memory set size on RT. However, the SS77 consistent mapping task is somewhat different from the consistent mapping condition that we used. In SS77, the memory sets were always numbers and the distractor items in the visual displays were always letters. Observers could do that task either by looking for the 1, 2, or 4 specific numbers or just looking for *any* number. SS77 also discusses a consistent mapping experiment by Briggs and Johnsen (1973) that uses a method closer to ours. If one looks at the replotting of Briggs and Johnsen's data in Figures 7 and 8 of SS77, one can, at least, imagine that the RT x memory set size functions are curved while the RT x visual set size functions are more linear, though any inferences about the shape of functions, based on set sizes of 1, 2, & 4, must be rather tenuous.

In addition to objects, categories, and alphanumeric characters, Boettcher and Wolfe (2014) used lists of words as hybrid search memory sets. There are two versions of this experiment. In one, observers memorized arbitrary lists of words. In the other, memory sets were the words from passages that the observers had previously committed to memory. These tended to be poems, song lyrics, or passages like the US Pledge of Allegiance, etc. The longest of these word lists was 86 words long. Again, we obtained linear effects of visual set size and curvilinear effects of memory set size. Our distractors were matched in word length to the target words. With the memorized texts, we had expected to get strong serial position effects. That is, it seemed likely that an observer who had memorized "London Bridge is Falling Down" would be faster to report the presence of "London" than "falling" in the visual display, but this was not the case. Any serial position effects were small and unreliable, clearly indicating that, even for highly ordered lists, observers were not searching memory by starting at the beginning and searching to the end, perhaps because these are long-term memory searches where, in Shiffrin's view, familiarity will be the driving force. In contrast, in variable mapping short-term memory search experiments, when a new list is presented in a specific order on each trial, position of the item in that memory list has a much stronger effect (Nosofsky, Little, Donkin, & Fific, 2011). It is worth mentioning that, in our experiments with these overlearned phrases, contextually salient words like "flag" in the US Pledge of Allegiance produced results no different from less salient words like "under".

WHAT DO THE CURVILINEAR RT X MEMORY SET SIZE FUNCTIONS TELL US?

Curvilinear functions in consistent mapping memory search experiments are actually quite common, though, because of the typically small and restricted range of memory sets used, the shape of the function is often somewhat ambiguous. In one of the few older studies to use larger memory set sizes, Burrows and Okada (1975) found a curve that they suggested could be modeled as either a bilinear or a log function. Schneider and Shiffrin (1977) described curvilinearity as by-product of practice; an intermediate step on the way from serial to automatic processing in short-term memory work (SS77, p8). In long-term memory, curvilinearity can be a very persistent feature of memory search data (Nelson & Shiffrin, 2013). Sometimes, the curvilinear function survives more or less unchanged over a great deal of practice, as in Kristofferson's (1972) month-long experiment. Indeed, curvilinearity seems be a natural outcome of the application of Shiffrin's ideas about the role of "global-familiarity" signals (Nosofsky, Cox, Cao, & Shiffrin; 2014—see further discussion below.)

Beyond memory search, in various tasks in which observers must choose between multiple alternatives, RT is often a log function of the number of alternatives. This is Hick's Law or the Hick-Hyman Law (Hick, 1952; Hyman, 1953; Nelson & Shiffrin, 2013; Schneider & Anderson, 2011). It is possible to see the comparison of a visual item with the N items in a memory set as just another N-alternative choice (or N+1 if none-of-the-above is an option). As SS77 and others note, a log function would fall out naturally if the path to the decision "proceeds according to a series of binary subdecisions" (p41). This is like the old game in which you might have picked a number between 1 and 100 and your little brother had to try to guess that number (a 100-alternative choice). If he picked at random, the average number of choices required would be 50 but if your little brother was allowed to ask a set of questions, like "Is it greater than 50", then he could make that set of "binary subdecisions" that would reduce the number of steps required to an average of log2 of the maximum number. In our hybrid search experiments, this would work if observers could implicitly determine that the current visual item didn't match anyone in *this half* of the memory set in step 1 nor in this remaining half in step 2 and so forth. However, it is not obvious that the space of arbitrary objects, used in our basic experiment, is organized in a manner that would permit this set of binary decisions (DiCarlo & Cox, 2007).

There are other routes to logarithmic functions. Curvilinearity, at least, would be a consequence of the global-familiarity mechanism in Shiffrin's SAM (Search of Associative Memory; Gillund and Shiffrin, 1984). In SAM, as more items are added to the memory set, the match of an item in the search display to its trace in memory would be more likely to get lost in noisy mismatches. Similar logic produces essentially logarithmic functions from diffusion (Ratcliff, 1978) or accumulator (Brown & Heathcote, 2008) models of recognition and decision (Donkin, Brown, Heathcote, & Wagenmakers, 2011). Suppose that identification of an item as the target of a search involves accumulation of information about that target to a decision boundary. If an item is the target, the information reaches the boundary and a positive decision is made. If the item is a distractor, information accumulates randomly or toward a negative boundary and the item is rejected at some point, unless the random accumulations reaches the decision bound by chance. This would produce a false alarm error. To avoid excessive errors, the decision bound must be placed far enough away from the starting point of the accumulation. However, since the distance to the bound determines RT, the desire for speed will push the bound back down. Now consider what happens if there multiple possible targets and, thus, multiple accumulations of information. This will increase the chance of error because any one of the accumulations might reach the decision bound by chance. Consequently, that decision bound must be moved further from the start point, slowing the RT. Our Monte Carlo simulations and rather more rigorous modeling by Leite & Ratcliff (2010) show that if the bound is moved in a manner that keeps error rates constant, RT must increase with increasing set size. This increase appears to approximate the logarithmic RT x set size functions found in hybrid search tasks.

Nosofsky, Cox, Cao, and Shiffrin (2014) hold that rate with which information accumulates decreases as memory set size increases. They argue that this produces a better fit to the combination of RT and error data than moving the decision boundary. Logan, Van Zandt, Verbruggen, & Wagenmakers (2014) propose a similar account of multiple alternative decisions: the rate of accumulation of information toward the decision boundary decreases as the number of alternatives increases. Under the assumptions of their model, this, too, can produce logarithmic increases in RT with increases in the number of alternatives.

INTERACTING COMPONENTS OF HYBRID SEARCH

How do visual search and memory search interact in hybrid search? Three steps are cartooned in Figure 4.3. Assuming a memory set already in memory, the first step will be selection of an item from the visual display as a possible target. That deployment of visual attention will be "guided". If, for example, all of the items in the target set are known to be red, attention will be guided to red items in the visual array (Egeth, Virzi, & Garbart, 1984).

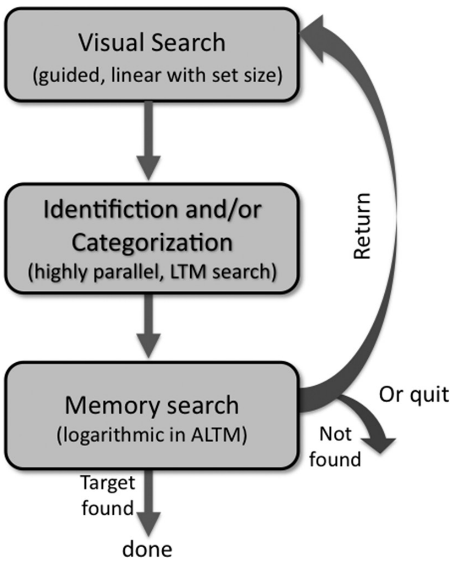

Figure 4.3 Three steps in hybrid search.

This guidance will be based on basic features of the items and not on their categorical status (Wolfe & Horowitz, 2004). This can be a little complicated because some categories of objects can be distinguished from others by their basic features. This complication can be illustrated in another hybrid search experiment. Cunningham and Wolfe (2014) had observers search for memory sets composed of specific instances of animals. That is, observers knew they were looking for *this* chicken or *that* alligator and they knew that all the items were animals. Some of the items in the visual display were non-target animals. Some were plants and clothing and some were picture frames and money. We found that RTs increased with the number of distractors from the plant or clothing categories but not with the number of distractors from the money and picture frame categories (Cunningham & Wolfe, 2014). The difference is that basic features of money and picture frames rule them out as possible animals and, thus, attention can be guided away from them. It was as though those items were not in the display at all.

The second step is to identify and/or categorize the item and the third step is to check the results of the second step against the memory set. This can be illustrated in the same, search-for-specific-animals experiment just described. Adding clothes or plants to the visual set increases the RT by a modest amount (about 25 msec) and that amount does not change with the number of animals in the memory set. The RT cost doesn't change because, in this experiment, all the targets are animals, so a non-animal can be rejected with the same simple decision whether there are two or twenty animals in the memory set. The categorical status is enough to do the task, even if, in the case of clothes and plants, the categorical status was not adequate to guide visual selection in Step 1.

In contrast, adding non-target animals to the visual set increases the RT by an amount that increases with the log of the number of animals in the target set. If an item selected from the visual display (Step 1) is identified as an aardvark (Step 2), it must then be checked against the specific animals in the memory set and, as we have seen, that step is logarithmic with memory set size.

Thus, hybrid search can be thought of as involving two memory searches. In Step 2, observers somehow search the entire contents of long-term memory to identify the object. In Step 3, if need be, they search the memory set. Recent evoked potential data seem to show just this sort of time course. Nako, Wu, and Eimer (2014) found EEG signals indicating the presence of the category ("letter") earlier than signals reflecting the search through a memory set of specific letters.

Where does that memory set reside? SS77 conceived of the memory set for those experiments as living in a short-term memory. However, short-term memory stores like visual working memory are extremely limited, with typical estimates of no more than four simple visual objects (Cowan, 2001; Luck & Vogel, 1997). Moreover, short-term memories are, by definition, short term. They do not persist, whereas the large memory sets used in hybrid search experiments (and in other massive memory experiments) can last for many minutes or more. They must be held elsewhere, in some form of long-term memory (LTM). This form of LTM must be a portion of LTM, not the whole memory. Otherwise, RT would not change with the size of the memory set. There are multiple candidates for the name of this portion of LTM. We favor the term "Activated Long-Term Memory" (ALTM) as a name for the portion of LTM that is relevant to the current task (Cowan, 1995). In a similar manner, in Shiffrin's SAM model, associations with the testing context (e.g. test probes) have the effect of activating the relevant items stored in long-term memory (Raaijmakers and Shiffrin, 1981).

If the memory set is in ALTM, how might it interact with the contents of working memory during the course of a hybrid search? The question arises because there has been considerable work in recent years on the interactions of working memory with visual search (Downing & Dodds, 2004; Han & Kim, 2004; Soto, Humphreys, & Heinke, 2006; Woodman, Vogel, & Luck, 2001) governed by various versions of the hypothesis that the search template is resident in working memory (Oberauer & Hein, 2012; Olivers et al., 2011). How could working memory hold the template when, for example, 100 items are in the target set? We had thought that hybrid search might involve rapidly shuttling items from ALTM in and out of working memory. If so, we reasoned, then if we disabled working memory, hybrid search would be crippled as well, leading to much slower rates of search through memory. In fact, to our surprise, we have been unable to find a significant cost of a working memory load on the rate of search through memory in hybrid search experiments, though we have tried a variety of ways to load working memory (Drew, Boettcher, & Wolfe, 2013).

Though the steps of Figure 4.3 have been presented as a sequence, the operations in each step are not entirely mutually exclusive. Based on the slopes of RT x visual set size functions, visual selection can proceed at a rate of 20–40 items per second. Yet, the identification step, Step 2, appears to take at least 100 msec (Kirchner & Thorpe, 2006). Thus, it would appear that multiple items can be in the process of being identified at the same time. This could reflect parallel selection of a group of visual items. We prefer to think in terms of "pipeline" operations, where items can be selected for entry into the identification pipeline but where multiple items can be in that pipeline at any one time. A carwash serves as a good metaphor. A new car might enter the carwash every minute (the selection step), while it might take five minutes for the car to get washed (the identification step). This can work because multiple cars are being washed "in parallel" (Wolfe, 2003).

Can observers conduct a memory search while carrying out other operations? We have an indication that this is the case from combining memory search with the attentional blink (AB; Drew, Sherman, Boettcher, & Wolfe, 2014). In the AB, observers typically look for two targets in a rapidly presented stream of single items. If the second target appears in a window of about 200–500 msec after the first, its detection is impaired (Raymond, Shapiro, & Arnell, 1992). If processing of the first target (T1) is made more demanding, the blink becomes more pronounced (Dux & Marois, 2009). In our experiments, identification of T1 required a memory search (Drew et al., 2014). T1 could be any of 1–16 objects. The second target in the stream of items (T2) was a number that needed to be identified. T1 difficulty increased as the memory set size increased. Interestingly, in a number of experiments, while this T1-T2 combination produced a strong AB, the size of the T1 memory set did not have an impact on the AB unless T1 was strongly masked. Apparently, observers could continue to work on checking the T1 object against the memory set while processing an ongoing stream of stimuli or were able to process this task offline after focusing on the T2 task. The memory set size did impact AB when all of the items in the stream were objects. By the way, in the spirit of this chapter, it is worth noting that something that looks quite a lot like an attentional blink can be found in SS77 in experiments with two targets (Schneider & Shiffrin, 1977, p46).

Boettcher, Drew, & Wolfe (2013) have found that observers can hold two memory sets in mind at the same time in hybrid search. In effect, observers can "partition" ALTM. The real world analog might be to have a shopping list that consisted of "tomatoes, celery, apples, milk, eggs, and butter". In a real shopping expedition, only the first three would be relevant in the produce section of the grocery store and only the last three would be relevant in the diary aisle. We wondered if observers could restrict memory search to a relevant portion of the memory set in one situation while keeping the other portion of the set available if the situation changed. In one set of experiments (Boettcher, Drew, & Wolfe, 2013), observers learned a 16-item memory set, split into two, eight-item partitions. The two sets were each associated with a different background (e.g. forest and beach). Observers were only responsible for reporting "forest" targets if they were present on a forest background. Although there was a small (100–200 msec) cost in RT associated with switching contexts, it is doubtful this cost would be problematic as you negotiate the transition from produce to dairy sections in the supermarket. Observers did respect the division of the list into two sections. On some trials, a lure item from the "beach" partition appeared on the forest background. That item, while it was in the memory list, would not be a target on that context (a misplaced gallon of milk in the produce aisle). Eye movement and RT data show some effect of the lure. It was noticed, but observers almost never produced a false alarm. Thus, both partitions must be stored in ALTM (or wherever these memory sets reside) along with rules that permit the correct partition to govern behavior at any given moment. This hybrid partition method could be useful as a way to study the ability to control competing behaviors.

In sum, hybrid search is a rich environment for examining the interactions of vision and memory. We are impressed by how many of those riches were uncovered by Richard Shiffrin and his colleagues long before we had considered these issues. We hope that our work with much larger memory sets and with a range of different stimulus materials will lead to a better understanding of how these two fundamental aspects of cognition interact with one another, and inevitably, it will add to the more than 10,000 citations recorded in Google Scholar for the two SS77 papers as of this writing.

Acknowledgements

We acknowledge support from NIH-EY010017, the National Geospatial Agency (HM0177–13–1–0001_P00001), and Google, as well as a NIH post-doctoral fellowship to Trafton Drew (NIH-NIBIB 1F32EB011959). We thank Rob Nosofsky for useful discussions.

REFERENCES

Baddeley, A. D., & Hitch, G. (1974). Working memory. In G.H. Bower (Ed.), *The psychology of learning and motivation*. (Vol. 8, pp. 47–89). New York: Academic Press.

Boettcher, S.E.P., Drew, T., & Wolfe, J.M. (2013). Hybrid search in context: How to search for vegetables in the produce section and cereal in the cereal aisle. *Visual Cognition, 21*(6), 678–682. doi: 10.1080/13506285.2013.844959

Boettcher, S.E.P., & Wolfe, J.M. (in submission). Searching for the right word: hybrid visual and memory search for words.

Brady, T.F., Konkle, T., Alvarez, G.A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proc Natl Acad Sci USA, 105*(38), 14325–14329.

Briggs, G.E., & Johnsen, A. M. (1973). On the nature of central processing in choice reactions. *Memory and Cognition, 1*(1), 91–100. doi: 10.3758/BF03198076

Brown, S.D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology, 57*(3), 153–178.

Burrows, D., & Okada, R. (1975). Memory retrieval from long and short lists. *Science, 188*(4192), 1031–1033.

Cowan, N. (1995). *Attention and Memory: An Integrated Framework*. New York: Oxford University Press.

Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavorial and Brain Sciences, 24*(1), 87–114; discussion 114–185.

Cunningham, C.A., & Wolfe, J.M. (2012). Lions or tigers or bears: Oh my! Hybrid visual and memory search for categorical targets. *Visual Cognition, 20*(9), 1024–1027. doi: http://dx.doi.org/10.1080/13506285.2012.726455

Cunningham, C.A., & Wolfe, J.M. (2014). The role of object categories in hybrid visual and memory search. *Journal of Experimental Psychology: General, 143*(4), 1585–1599.

Czerwinski, M., Lightfoot, N., & Shiffrin, R.M. (1992). Automatization and training in visual search. *American Journal of Psychology, 105*(2), 271–315.

DiCarlo, J.J., & Cox, D.D. (2007). Untangling invariant object recognition. *Trends in Cognitive Science, 11*(8), 333–341.

Donkin, C., Brown, S., Heathcote, A., & Wagenmakers, E.-J. (2011). Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? *Psychonomic Bulletin & Review, 18*(1), 61–69. doi: 10.3758/s13423–010–0022–4

Downing, P., & Dodds, C. (2004). Competition in visual working memory for control of search. *Visual Cognition, 11*(6), 689–703.
Drew, T., Boettcher, S.E.P., & Wolfe, J.M. (2013, November). *Hybrid visual and memory search remains efficient when visual working memory is full*. Paper presented at the Annual Meeting of the Psychonomic Society, Toronto (Abstract 207).
Drew, T., Sherman, A.M., Boettcher, S.E.P., & Wolfe, J.M. (2014). Memory search for the first target modulates the magnitude of the attentional blink. *Memory & Cognition, 42*(8), 1333–1344.
Dux, P.E., & Marois, R. (2009). The attentional blink: a review of data and theory. *Attention, Perception, & Psychophysics, 71*(8), 1683–1700. doi: 71/8/1683 [pii] 10.3758/APP.71.8.1683
Egeth, H.E., Virzi, R.A., & Garbart, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception and Performance, 10*, 32–39.
Gillund, G., & Shiffrin, R.M. (1984). A retrieval model for both recognition and recall. *Psychological Reviews, 91*(1), 1–67.
Han, S.-H., & Kim, M.-S. (2004). Visual search does not remain efficient when executive working memory is working. *Psychological Science, 15*(9), 623–628.
Hick, W.E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology, 4*, 11–26.
Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology, 45*(3), 188–196.
Jolicoeur, P., Gluck, M.A., & Kosslyn, S.M. (1984). Pictures and names: making the connection. *Cognitive Psychology, 16*(2), 243–275.
Kirchner, H., & Thorpe, S.J. (2006). Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vision Research, 46*(11), 1762–1776.
Konkle, T., Brady, T., Alvarez, G.A., & Oliva, A. (2010). Scene memory is more detailed than you think: the role of categories in visual long-term memory. *Psychological Science, 21*, 1551–1556. doi: doi:10.1177/0956797610385359
Kristofferson, M. (1972). When item recognition and visual search functions are similar. *Perception & Psychophysics, 12*(4), 379–384. doi: 10.3758/bf03207225
Leite, F.P., & Ratcliff, R. (2010). Modeling reaction time and accuracy of multiple-alternative decisions. *Attention, Perception, & Psychophysics, 72*(1), 246–273.
Logan, G.D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review, 121*(1), 66–95.
Luck, S.J., & Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390*(20 Nov), 279–281.
Nako, R., Wu, R., & Eimer, M. (2014). Rapid guidance of visual search by object categories. *Journal of Experimental Psychology: Human Perception and Performance, 40*(1), 50–60.
Nelson, A.B., & Shiffrin, R.M. (2013). The co-evolution of knowledge and event memory. *Psychological Review, 120*(2), 356–394. doi: 10.1037/a0032020 2013–07503–001 [pii]
Nosofsky, R.M., Cox, G.E., Cao, R., & Shiffrin, R.M. (2014). An exemplar-familiarity model predicts short-term and long-term probe recognition across diverse forms of memory search. *Journal of Experimental Psychology : Learning, Memory, and Cognition*. Published online April 2014.
Nosofsky, R.M., Little, D.R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review, 118*(2), 280–315. doi: 2011–04120–001 [pii] 10.1037/a0022494

Oberauer, K., & Hein, L. (2012). Attention to information in working memory. *Current Directions in Psychological Science, 21*(3), 164–169.
Olivers, C.N., Peters, J., Houtkamp, R., & Roelfsema, P.R. (2011). Different states in visual working memory: when it guides attention and when it does not. *Trends in Cognitive Science, 15*(7), 327–334. doi: S1364–6613(11)00085–4 [pii] 10.1016/j.tics.2011.05.004
Raaijmakers, J.G.W., & Shiffrin, R.M. (1981). Search of associative memory. *Psychological Review*, *88*, 93–134. doi:10.1037/0033–295X.88.2.93
Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*(2), 59–108.
Raymond, J.E., Shapiro, K.L., & Arnell, K.M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18*(3), 849–860.
Rosch, E.H. (1973). Natural categories. *Cognitive Psychology, 4*, 328–350.
Schneider, D.W., & Anderson, J.R. (2011). A memory-based model of Hick's law. *Cognitive Psychology, 62*(3), 193–222. doi: 10.1016/j.cogpsych.2010.11.001
Schneider, W., & Shiffrin, R.M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84*, 1–66.
Shepard, R.N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior,* 6, 156–163.
Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84*, 127–190.
Shiffrin, R.M., & Steyvers, M. (1997). A model for recognition memory: REM-retrieving effectively from memory. *Psychonomic Bulletin and Review, 4*(2), 145–166. doi: 10.3758/BF03209391
Soto, D., Humphreys, G.W., & Heinke, D. (2006). Working memory can guide pop-out search. *Vision Research, 46*, 1010–1018.
Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology, 25*, 207–222.
Standing, L., Conezio, J., & Haber, R.N. (1970). Perception and memory for pictures: Single trial learning of 2500 visual stimuli. *Psychonomic Science, 19*, 73–74.
Vickery, T.J., King, L.-W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision, 5*(1), 81–92.
Wolfe, J.M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Science, 7*(2), 70–76.
Wolfe, J.M. (2012). Saved by a log: How do humans perform hybrid visual and memory search? *Psychological Science, 23*(7), 698–703. doi: doi:10.1177/0956797612443968
Wolfe, J.M., Alvarez, G.A., Rosenholtz, R., Kuzmova, Y.I., & Sherman, A. M. (2011). Visual search for arbitrary objects in real scenes. *Attention, Perception, & Psychophysics, 73*(6), 1650–1671. doi: 10.3758/s13414–011–0153–3
Wolfe, J.M., & Horowitz, T.S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience, 5*(6), 495–501.
Woodman, G.F., Vogel, E.K., & Luck, S.J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science, 12*(3), 219–224.