Reflections on Context and Episodic Memory

Gordon H. Bower, PhD.
Stanford Professor, Emeritus
(www.Stanford.edu/~gbower)

Keynote Address
”Context and Episodic Memory Symposium”
University of Pennsylvania
May 8, 2014

Thank you, Ken, for that introduction, and I thank this society for inviting me to speak on a topic that has interested me for a long time, namely, the role of context in episodic memory. In their invitation, Mike Kahana and Ken Norman suggested that I review some of my early memory research from the 1960s & ‘70s based on the ideas of contextual associations. I am honored and happy to do so. I hope my review will connect with the exciting research that your group is doing these days.

Here is the way I’ll organize my talk. (#1-OUTLINE SLIDE). This is also printed on page 7 of your Program. After a brief review of the kinds of contexts people in lab experiments remember, I’ll introduce a particular perspective I call the paired-associate view of context associations, and proceed to describe its applications. As I go along, I’ll indicate where I am in this organization so you can follow along. (SLIDE OFF)

I first encountered the importance of context in performance during my first year in graduate school at Yale. I was training rats to get a food reward in a T-maze by choosing a white alley over a black alley. The T-maze had a wooden floor and sides, and over many days of running my rats through that maze, the wood floor got pretty dirty, messy, and smelly. I decided I would wash out the maze so that my rats could perform in a clean well-lighted place. However, when I put my rats back into that sparkling-clean maze, they froze up and refused to run! They reacted as if they were saying, "What's this? Where's all the crap that I'm used to seeing and smelling in this alley? I'm not going to venture out into this strange place where I can't smell all the pee I left behind when I was last here. "

That experience taught me that the stimulus to which the animal was responding is not just the black vs. white maze arms that I had been focused on. Rather, it included a lot of extra stimuli ---- extraneous sights, smells, and sounds --- beyond what I thought were the important elements of the rats' environment.

So I was well prepared when I moved into studies of human learning and began to appreciate the role of environmental context in human memory --- the idea that people are recording in memory something about the on-going milieu in which events are taking place.

In those days we were studying event memory by using a very simple task. We'd have college sophomores learning lists of unrelated words presented one at a time. We can consider each word presentation as an event to be remembered. After that list exposure, people's memory for the words would be tested in some manner. I presume you are all familiar with the routine.

By 1971, work by my students, Doug Hintzman and John Anderson, and I had uncovered the fact that college students in these list-learning experiments can remember at least a little bit about many things regarding the context in which verbal items occurred. They can remember the physical properties of the presentations, their modality, their sources or locations, surroundings, and how they had dealt with the item during study. Hintzman had also discovered that people could remember in which of several lists an item had occurred as well as their frequency and locations within a list. These judgments, often gathered from incidental memory, are somewhat inaccurate, of course. Nonetheless, the sheer amount of material that adults could remember about the surroundings of an item's presentation is impressive.

As a first pass in dealing with some of this complexity, John Anderson and I supposed that upon presentation of a known verbal item in a learning experiment, call it Word A, the subject's internal representation of that word (in long-term memory) would become associated with a "context marker". (#2-SLIDE MARKER) The marker, denoted M1 here, was conceived as a new node established in memory and pointing to a collection of descriptors of the external stimuli and internal thoughts surrounding that item's presentation on that occasion. These descriptors, denoted E1, and E2 etc, here, would include not only the item's physical appearance and its position within an array, but also their thoughts about the item, and whatever internal elements might permit "time tagging" of the presentation. Prominent among these would be a special element marking the collections of items grouped in time, such as First List, or List 1 and List 2.

This S**LIDE** diagrams the internal representation in memory of three prototypical words that occurred in three different lists. Each word in the list would have such markers associated to it. To elaborate, presentation of word A has caused it to become associated to one new marker, M1, recording the fact that it occurred along with a swarm of contextual elements arising during presentation of List 1. Later presentation of Word B has caused it to become associated with two markers, M2 and M3, recording the facts that it occurred among two swarms of elements, some identifying List 1, others coming along later and identifying List 2. Word C has been associated to two markers recording the fact that it occurred at two different positions within List 3.

I'll take a minute to insert a necessary comment here. I know that many of you are neuroscientists who are looking for brain structures that might instantiate these notions of internal representations of items, contexts, and associations. I sympathize and say, more power to you. I suppose that there is an isomorphism or correspondence between the node diagrams I'll be using and the brain diagrams you might use. For example, here is a schematic brain diagram **(#3-SLIDE** brain diagram**)** from a paper by Sean Polyn and Mike Kahana. They supposed that an input item like a word would be represented by a neural pattern of activation in a sensory area such as the word-form area of the temporal contex; furthermore, the ongoing context would be represented by a neural pattern in the prefrontal cortex; and they supposed that the binding of the two would occur when the two patterns arrive together in the hippocampus and interact to form an episodic association. That association would essentially encode the fact that this item occurred in this particular context. Polyn and Kahana also supposed that the context would be updated as new items are presented to the person, and that the context via these contextual associations has an important role to play in guiding retrieval of memories. I say Amen to all of that.

Now, all of these processes can be represented in my diagrams of nodes with associative connections. It's just that I use a single node rather than a sparse-coded network to represent a word. With that bow to the neuroscience of context memory, I will proceed with my old notation.

I'll now move to the Program Outline Section 3, on implications of this view.

(#4-MARKER SLIDE again) Given my paired-associate framework for associating words to context markers, let's consider a few implications. First, the approach can deal with simple recognition memory, wherein the person judges whether or not an item like word A occurred at all during presentation of, say, List 1. The model essentially looks for an associative connection between the internal representation of word A, and those aspects of the list-context being asked about, such as LIST 1. One can elaborate a probabilistic signal detection model for such Old-New judgments. I will spare you the details here.

Second, the approach is set up to deal with list or source discrimination, for people remembering in which of several lists an item had occurred or which source, say, male or female speaker, had spoken the item. We see this for WORD B in this slide. You can see how source identifications would improve with stronger learning or repetitions of a word from a given source.

These source judgments can be in error, of course, due to poor learning or simple forgetting over time. Moreover, judgments will surely be contaminated with inter-list confusions. Because the list markers refer to memory for bundles of stimulus elements, the context-elements for two different lists will overlap somewhat in content. We can schematize those shared elements with Venn diagrams like this (#5- 2 VENN SLIDE). The greater the overlap of shared elements, the greater the similarity of the two contexts, and so the more often subjects will confuse the sources where they saw or heard an item.

Let's consider the influence of those list-similarities in learning, say, three successive lists of words distinguished only by their order of presentation in time. The overlap of contextual elements for three successive lists might be illustrated as in this SLIDE (#6- 3 sets). To reflect this change in contexts over time, this diagram shows that the context elements associated with List 1 overlaps a proportion omega, w, with those elements identifying List 2, and List 2 in turn overlaps a proportion omega with List 3. The consequence of those context-similarities of these 3 lists is that subjects will later make the most generalization errors in identifying items from the middle-list (#7-SLIDE of 3 lines). That is, the middle list will have the poorest accuracy score. This, of course, is what Hintzman and others have routinely found. (SLIDE OFF)

So, referring again to your Program Outline, I'm about to talk about Item 3B, contextual interference.

John Anderson and I reasoned that if item-to-context associations act like standard paired-associates, then we should be able to create a lot of negative transfer and interference when we have subjects try to set up associations of a given word to several different contexts in which it occurs. These negative effects are, after all, what one routinely finds in standard interference paradigms in which we associate the same cue to different responses.

Not being content with studying discrimination between just two- or three lists, Anderson and I tried to bend our subject's heads around the task of keeping track of 4 or 5 lists in which different words occurred or not. Here's a simplified illustration (#8 +/- SLIDE) for a 3-list experiment. Here, unrelated nouns (represented by rows) were presented or not in each list. With 3 lists, you can have 8 patterns of appearances (that is, the pluses) and nonappearances (minuses); with 4 lists, you can have 16 patterns; with 5 lists, you can have 32 patterns of appearances or non-appearances. You can assign as many distinct words (denoted as n here) as you like to exemplify each row. In our experiment, we would announce the name of the list before showing its randomized words 1-by-1 on distinctive backgrounds for subjects to study. At the end of studying the 4 lists, subjects rated their confidence that a given word had appeared or not in each of the 4 separate lists.

So, what did we find? First, we found a ton of generalization errors among adjacent lists. This SLIDE shows (#9 -GENERIZATION SLIDE) data from the two experiments, one with 4, one with 5 different lists. It brings out the graded generalization of list-identification errors among adjacent lists. For example, items appearing in List 3 were most likely to be falsely identified as having appeared in Lists 2 or 4. List 2 items were most likely to be misidentified as occurring in List 1 or 3. And so on in an orderly manner. In fact, the more lists in which an item appeared, the more likely subjects were to claim that it had appeared in some other list even it hadn't. Thus, we're seeing a heightened false alarm rate. Those false alarms lower the accuracy of the person's specific list memories. This decreasing identification accuracy, indexed by d-prime, is shown in this slide (#10-d' SLIDE). These list-identification data reveal strong negative transfer. Pairing the same word with more lists created strong interference among different associations of lists to that particular word. This is the same as one would find in simple paired-associate learning, where we pair a given stimulus with the learning of more interfering responses. So that's one consequence of viewing context memories as though they are paired-associates. (SLIDE OFF)

Moving on, the next item on your Program Outline is 4C, context-changes. Beyond interference effects, what about simple passive forgetting of contextual associations? Loss of associations can be attributed to context elements that vary and change over time. Many years ago Bill Estes proposed that contexts might be composed of populations of stimulus elements that would fluctuate or change their availability from moment to moment, with the elements coming and going, appearing, disappearing, and reappearing over time.

This SLIDE (#11-fluctuation) illustrates the general idea: there's a huge pool of possible contextual elements for any situation, only some of which are available for the subject's attention at any one time. I show the available set at time 0, and how at time 1 it is composed of some elements persisting from time 0, some elements that have diffused out and become unavailable, whereas other elements that were previously unavailable have diffused into the available set at time 1. This fluctuation process goes on continuously so that by a much later time X, the set of elements may have changed to a great degree.

I developed this idea a bit more in a 1972 paper, and it's been elaborated further by Mensink & Raaijmakers in 1989, and by Marc Howard and Mike Kahana in a 2002 paper on their Temporal Context model.

If we just focus on two points in time (#12-SLIDE of 2 circles), you can see that the contextual fluctuation process provides a sort of crude interval timer. The more time that passes, the greater will be the change in the composition of the context compared to those elements that prevailed at an earlier time.

We can use this observation to try to understand how people might be estimating how recently a particular item might have occurred in a uniform stream of undistinguished experiences; that is, they're estimating the time interval that has elapsed since they saw it. Consider an experiment like this (#13-SLIDE of item stream)) in which subjects are studying unrelated words presented one at a time at a steady pace in a continuous stream. This is a laboratory analog, say, of customers' orders coming into Amazon's receiving desk in a steady random series, or it might be sports fans arriving in random order to watch a basketball game. In the experiment, we occasionally stop and ask subjects to estimate how far back in the series an earlier item was presented. For example, in this illustration, word P occurred 4-back in the series before the probe question, and word M occurred 2 back from its probe.

To model this, let's remember that the earlier item had been associated to a collection of contextual elements back when it was first presented. Let's assume that those elements are now called up from memory, (#14-SLIDE overlapping sets), and the retrieved elements are compared to the elements of the current, ongoing context. The contextual elements will have drifted over time, slowly altering the composition of the current context. By comparing the elements of the current context with those retrieved from his memory by the target item, the subject can assess the similarity or overlap of elements in these two sets. Let's suppose that he bases his judgment of recency of the target item on the percentage overlap of the two sets.

If we suppose that contextual fluctuation causes the overlap percentage to decrease exponentially with Lag, and that the subject's Judgment of Lag will be inversely related to the overlap percentage, then the subject's Judgment of Lag will approximate a simple power function of the actual Lag. (SLIDE OFF).

We can compare this power-function to data on Lag judgments collected in my lab in 1970 by my student Jim Hinrichs. We were trying to separate people's memory-strength for words in a series from the decision process they were using to make their lag judgments. A simple way of doing that was by asking subjects to use different numbers of judgment categories: although the lags in the actual series subjects saw varied from 1 to 9, three different groups of subjects were told that the lags in the item series they were dealing with were at most either 6, or 9, or 12 --- effectively giving them 6, 9, or 12 judgment categories. Memories are so bad in such experiments that our subjects couldn't tell that we were deceiving them about the maximum lags.

Here are the lag judgments data that Hinrichs collected. (#15-Hinrich's SLIDE). As expected, the more response categories people have, the higher are the judgments they make. Also as expected, judged lags exceeded actual lags at short lags, but they crossover, begin to decline and level out at longer lags where subjects' memories become poor and they are guessing more. The three curves can be fit by a power function with just the multiplying constant increasing with the different instructed maximums of 6, 9, and 12.

This context-similarity model can also compare the recencies of two items from the earlier input stream, as was done in an old experiment by Yntema & Trask back in 1963 --- 51 years ago (#16-Stream SLIDE). Suppose the more recent item was presented N items back in the stream; and the other item to which it is compared was presented K items farther back than that. The exponential-decay model of context changes yields predictions like those shown in the next slide (#17 –Yntema curves). Choice of the more recent item is more accurate the closer it is to the present (that is, for smaller N), and is more accurate the greater is the distance (K) from it to the earlier item in the input stream. These curves closely resemble the comparative recency data reported by Yntema & Trask. The curves here are similar to those expected by a Weber-Fechner law for sensory comparisons, that arise with comparative judgments of any physical stimuli such as two weights, or loudness of two tones, or more to the point, the durations of two time-intervals. (SLIDE OFF)

Let's move on from recency judgments to the next topic 4D in your Program Outline, namely, item frequency judgments. Consider a classic experiment conducted by Doug Hintzman in 1969 on frequency judgments from memory. In his experiment, individual unrelated words appeared either 0, 1, 2, 4, 6, or 10 times scattered randomly throughout a very long list. Think of this as a given word-type laying down in memory a number of tokens of itself. Afterwards, subjects were shown each word type and they estimated how often that word had appeared --- that is, how many tokens of it they had seen in this long list.

Our paired-associate model has some rudiments for dealing with item frequency judgments from memory. Recall this earlier slide (#18-SLIDE), where word C occurred twice in List 3. The multiple list-markers associated to item C would allow the person to more or less count and report from memory that it had occurred twice in List 3. And, maybe that counting process could be extended to remember 3, 4, or more presentations of the item. But surely, however, this frequency-estimating task would quickly become very difficult because as subjects associate more list markers to many different repeating words, the associations would interfere with one another and become confused.

One might be able to develop our marker model by assuming that each presentation of the same item sets up a memory trace, and some of these persist to the point at which the frequency estimate is taken. The model would have to be complicated by interference and blocking among the various traces, so lots of errors in counting would occur. I haven't worked out a model for the process; perhaps one of you can. But the results in Hintzman's experiment are crystal clear (#19-Hintzman SLIDE) Subjects frequency judgments were a simple logarithmic function of the actual frequency of an item's presentations. This graph plots median frequency judgments as a straight line on semi-log graph paper against actual presentation frequencies.

This counting process will surely cease as more traces of an item accumulate and interfere with one another. At that point, subjects will increasingly just guess at some maximum frequency. That kind of guessing would create an over-estimation of low actual frequencies and an under-estimation of higher frequencies, rather like the logarithm function that Hintzman observed. (SLIDE OFF).).

I'll examine one more implication of our paired-associate theory, and that is how two presentations of an older item can cause it to appear more recent than a later, once-presented item. Consider the recency comparison task of Yntema & Trask where subjects judge which of two words has occurred more recently.

Here's the set up (#20-Distribution SLIDE). It's kind of complicated, so I'll describe its parts. The arrow on the right indicates the NOW time of the choice test, and the horizontal axis is the apparent recency or pastness of memories of particular items. This shows distributions for some items that had been presented, let's say, 10 back from the probe test, some 13 back, and some 16 back from the test. Recall, these apparent recency distributions are based on the similarity of their contexts to the NOW context prevailing at the time of the judgment. The 3 are plotted as normal probability distributions to recognize the fact that fluctuation is a statistical process; moreover, we are pooling together the apparent recencies for many different subjects and different items that exemplify this specific arrangement of lags.

Now, item B is the most recent and we're supposing it occurred 10 items before NOW; and the first two rows show that its earlier competitor, A, was presented either close (at 13 back) or a bit farther back (at 16). Moreover, before this final B, we play an underhanded trick by presenting earlier A's or B's in the arrangements shown in the last three rows here. The subject could receive either A then B before the final B, or receive B then A before the final B, or A then another A before the final B. How should these arrangements affect the model's choice of B as the more recent item? Recall, chance responding is 50% choice of item B as most recent.

Well, the first two lines show that, as expected, choice of the more recent word B over a singly presented competitor, A, is more accurate with a greater distance between the ages of the two items --- that is, choice of lag 10 over 16 is greater at 64% than choice of lag 10 versus 13 at 61%. Moving down to the 3rd row, item B should be chosen the most often since it will win out if the recency of either of the two recent B's exceed the apparent recency of the one older A. Indeed, that row yields the highest choice of item B at 77%. Moving to the 4th row, item B will be chosen as more recent if the apparent recency of either B exceeds that for the one A. We see that the earlier B does help a little, raising this choice to 69% compared to the baseline 61% which had just one presentation of A in the competition. Finally, in the last row, the recency of B must exceed the recencies of each of the two A's in order for B to be chosen, so ganging up two A-traces against one-B trace knocks down the B-choices to near chance, 54%, which indicates the most errors of A being chosen over B as most recent.

In order to make predictions, we assumed that the apparent recencies had normal distributions with equal variances. We estimated the distance between adjacent distributions as three-tenths of a standard deviation. That is, the means of 10-back and 13-back differed by 3-tenths of a standard deviation as did the means of 13 -back and 16-back. So with that one parameter, we predicted the percentages in the last column. You can see that the model with one parameter fits the 5 rows of numbers quite well.

So those are the kinds of predictions one can make by assuming that presentation of an item establishes an association of it to the context and moreover, those contexts fluctuate over time.

Now, I'll move on to the next topic, #4, of your Program Outline, namely, implicit memory.

In an article published in 1984 and elaborated in 1990, I tried to account for many of the facts of implicit and explicit memory by highlighting the crucial role played by contextual associations. Suppose that when a person reads a word like HARE (#21-Hare SLIDE), his visual system uses the fragments of the graphemes to identify letters, and those letters-in-position are associated with word nodes in long-term memory. The exact details of this picture are not important just so long as we retain the idea that letters-in-position are associated to words. These links here are all old associations that are well established in long-term memory for English readers.

I will make one assumption, namely, that whenever an old association, like this one from H-in-position-1 to HARE, is reactivated, it will be strengthened and be more readily available for a while.

A simple consequence of this strengthening assumption is simple perceptual priming: that is, the primed person will see HARE more readily in a brief flash or in a degraded visual display, and he'll rapidly judge that HARE is an actual word. He will also complete a fragment like HA— or H - - E more readily with HARE. Moreover, presenting beforehand closely spelled words like HERE, or HIRE or HART will create interference for subjects trying later to see HARE in a brief flash.

By substituting auditory input of spoken words and extraction of phonemic features, as illustrated in this **SLIDE**, (#21-auditory prime) the same strengthening assumptions will produce auditory priming. For example, the model would more readily hear HARE in a masked auditory test, and it'll spell HARE to the sounded word rather than HAIR after reading about rabbits.

Similar considerations can be extended to explain picture priming and conceptual priming. That is, conceptual as well as perceptual priming are implied by this strengthening assumption.

How does this relate to explicit memory? Well, explicit memories depend upon associating a presented item to its prevailing context, as shown in this SLIDE (#22 HARE in List). There are two operative types of associations: the old letter to word associations that are already there and are strengthened by reading or hearing the word; and the new association that was formed from presentation of the word in its study context.

Let's suppose that the formation of that contextual association is independent of the amount of strengthening of the old letter to word associations.

That independence assumption has far reaching consequences including the famous dissociation between measures of implicit and explicit memory. For example, Tulving showed that fragment completion of primed words is statistically independent of explicit recognition memory of them. Also, semantic processing of a word enhances explicit memory for it but has no effect on its perceptual priming. In contrast, switching from an auditory prime to a visual test word causes a big decrement in implicit memory tests but has practically no effect on explicit recognition memory tests. So these are all dissociations of the two types of memory tests.

Let's apply the independence idea to another domain. Suppose that damage to the medial temporal cortex and hippocampus makes it harder for the person to establish contextual associations, but does so without affecting the way repetitions strengthen old associations in long-term memory in the cortex. If so, then patients with hippocampus damage will suffer deficits in explicit memory, whereas their implicit memory for words and pictures will remain intact. I could go on --- in fact, my publications did so --- but you can appreciate that this small hypothesis about the use of context can explain many results on implicit versus explicit memory with a very simple story. (SLIDE OFF)

I'll move on now to the last topic, #5, on your Program Outline, on changing emphases regarding sources of context. So far I have been treating episodic memory for a context in a rather sterile fashion, as the subject simply retrieving a proposition like "The word HARE appeared in List 1". While that's fine for accounting for a lot of data, I'd like to extend the notion a bit.

I will use that as a segue into my next topic, namely, the relative contributions of external versus internal sources as providing the effective elements in the learning context. (#23-Intro SLIDE re internal)

The old style approach treated the subjects' external environment as the major determinant of the context of learning. This led to the frequent experiments on changing the room in which people learned and were tested, say, moving from one classroom or lab room to another. An embarrassing secret of this field is that such external room changes often had only a small effect on memory retrieval. (SLIDE OFF)

Various attempts have been made to beef up these external context effects. One way to do it is to instruct subjects explicitly to associate list items to objects in the room where they're learning. For example, one could associate the list-word rabbit with, say, the wastebasket in the corner, by imagining a bunny peeping out from the basket. Then at recall, you just look around the room, attending to the objects out there to cue your recall of items you've associated to them. That'll surely enhance recall for same-room subjects compared to subjects who are tested in a different room with different props and places.

While that result would be understandable, it seems somewhat unfair since it provides too much coaching, telling subjects explicitly to do what we hoped they might be doing spontaneously on their own. Moreover, we'd be sticking the retrieval cues directly under their nose.

Another way to modify context effects in the two-room experiments is to ask subjects during recall to imagine themselves back in the old room of their original learning. These imaginings can call up memories of things in the original room, and those can be used to prompt a few more memories of items learned there. And we find that that maneuver indeed reduces the impact of the changed environment.

We can take these matters even further and use a method that increases the subjects' recall immensely by providing them with distinctive mental contexts for learning each word in a list. Moreover, the contexts are portable, so that subjects can carry them around in their head and reinstate them for recall whenever and wherever they are. A way to do that uses the familiar mnemonic method of loci or peg words such as these used in one of my experiments (#24-**SLIDE** of pegs --- read a couple). I suppose most of you are familiar with this method. Subjects first learn this list of 20 pegs --- one gun, two shoe, and so on ---and are instructed to learn any new list by associating each list item with its corresponding peg. When it comes time to recall, the person runs through his mental pegs and outputs the name of the image he's attached to each peg. As we know, this can be an exceedingly effective method. For example, in this Reitman & Bower experiment, subjects studied 5 lists of 20 words, each for one trial, and at the end of an hour's learning session, they recalled 85% of the 100 words.

So, in this case we can say that subjects have a method for creating their own unique context for each item on the list being learned. Moreover, they can easily regenerate or reinstate each specific context at the time for recall of the items. Moreover, because subjects know the peg words, they can recall the list in either forward or backward order, and recall the list position in which any item appeared. And the contexts are portable, and will not be disrupted by changing environments. [SLIDE OFF]

While the pegword method illustrates the power of reinstating during recall the context of learning, it is a bit too heavy handed and circumvents the point at issue, namely, whether people spontaneously or automatically associate what they're learning with their surroundings, and utilize those associations for later recall.

In a 1972 paper, I wrote that a critical element of any context is the experiencing person --- the self or ego --- the internal sense or feel of oneself as a witness or participant in an event, although in lab studies the events were rather trivial. In more flamboyant cases, such recall of episodic memories has been called mental time travel, going back into the past, typically with awareness that one is remembering "being there then." These contextual associations provide the basis for autobiographic memories. They put our current mental state into contact with our prior mental states, and that contact underlies our sense of personal continuity and coherence.

This viewpoint corresponds to the evolving view that emphasizes the subject's internal experience as a more potent determinant of his context. That is, the important stuff is the free-flow of people's conscious thoughts and feelings that accompany their experiences with the material.

These internal effects arise when people are put into different physiological states during learning and retention. (#25-SLIDE of Physiology). One way to dramatically shift people's experiences is by giving them "mind altering" drugs. In this literature, a recall decrement due to drug-state changes has traditionally been called "state-dependent memory". State-dependent memory can be induced with psychoactive drugs such as alcohol, marijuana, heroin, cocaine, methamphetamine, light anesthesia, and so on. State-dependent memory also arises in cases of so-called multiple personality, fugue states, and deep hypnotic amnesia as the person moves from one state to another. The dissociation of memories when states are switched can sometimes be quite profound, as we have all seen dramatized in movies.

It is significant that in all these cases the drop in memory caused by state-changes shows up most strongly with free-recall. I think that's because free recall depends heavily on reinstating the learning context as cues, but that's what's missing when we switch the drug state. However, such state-changes cause no drop in memory whatsoever when we test with implicit memory measures. The earlier priming model expects this: presentations of familiar words strengthen old cortical associations that produce perceptual priming. But the drug-altered state messes up learning or recall of contextual associations. (SLIDE OFF)

An important finding regarding internal context in memory was Eric Eich's demonstration that the transfer of memories between contexts was largely determined by people's subjective experiences. Eric had subjects learn verbal items either in a lab room or outside in a lovely park; some subjects were tested for retention in the same setting and others in the different setting. Along with recall measures, he asked subjects to compare their feeling state during recall to their state during initial learning --- that is, to indicate how similar was their experienced emotional affect and arousal states during recall compared to what they had felt during their earlier learning. Eich conducted three different experiments of this kind

Here (#26-SLIDE of Eich) is what he found. Along the horizontal axis is the similarity of the subjects' current feeling state during recall to what it had been during learning. On the vertical axis is the amount recalled. Each dot represents a subject's recall plotted against his rated subjective similarity of the two feeling states. In three different experiments Eich found a significant correlation between the similarity of feeling states and the recall. This was true regardless of where the subject learned or recalled, whether the learning and recall venues matched or not, and whether the person had been given the same or a different mood induction in the two settings. What was important was the similarity of the person's experienced feeling during learning compared to recall.

Here is another summary of this data (#27 SLIDE of recalls), dividing subjects into those who judged high versus low similarity of feeling between original learning and testing. The data are also divided according to whether subjects were recalling in matching environments or in mismatching ones. You can see that although mismatch of environments had some effect (about 3%), that effect was small compared to that associated with the similarity of the person's subjective experience as he moved from the learning to the recall stage.

The important message is that it's not the external environment that is controlling recall; rather, it is the extent to which the person's subjective feeling at the time of recall reinstates how he was feeling at the time he learned. Eich also has shown that the psychoactive drugs mentioned earlier are likely having their state-dependent effects via changes in the subjects' experiential state. I think these are very important findings regarding what are the critical elements for reinstating context during recall. (SLIDE OFF)

Well, enough. Time I wrap up this talk. Rather than summarize it (Organization Slide), I'll just direct your attention to the summary Program Outline you have.

In conclusion, I would emphasize again how much episodic memory is bound up with context associations. This conference proves that contextual ideas are still alive and well. I am happy to have contributed to the early years of that research. I look forward to your research adventures into this field. I am confident that you will advance the field towards a better understanding of how context and its associations contribute to human learning and memory.

Thank you for your attention.

(#28-SLIDE of website, Stanford.edu/~gbower).