Seminar: Behavioral and Cognitive Neuroscience Seminar (603)

Meeting Date: December 6, 2001

Chairs: Herbert S. Terrace, Peter Balsam, Jon Horvitz, Yaakov Stern

Speaker: Mike Kahana

Topic: Associative Properties and Episodic Memory

Participants:
Emily Stern, Department of Psychology, Columbia University
Chris Summerfield. Department of Psychology, Columbia University
Eduardo Mercado, Rutgers University
Herman Buschke, AECOM
Robert L. Thompson, Hunter College, CUNY
Anja Soldan, Department of Psychology, Columbia University
Maggi Fizelli, Padova, Italy
Brigid Finn, Department of Psychology, Columbia University
Jennifer Mangels, Department of Psychology, Columbia University
Jim Magnuson, Department of Psychology, Columbia University
Tammy Moscrip, Department of Psychology, Columbia University
Dana Heitmeyer, Columbia College
Vince Brown, Hofstra Universirty
Hadassah Paul, Hofstra University
Chris Habeck, Sergievsky Center, Columbia University
Elise Caccappalo, Sergievsky Center, Columbia University
Brady Butterfield, Department of Psychology, Columbia University
Jacqui Rick, Department of Psychology, Columbia University
Dustin Merritt, Department of Psychology, Columbia University
Johannes Schwaninger, Department of Psychology, Columbia University
Janet Metcalfe, Department of Psychology, Columbia University
Herbert S. Terrace, Department of Psychology, Columbia University
Jon Horvitz, Department of Psychology, Columbia University
Peter Balsam, Department of Psychology, Barnard College
Yaakov Stern, Sergievsky Center, Columbia University

Rapporteur: Kate Lynch

Summary:

The main objective of Dr. Kahana's talk was to introduce a new quantitative approach to modeling Free Recall called the Temporal Context Model. In the first part of his talk, Dr. Kahana gave a brief overview of the Free Recall task and discussed phenomena found using the task. He discussed how associations can be measured using contiguity effects and how Free Recall can be decomposed into two underlying processes described in the classic modal model of episodic memory, short term (STM) and long term memory (LTM). He then introduced findings from his laboratory at Brandeis that show the timescaled invariance of association in Free Recall, which will be described in more detail below. In the second part of his talk, Dr. Kahana introduced his temporal context model. In the last part of the talk, he briefly discussed free recall in normal healthy older adults and the applicability of the temporal context model to Serial Recall Tasks.

While free recall has simple and straightforward instructions, the results can be complex. Free recall data can be analyzed for effects of input position, output order, transition probabilities and transition latencies, where the subjects begin the recall process, and how they proceed. Dr. Kahana discussed some classic phenomena examined with the free recall task, recency effects, primacy effects and serial position effects. The recency effect refers to the phenomena that recently presented items are recalled first and with the greatest probability because recent items are still in the short-term memory store. The primacy effect is a smaller effect than the recency effect and refers to the advantage of the first couple of items in the list, due to rehearsal and storage in long-term memory.

Dr. Kahana's research focuses on the flat/asymptotic portion of serial position curve where items are recalled with equal probability. Conditional response probability (CRP) functions are used to quantify the degree to which temporal proximity of input influences the probability of a successive transition in output. Conditional Response Probabilities functions are also referred to as lag recency functions because what is being measured is the recency from the just recalled item. Eleven studies from different laboratories have used CRP functions to show that adjacent items (lag 1) at input are recalled successively during recall. The rate of association falls off continuously as the input lag increases. There is an asymmetry in the response probability function, with the tendency to look forward in the list being greater than the tendency to look backward. For example, if input item 5 is recalled, item 6 is most likely to be recalled next, and item 4 is slightly less likely. In normal healthy subjects the probability of recalling an item again after it was just recalled (zero lag) is zero. While conditional response probabilities have been used extensively in previous research, the application to free recall tasks is new.

The latest mathematical incarnation of the classic two-store model of human memory based on the division between STM and LTM is the SAM model. This model captures the associative effect found using conditional response probabilities. Items that are held in the short-term rehearsal buffer are associated with one another. This association explains the successive retrieval of items that are adjacent in a study list. This model postulates ad hoc that forward associations are twice as strong as backward associations, in order to explain the asymmetry effect.

The effects found with latency of response at retrieval are similar to those found with CRP. Five studies from different laboratories have shown that graphing response latency by input lag produces a function that is the mirror image of CRP functions. The fastest response latencies occur for words that were in the forward adjacent input position to the item just recalled. In summary, nearby items at input are recalled together and more quickly in free recall. In response to a question posed by Dr. Terrace, Kahana says that CRP functions can be analyzed for lag positions beyond +/- 4; however, very long study lists are necessary to do this.

CRP functions provide a way to look at transitions in free recall. However, another aspect of free recall that is not addressed using CRP functions is the way in which retrieval of the first item is initiated. Kahana used a fast presentation rate to discourage rehearsal and thus minimize the primacy effect; however, a large recency effect is evident in the serial position curve. Kahana has found that the last item in a study list has a .5 probability of first recall, and the second to last item has a .25 probability. The large probability of first recall for the items at the end of the list is exactly what creates the recency effect. If the last items were not recalled first, this effect would not be found. After about the third transition, the function flattens out and so data from the first three transitions are thrown out.

Two pieces of data prove fatal for classic modal models of memory: scale invariances of recency and lag recency (or of association). In a delayed free recall task, in which subjects are distracted at the end of the list for 10-15 seconds prior to recall, end of list items are recalled with a much lower probability. However, if the same delay (i.e.15s of arithmetic) is given not only at the end of the list, but between each item in list, the recency effect returns. The CRP functions are unchanged as long as the relative separations are kept constant.

Dr. Kahana presented the analogy of the list of items as points in space. If all the items are stretched out in space, the tendency to recall the most recent items first is unchanged. However, if a distractor is given at the end of the list, but the items in the list are not spaced, the tendency to begin at the end of the list is significantly reduced.

Mark Howard and Kahana discovered the same invariance for the lag recency effect. The tendency to report items that co-occur in the list in adjacent output positions is unchanged by interpolating an arithmetic distractor between them. This result poses a problem for the classic model that suggests that associations resulting in the lag recency effect are formed in the short-term store. The same distractor that should be pushing list items out of the short-term store doesn’t diminish the tendency for temporally related items in a list to be recalled in adjacent output positions. Thus, it can be concluded that recency effects depend on relative timing, not absolute timing. STM is critically dependent on the idea of a time bound store. Without the assumption of a temporally limited store, the distinction between STM and LTM is no longer justifiable.

Dr. Kahana concluded part I of his talk by summarizing that in his analysis of whether and when items are reported in a free recall task, he has discovered two scale invariances of recency and lag recency. These invariances post a problem for the classic modal model of human memory, because it assumes a division between STM and LTM.

Dr. Terrace asked whether any studies have been done that show similar results in children?
Dr. Kahana replied that he is not aware of any similar studies with children.
Dr. Terrace said that he suspects the assignment of items to ordinal bins must be learned.

Dr. Kahana followed up that he suspects similar result would be found in children. His statement is based on the fact that similar effects were found in studies done in Dr. Terrace's lab with rhesus macaques.

In Part II of his talk, Dr. Kahana discussed his Temporal Context Model. In this model each item is stored as a distributed representation in the F or item layer. The item then becomes associated or linked with a representation of temporal or spatial context. Each word becomes associated with a place in space, as we move through space or time. The most recent items are located most closely in space or time. The model suggests that space/time must be coded, associated with items, and then retrieved. Items, or words, are assumed to have a rich multivariate, distributed representation and can be thought of as a point in space as in Minerva, Charm and other classic memory models. The component to the model is that we are forming a distributed representation of time. But where does this representation of time come from?

Dr. Kahana assumes that there exists a distributed representation of time, with the property of drift as we move through space and learn information and that the distributed representations of items are linked to distributed representations of time. Likewise, there is the distributed representations of time are linked back to the distributed representations of items. The linkage, which can be conceptualized as an association, is formally coded in the memory model using outer product matrices, or Hebbian learning rules.

The key to the model is the mechanism for the evolution of context or the evolution of this time signal. Dr. Kahana uses a simple drift process to explain this evolution.
The equation describing this process begins with an initial representation of time and as temporal evolution occurs, the temporal context is represented as a reduced weight of its previous state, added to the new input to the system. As new information is retrieved, the just retrieved information, rather than noise, pushes the representation of time along. The retrieved information after an initial encoding is the input that leads to the temporal evolution of context.

Dr. Kahana provided a concrete example to explain this idea. Consider the following list of words: *absence, hollow, pupil*. When the word *absence* is encountered, things that are associated with absence prior to the list (pre-experimental context) are retrieved. Because the words in the list are selected randomly, the retrieved pre-experimental information associated with *absence* is assumed to be uncorrelated with the other words in the list, *pupil* and *hollow*. The ordering of the items produces a drift and the context signal will always move away from its previous position.

Dr. Kahana explained how recency effects can be produced by retrieved context.
When a context signal is retrieved, items that are associated with that context, in memory, are retrieved. The most recent items will be the most similar in memory, producing the recency effect. However, there are two components to the retrieved context, the list context for the just retrieved item, and the pre-experimental context for that item. For example, given the study list *absence, hollow pupil*, when *absence* is retrieved, both the context of *absence* in the list, and the pre-experimental context for the word *absence*, will

be retrieved. Though the model assumes that the two components of context are retrieved in equal proportion to one another, these two components of context have different effects.

Retrieved context from the list is equally similar to nearby items originally presented both before and after. The pre-experimental context, on the other hand, is only in the list for the items that came after the studied item. This is because the pre-experimental context for an item must necessarily occur in the list only after that item is presented and thus this pre-experimental context will only be similar to those items occurring afterwards. Suppose a subject studies the word *house*, and then retrieves the pre-experimental context associated with the item *house*. Suppose the next word in the list is *absence,* and then the pre-experimental context for the item *absence* is retrieved. Then the word *hollow* is presented and so on. At test, when *absence* is retrieved, the pre-experimental context associated with absence will be retrieved as well. This pre-experimental context will be associated in memory with *hollow*, but not *house*. Thus hollow will be retrieved next.

In fitting the model to the data from the immediate free recall task, it is evident that the experimental and model induced recency and lag recency effects are identical.
In delayed free recall, both the model and experiment show reduction in the recency effect. However, in the continuous distractor free recall task, the model and the data show a return of the recency effect. Thus, the model is capturing both the scale invariance of the recency effect and the asymmetry in the lag recency effect.

The asymmetry results from the forward bias occurring with the retrieval of pre-experimental context. The list context produces a symmetrical effect. The model predicts an approximate 2:1 ratio, of the forward going to the backward going dissociation, which is exactly what is found in the continuous distractor task. In part two of this talk, it was shown that the Temporal Context Model is able to explain both recency and contiguity effects, the asymmetry in contiguity (or lag recency) and the tendency for both of these effects to be invariant with the overall relative spacing of the items.

Dr. Terrace asked whether the model makes any specific prediction in regards to the shift from the end of list items to the items occurring earlier in the list, during free recall.
Dr. Kahana responded that the primacy effects described by Dr. Terrace only occur when the early items in the list are well rehearsed. The distractor task used in Dr. Kahana’s studies is meant to reduce the complications introduced when rehearsal occurs. However, Dr. Kahana noted that even without rehearsal it is difficult to eliminate the primacy effect for the first item in the list. This effect can be explained by assuming that there is a big jump forward as the list context begins and thus interference of pre-experimental information on the first item is reduced.

In part three of his talk, Dr. Kahana discussed the effects of normal aging on free recall. Older subjects have more difficulty than young subjects on serial position tasks. However, although the older subjects remember fewer items than the young subjects in general, when a delay is imposed at the end of the list, both older and young subjects show attenuation of the recency effect.

Dr. Kahana discussed performance of older and younger subjects on free recall tasks when initiation (recency) and transition (lag recency) processes are considered separately. The recency data are identical in the older and younger subjects, even after a delay. Transitions were analyzed to investigate differences in the lag recency effects (involving associative processes) of younger and older subjects. Older subjects show the same asymmetric CRP function found in younger subjects, with forward item recall better than backward item recall. However, these CPR functions are not as peaked in older subjects as they are in younger subjects, and thus there is a decrease in the associative tendencies in older adults compared to young. The Temporal Context Model predicts the data quite well if you make the assumption that older subjects have trouble retrieving the temporal context that went along with the items in the list.

Older subjects have no trouble using the temporal context that is currently available in memory and thus have no trouble recalling items that are recent in the list. However older subjects have a hard time jumping back in time and retrieving the context that was associated with the recalled item, which produces retrieval noise. Similar evidence of dissociations between end of list and within list phenomena has been used to argue in support for the modal model and the distinction between STM and LTM.

Dr. Stern asked for clarification of whether older subjects have difficulty encoding the context initially or recalling the context.
Dr. Kahana responded that older subjects have no difficulty with the initial coding, but have difficulty with the retrieval of context.

Lastly, Kahana discussed the Temporal Context Model in the context of serial recall.
In serial recall tasks, subjects are asked to free recall words that were presented in a series of list. Subjects were instructed to encode the order of the words in the lists (serial encoding task).that were encoded for order. Errors were analyzed to determine whether forward errors occur more frequently than backward errors, as would be predicted by a lag recency effects found in the free recall task discussed earlier.

In probe serial recall tasks, Dr. Kahana found identical functions with similar asymmetries to those found in free recall, and serial recall tasks. When subjects are asked to indicate which word came after a particular item in a list (i.e. house), the same asymmetry is found regardless of whether a forward of backward was given. These effects from serial, free and probed recall are extremely robust, but do not hold for paired associate recall, which produces symmetrical functions.

In conclusion, Dr. Kahana summarized several key ideas from the Temporal Context Model. The evolution of context is not driven by random noise but by information retrieved as each item I the list is encoded. The model assumes that temporal context is retrieved both when an item is studied and when it is recalled and that this retrieved context pushes the list context along. The effects found with free recall are explained by assuming that the temporal context that is present at the end of the list leads the subject to sample items to initiate recall. With the retrieval of those items context is recovered which allows recall transitions to be made. The two mechanisms of initiation and transition explain the problem of scale invariance in recency and contiguity. These mechanisms also explain the dissociation between recency and associative phenomena, which had previously been thought to be pillars on which the classic modal dual store model of memory was based.

Dr. Kahana made one last comment on the evolution equation presented above. In the example provided above, retrieved context is the input that moves the list context forward. However, the input can be something other than retrieved context. For example, with animals the input could be a velocity vector as the animal moves through a spatial field. If the input is the velocity information as the animal moves around, than a representation of space will result.

Questions

Dr. Stern asked how the shift from the most recent items at the end of the list, to items within the list occurs?

Dr. Kahana responded that the CRP functions described above explain only part of the variance. Other factors, such as semantic relatedness, contribute as well. The metric of semantic space can be used to build new CRP functions in which semantic and temporal similarity hold equal weight. The incorporation of semantic relatedness allows one to explain the jumps in recall to related items. These jumps will be totally random and thus produce a flat distribution.

Dr. Paul asked what would happen if the distractor interval between items in the free recall task was not fixed?

Dr. Kahana responded that this will cause a shift in the CRP function. If the distractors get longer throughout the task, function will be biased backwards. Shortening the distractor intervals throughout the task, will shift the function in the other direction. Varying the distractor intervals shifts the associative effects and the model predicts those shifts

Dr. Paul asked Dr. Kahana to comment on the effects of similarity of distractor task used to prevent rehearsal and the target task.

Dr. Kahana responded that a distractor task other than the arithmetic task could be used. For example one could use categorized word lists composed of animals, flowers and gemstones. Then, subjects would be asked to recall all animals, making the other items to be remembered the distractor between two of the animal items. Dr. Kahana does not know whether the time scale invariance found in his studies depends on the use of an unrelated distractor task, but thinks that this is an important question.

Dr. Metcalfe described some old data from a serial recall study. It was found that when people made an error, the usually reported an item that was adjacent to the target item, as would be predicted by the transition probabilities. However, occasionally an error would be made that wasn't an adjacent error, but rather an item from the same serial position in another list. Two items in the same serial position from different lists should not be semantically related. Dr. Metcalfe asked how Dr. Kahana's model, which is entirely based on semantic context, would explain this data?

Dr. Kahana agreed that his model does not predict this finding. Dr. Kahana reiterated that what this question is challenging him to do is to take the temporal context model of free recall and apply it to serial recall. The current version of the temporal context model will not work very well with serial recall data, even though serial recall has output gradients similar to those from free recall tasks. The temporal context model requires that you sample items from memory based on temporal proximity. Serial recall dictates that you must recall an item in from a particular serial position. Serial recall tasks, perhaps train, subjects to rely on a more ordinal representation. Dr. Kahana suggested that it is possible to imagine some version of the temporal drift idea that could code time directly rather than relying on randomly retrieved semantic information.

Dr. Metcalfe followed up by suggesting that the analogy of items represented in space or time, presented by Dr. Kahana, seems to almost provide a concrete representation of the coding of either space or time. Is it the case that you have a concrete representation of this sort?

Dr. Kahana responded that while you have time or space coded for the current list, you do not have it for the prior lists. This poses a challenge to the model. Dr. Kahana suggested that the next step in this program of research is to determine how this kind of approach could be extended to the retrieval of errors from previous lists.

Dr. Stern asked whether Dr. Kahana think s that the lessening of the temporal contiguity effect in aging is really what explains the result that older subjects recall fewer words than younger adults?

Dr. Kahana said that the lessening of the temporal contiguity effect can account for the flattening of the gradients in output from free recall. Dr. Kahana suggested that the reason older subjects recall fewer items is perhaps that older subjects sample more diffusely in time. Another possible explanation involves the point at which subjects stop attempting to recall more items. A subject might stop attempting recall if an item is recalled incorrectly, and the subject is able to use post recognition processes to determine that the item was not actually in the list. Experiments were conducted in Dr. Kahana's lab that attempted to look at intrusions in free recall. However, both older and younger subjects hardly made any intrusions. A more recent study in his lab is using an overt recall procedure in which subjects are told to recall everything they think of and then press a button if they think a recalled item was actually not in the list. Since not all sampled items are actually recalled, there must be some mechanism that occurs in the transition from sampling to actual recall. Dr. Kahana suggests post hoc, that this mechanism might be used to help the subject decide when to terminate recall. If several items are sampled, but determined not to be in the list, this could be a signal to the subject to stop attempting recall. If subjects free associate based on temporal and semantic context, they would never get anywhere. In order to understand the deficit in elderly, the mechanism that translates from the sampling rule to the stopping rule must be specified.

Dr. Terrace asked whether Dr. Kahana has applied this model to serial learning per se?

Dr. Kahana responded that he has not yet applied the model to serial learning. However, as discussed earlier in the talk, he has shown that certain key effects found in free recall, that were used to build the model, do hold for serial learning. Dr. Kahana knows of four models of serial recall that have been proposed over the past two years. Each of these four models indicates that subjects associate the item in the list with ordinal position. Thus, these models would predict the prior list intrusion effect. However, these models have only been verified with short lists. Results of serial recall tasks using longer lists show a deviation from the predications of the ordinal code based models.