Educational Psychology Review, Vol. 14, No. 4, December 2002 (© 2002)

Interview

An Interview with Henry L. Roediger III

Michael F. Shaughnessy[1,2]

Henry L. Roediger III is presently the James S. McDonnell Distinguished University Professor of Psychology at Washington University in St. Louis and Chair of the Psychology Department. Dr Roediger has published about 150 articles and chapters, mostly on issues concerning human memory. He is coauthor of three textbooks (which have been through a combined 18 editions) and has coedited four other books. He was Editor of the *Journal of Experimental Psychology: Learning, Memory and Cognition* from 1985 to 1989 and was the founding Editor of *Psychonomic Bulletin & Review* (1994–98). He has served as President of the Midwestern Psychological Association, Chair of the Governing Board of the Psychonomic Society, President of Division 3 (Experimental Psychology) of the American Psychological Association, and on the Board of Directors of the American Psychological Society. He is a Fellow of the American Association for the Advancement of Science, the American Psychological Association, the American Psychological Society, and the Canadian Psychological Association. He was elected to the Society of Experimental Psychologists in 1994 and held a Guggenheim Fellowship for the 1994–95 academic year. According to a study by the Institute of Scientific Information that covered the years 1990–94, papers by Dr Roediger had the greatest impact (average citation count) of those of any psychologist during this period. Roediger was recently elected President of the American Psychological Society and will serve his term during 2003–2004 academic year.

In this interview, he discusses his current research and examines crucial concerns regarding memory, false memories, hypermnesia, educational psychology, and the field of education in general.

[1]Eastern New Mexico University, Portales, New Mexico.
[2]Correspondence should be addressed to Michael F. Shaughnessy, Eastern New Mexico University, Station #25, Portales, New Mexico 88130. E-mail: michael.shaughnessy@enmu.edu.

1040-726X/02/1200-0395/0 © 2002 Plenum Publishing Corporation

MS: What kinds of problems are you presently working on in your research?

HLR: Most of my present research is involved with issues of memory illusions and false memories. In particular, I am interested in how people can come to remember an event very vividly when the event never happened to them. That type of occurrence is the most dramatic case of false memories. In other cases, people can remember the core event correctly, but get the details wrong. In some real-world cases, such as those of eyewitness testimony or identification of suspects in a lineup, the details can be quite important and provide the difference between convicting the right person and convicting an innocent person.

Another research interest is concerned with the topic of implicit memory, or expressions of past experience influencing ongoing behavior. These influences can occur when people are not trying to remember and, indeed, when conscious recollection of the experiences does not occur.

MS: Why has there been so much interest in false memories, both in psychology and in the public arena?

HLR: There are many reasons. First, one great question that has guided philosophy and science throughout western civilization is how we can know what is true from the phenomena of the world around us? How do we distinguish illusion from truth, appearance from reality? In the study of visual perception, the topic of illusions has a long and venerable history. Aristotle identified the moon illusion and proposed a theory of it over 2000 years ago. From 1850 to 1900 many scholars identified and tried to explain visual illusions, and the entire twentieth century witnessed steady research on this topic. I like to think of false memories as memory illusions, but memory illusions have only been studied systematically for the past (perhaps) 30 years. Just like perceptual illusions, many memory illusions arise from processes that actually are adaptive in some situations—the illusion is a natural by-product of some process that often aids retention but that, in certain circumstances, can lead to illusions.

We all have this belief that our minds are fundamentally sound and free of error. We see the world as it is and we remember it as it was; our perceptions and memories are veridical. Thus, examples of systematic bias and error in these cognitive processes are fascinating to psychologists and intelligent laypeople alike.

Another reason for the public fascination with false memories is that the testimony of memory is so critical in so many situations. There are problems of eyewitness identification in the courtroom, problems of accuracy of children's testimony (where research shows that children are more prone to errors of memory), the whole recovered memory/false

memory controversy arising in certain therapeutic contexts, and the issue of every sort of testimony made in the public arena. Whenever there are tense congressional hearings on some issues like the Whitewater or Watergate scandals that seek to find the truth of what happened, one of the most often-heard phrases is "I can't recall." Memory problems are everywhere.

I don't think it is too far a stretch to say that issues of memory have led to the recent questioning of the death penalty. Now that DNA evidence can be used routinely in court cases, even in ones long-decided, convicted inmates on death row are being exonerated. The DNA evidence shows they were not guilty. After 11 death row inmates were freed a few years ago, Attorney General Janet Reno began an inquiry into how they were convicted and given the death sentence in the first place. Nine of the cases hinged on eyewitness testimony, which was obviously erroneous in these cases. This determination led to a conference called by Reno, which included psychologists, that led to improved guidelines on how to conduct lineups to identify suspects. So, the issue of false and erroneous memories has critical real-world implications.

MS: Now can you tell me why you think there has been such a flurry of interest in implicit memory?

HLR: The story here is completely different but has its own fascination. Research on what is today called implicit memory came from studies in neuropsychology. Patients with damage to the hippocampus and surrounding areas of the medial temporal lobes often exhibit a profound amnesia, both on standard memory tests (recall a paragraph after a brief delay) and in daily lives. The conclusion was that they basically could not learn new information and retain it over a brief delay (although very short-term memory was intact). However, in the late 1960s, systematic studies by Elizabeth Warrington and Larry Weiskrantz in England showed that, on some types of memory tests, these patients performed normally. In retrospect, these tests all provided cues (such as a fragmented form of a word that had been studied) and instructions to respond to the word with the first thing that came to mind. If the word had recently been studied, normal subjects with healthy brains show priming, a form of memory. They are more likely to complete the fragmented form of the word if they had recently studied the intact form of the word, a phenomenon called priming.

The remarkable aspect of these studies was that the brain-damaged patients also showed priming and to the same extent as the normal controls. This finding seemed hard to believe at first—after all, the patients were densely amnesic—but after a time other groups replicated these findings and today they are generally accepted.

Interestingly, the patients often exhibit little or no conscious recollection of the experiences. They show priming, but don't know that they do, and later they cannot remember the experimenter or having been in the experiment. They perform dismally on tests of recall or recognition of the words, but nonetheless show priming. The patients show retention without awareness. Much research on amnesic patients has been carried out by Larry Squire, Endel Tulving, Dan Schacter, and others. All this research reveals a dissociation in the patients (and, by extension, in all of us) between different forms of knowledge.

Why did people like me—mainstream experimental psychologists interested in learning and memory—get interested in this topic? The basic reason is that the tests used as implicit memory tests undeniably test memory, but their properties seem totally different from the properties of standard memory tests—what are now called explicit tests. (Explicit tests are those in which memory is tested deliberately; people are told to try to recall or recognize experiences from the past. Implicit tests are those in which cues are provided and subjects are told to respond with the first thing to come to mind, with no instructions to try to remember the past.) When independent variables are manipulated, priming on implicit memory tests can be shown to be uncorrelated with, or even negatively correlated with, performance on explicit tests in college students as well as in amnesic patients. Therefore, the "laws of memory" seem to operate differently for explicit and implicit tests.

MS: We can talk more about implicit memory later. Let me switch gears. Your work on repeated testing, in which you showed the phenomena of reminiscence and hypermnesia, has also received much recognition. Why are these phenomena important?

HLR: When people take multiple memory tests, they usually recall information on later tests that they could not recall on earlier tests. This recovery of information on a later test that could not be recalled in an earlier test is called reminiscence and it occurs despite the fact that no further study of the material was permitted. Therefore, this research shows the critical role of retrieval processes in the study of memory. Indeed, in those cases where there is greater overall recall of a set material on a test given later than on a test given earlier, the result is called hypermnesia. In some ways the effect is quite counterintuitive because we normally expect greater forgetting over longer retention intervals. That is, the longer the time since people studied some material, the worse (in general) is their memory for that material. That is certainly true when people take only single tests after varying retention intervals, but the interesting finding is that when repeated tests are given to the same people, information can be produced on a later test (at a greater

delay) than was produced on an earlier test (at a shorter delay) even without intervening study of the material.

MS: Why does this happen?

HLR: There is no generally accepted account, although there are several competing theories, including my own. Basically, events that have been processed distinctively or given "item-specific processing" in the jargon of the field seem to be those that are recovered across repeated tests. However, let's not go into the theoretical complications here, but instead let me tell you about three general points to emerge from this research literature that, I believe, are important.

First, this work shows that no single test of memory, of any type, is a perfect barometer of what a person knows. Test them again later—even with the same test—and you may get a different answer. Of course, that fact should not be a surprise to psychologists of any stripe. However, researchers often write and act as if a particular test of memory or scholastic achievement they are using is a perfectly reliable assessment of what a person remembers or knows, but of course there is variability in all measures. In measures of memory, particularly free recall of a set of information, items that cannot be recalled at one time may well be recovered on a later occasion. Some forgetting of items between tests will occur, too, but often the recovery outweighs the forgetting.

A second general lesson to emerge from this research concerns the whole controversy over "recovered memories." Many people seem to take the fact that memories can be recovered after a period of forgetting as indicating that the memories must have been "repressed." That is, the fact of recovered memories at a later point in time must indicate that memories were somehow repressed in the meantime. However, the hypermnesia literature shows that recovered memories are actually commonplace. Give people repeated recall tests after they have been given stories or pictures or other types of material, and you almost always find items recovered on a later test that could not be recalled on an earlier test and, in some cases, more total items can be recovered on the later test. I do not suppose anyone would want to claim that the facts or pictures remembered on the second test that could not be recalled on the first test were repressed in any way. Rather, what these studies of repeated testing and hypermnesia seem to show is that retrieval can be highly variable. The information that a person can retrieve at one point in time may not be the same as what the person can retrieve at another point in time. Recovered memories are the general rule and not some strange exception in repeated tests of memory, especially tests in which very few cues are given (such as free recall).

The third general point is that retrieval processes are a critical and often-neglected component of research into human memory. Endel Tulving began arguing this point in the mid-60s, and much credit belongs to him for beating this point into psychologists' thinking. However, I am still surprised that the critical role of retrieval is still so often neglected. One focus of interest, especially among psychologists, should be how stored information is somehow converted by the mind/brain into conscious form when a retrieval query is directed to the system. For example, what was the first point I made in this section? You may remember the gist of it—it was just a few moments ago that you heard it or read it—but the mystery is exactly how people access information when they need it. The key to understanding memory is understanding retrieval.

MS: What harm is done if retrieval processes are neglected? Can we pick a memory test, use that test, and make firm conclusions about how memory works from that test?

HLR: Yes and no. Let's say you do all your experiments on recognition memory tests (like the multiple-choice tests often used in education). For example, in studies of the neuroimaging of memory, I think about 80%–90% of all research uses recognition memory because the test can be most easily adapted for the scanner. Many behavioral psychologists have also essentially spent their entire careers studying recognition memory processes. That is fine as far as it goes and these researchers are learning many interesting facts from these studies. However, the conclusions being drawn are usually generalized to "memory." Implicitly or explicitly, the assumption is that the conclusion found about a study of recognition memory can be generalized to "this is how memory works" rather than to "this is how recognition memory works under this set of testing conditions." The distinction between these statements is fundamental because memory can be tested in many different ways and the "laws" of memory can be quite different depending on how memory is tested.

For example, take the generation effect. If information is generated, it is better remembered on some explicit memory tests than if it is read. Consider the task of subjects generating antonyms of words in one condition (opposite of hot—????) or reading words in another condition in a long list. On a free recall test, words that were mentally generated are better recalled than words that were read. However, on other types of tests there is no difference and on certain types of implicit memory tests, the finding is actually the opposite. Reading words produces much more priming on a later test than does generating words. So, we cannot simply take one memory test, whether it be recognition, free recall, cued recall with meaningful cues, implicit tests or what have you, and assume that what we find will permit us to make generalizations about "how memory works" that transcend the

tests we have used. The situation is much more complicated than that, and we must understand interactions between study conditions and test conditions.

Many researchers hit upon a favorite technique—a "preparation" as the biologists call it—that measures a particular process and then invest all their effort into studying that one paradigm. That might work well in some spheres of research, but it is a recipe that would guarantee an incomplete understanding of memory because the many different abilities that all reflect "memory" for experience can be measured in numerous ways and the measures do not always agree.

MS: Let me widen the topic slightly. Why is there so little research on long-term memory?

HLR: Well, in some sense I would disagree with the premise of this question, although I understand what you mean. I suspect that you mean, "Why is there so little research on effects of education years later, after the educational experience has been completed?" In this sense of long-term memory, you are undoubtedly correct—there is little research. Harry Bahrick of Ohio Wesleyan University has conducted pioneering studies on what students remember of foreign language vocabulary years after they have left school (and when the students had little exposure to the language in the meantime). He showed beautiful forgetting curves that reached a relatively stable asymptote that he called the "permastore" of information retained over long periods of time.

Part of the problem in your asking that question is that, for cognitive psychologists, the term *long-term memory* is used in far too many ways. Long-term memory has been used by researchers in different traditions to refer to words presented in a list 10 seconds ago, to autobiographical memories from early childhood, to events that happened over the past week, and to events permanently stored in memory (what Tulving calls semantic memory—our general storehouse of knowledge).

All of these uses, in the right context, can justifiably be called long-term memory, but obviously they differ tremendously. We could add to them, too: How you remember to ride a bicycle, or tie your shoes and other motor skills (or procedural memories) could also justifiably be referred to as instances of long-term memory. In some sense, then, there is a tremendous amount of research on long-term memory, but in the specific sense that I believe you have in mind—long-term effects of education—there is very little.

MS: What are the main issues currently facing educational psychology and educational psychologists?

HLR: I am a cognitive psychologist and I am not really qualified to speak directly on the main issues facing educational psychology and its practitioners. However, I think it is safe to say that one aspect of this issue would be

to raise the entire image of educational psychology. I am constantly amazed, as I listen to the debates on what should be done about our schools and educational systems across the country, that educational psychology seems to have so little influence in these debates. There is a tremendous amount of research directly relevant to the practice of education generated by educational psychologists, and I fear it is somehow not being translated into actual practice in the schools.

MS: Cognitive psychologists have had some influence on the research and practices of educational psychology. Do you believe the influence has been strong enough? What further advances would you hope to see?

HLR: Again, because I am not an educational psychologist, I do not feel safe making pronouncements about that field. However, from my limited knowledge I do believe that some of the findings of cognitive psychology have permeated educational psychology. The converse is certainly true, too. Indeed, much of the early research on prose learning and retention that today forms part of cognitive psychology was originally conducted by people who were affiliated more with educational psychology than with cognitive psychology. At least part of educational psychology might be perceived as cognitive psychology applied to classroom and textbook learning experiences. Still, I do believe there is room for improvement in communication between cognitive and educational psychologists and that many of the more recent findings in cognitive psychology might still apply to and improve educational practice. Educational psychologists would be the natural intermediaries in this process.

MS: Could you give an example of findings that might be applied?

HLR: Here is one example, which I borrow from writings by Bob Bjork of UCLA. There are many findings in various areas of psychology that show an interesting pattern: methods of teaching and training that lead to relatively rapid learning (as compared to other methods) are often favored in training and in education. That seems natural, of course. After all, if learning occurs more quickly (when measured objectively), then the techniques producing rapid learning ought to be used rather than ones producing slower, more laborious learning. However, Bjork argues that there can be a negative side to this practice because certain methods of training that promote quick learning produce less stable retention than do methods that produce slower learning.

Let's take a hypothetical example of learning skills for a job. Suppose I want to teach a group of people how to practice five different skills for their jobs (say, making widgets or something as part of a manufacturing process). One way to do so would be to teach each component skill and have it be practiced over and over until the person reaches some criterion

of mastery. Then the instructor could go on to the next skill and do the same thing, and so forth and so on, until all five skills are learned. If one practices the same skill repeatedly until some point of mastery is reached—a condition called massed practice—learning is relatively quick on each component skill.

The alternative to this kind of massed practice would be to have distributed training, such that one practices for a while on one skill (but does not get very good at it), then practices another skill for a while and so forth, cycling through the five skills. The five skills might each be practiced for a while in a more or less random order, recursively, until the same criteria of mastery are met. When acquisition curves of learning the skill are plotted, it takes many more trials to reach criterion with this form of distributed, random practice than with massed practice. Therefore, in industry and in many other settings, massed practice is often used to teach new skills. Learning is quicker and when asked to assess the two types of training regimens, the people who are being taught greatly prefer massed practice to distributed practice. Distributed practice seems "harder." Indeed, if you look at the research literature, massed practice does generally promote faster learning than distributed practice on relatively immediate tests. So, use of massed practice would seem to be a win–win situation, because learning is faster and more efficient and the students like this process better than when they have distributed practice training regimens.

But there is a downside here, one rarely appreciated in the field. If one uses a delayed retention test—one that tests those skills a month or two later—distributed practice is usually shown to create longer lasting and more durable learning than does massed practice. Long-term retention is better following distributed than massed practice. Put another way, forgetting is greater following massed practice and so a person may forget the skill after massed learning before he or she needs to use it.

From one point of view, this makes sense. After all, in the long-delayed test a person must suddenly practice the skill out of context when it has not been done in a while. Distributed practice fosters this type of "out of context" performance whereas massed practice does not. However, as Dr. Bjork also points out, practice in industry and often in education (say, learning to write letters in first grade, or to sound out names of words, or to do a certain type of arithmetic problem) favors massed practice because learning is quicker and student ratings are higher. I believe a series of studies that examined the massed/distributed practice effect in direct educational settings, using both immediate and delayed tests, would be a welcome addition to the literature. Perhaps such studies exist now with prose materials, arithmetic facts, and others, but I know of relatively few studies that generalize the findings of the laboratory in more educationally relevant contexts.

In some ways, the findings I am describing are counterintuitive because most educators (and psychologists) believe that all measures of memory should correlate positively. A form of practice that produces good learning immediately should spill over and be more durable on a delayed test. However, in many cases the opposite is so—what makes learning easy in the first place also makes forgetting more rapid. If learning is made slow and somewhat more laborious, it may last longer. And, of course, the type of test used to assess retention also matters—not all tests show the same effects.

MS: That is an interesting example of research from cognitive psychology. Can you give another one?

HLR: Sure, although in some ways it is another example of the same type. We generally believe that good, coherent organization of materials fosters better learning and memory. In fact, many experiments show conditions under which that statement is true. In practice, textbooks typically contain advance organizers within chapters. Certainly, in many circumstances, providing organization is to be encouraged. However, some fascinating work by Mark McDaniel at the University of New Mexico and his colleagues (as well as by other researchers) shows that often learning can be better when it is made more difficult. For example, in one of McDaniel's studies he showed that if letters are left out of words in text passages, so that people must expend more effort in reading, then people remember the material better. The words are not so degraded that they cannot be deciphered, but reading is made slower and more effortful and the reader must work harder to make sense of the whole passage, relative to normal reading.

In addition, other studies have shown that if an advance organizer of a prose passage does not match exactly the form of the text, learning is also better, which seems surprising. Students have to work to figure out the text and its relation to the organizer. With letters omitted and with advance organizers that are ambiguous, people have to work harder to understand the material, and if they do make that effort, they remember the material better. Both these conditions seem to create what Bob Bjork calls "desirable difficulties" for the learner.

Now, no textbook company is going to create "advance disorganizers" for the chapters of their textbooks and then leave letters out of words in the books, but perhaps for some students this would actually help them learn the material better. Sure, it would slow reading, but if students would make the extra effort to go through the material, they may understand it better and remember it better later on. They have had to do the work to organize the material themselves and pay more attention to their reading, rather than be handed the organization by the textbook writer.

Speculating, one could hypothesize that, for highly motivated students, this type of "difficult" presentation format would be more effective for learning and memory, whereas for less motivated and less engaged students, the traditional "easy" organization might be better. In fact, there is support for these speculations in some of McDaniel's recent work. Very good students (skilled readers) profit when words are degraded in text by leaving out letters, but poorer students who do not have good word-decoding skills are harmed by this manipulation. On the other hand, when sentences are scrambled and good and poor students are asked to put them in a coherent order, this manipulation helps poor students' retention but does not benefit good students' retention (relative to reading normal text in both cases). The reason seems to be that the good readers can already recognize and benefit from the organization of a text in normal reading whereas poor readers don't naturally do this. Reordering the scrambled sentences forces them to think about the organization of the story and to integrate it in a coherent fashion. In short, reordering scrambled sentences creates a desirable difficulty for poor readers whereas leaving out letters of words creates an undesirable difficulty for them. For good readers, the situation is the opposite—good readers do not profit from reordering scrambled sentences but do benefit from the letter deletion manipulation in reading coherent text.

MS: Is there a general principle from what you are saying from all these examples?

HLR: Yes, to sum up, I am saying that some educators seem to have forgotten that good learning often requires effort and that creating a challenge in education is not necessarily a bad thing. Instead of continuing to make situations easier, so that text materials are written at increasingly more basic (that is, low) levels, we need to provide challenges that provoke effortful and effective learning. However, challenges need to be geared to students' current levels of expertise.

To put it another way, the more students work at a subject and the more effort they expend, often the better they learn and remember. That sounds like a platitude, but it is remarkable how frequently it is forgotten in education. It seems to me that many educators are working hard to make learning simple and fun. They believe, "We need to engage the students and keep them entertained." I believe that is true—we need to make material as interesting and as relevant as possible—but often we need to challenge them to make them think. Of course, we cannot simply overwhelm students by presenting material of great difficulty for which they are unprepared. Rather, one could imagine "graded challenges" in teaching, such that the first hurdles would be relatively easy and they would increase in difficulty.

We need to create desirable difficulties, not undesirable ones, and discovering the difference should be a great challenge for educational research.

MS: Can you give an example of an educational practice that fosters learning that is "too easy?"

HLR: The emphasis on massed practice forms of instruction is one case. I believe another is represented by the work on programmed learning that came out of behavioristic analyses of learning. Programmed learning used to be a big deal and I can recall many texts from the 1960s, including some I used in college, that embodied the principle of giving students a statement and then, soon after, giving them a sentence frame in which they had to demonstrate their knowledge and have it "reinforced" by getting a right answer. However, the programmed learning materials were made too simple. People could read the statement, hold it in working memory for a brief while, and then give it back. I had a textbook like this and it was mindlessly simple. I would just read the statement, then read the sentence frame following it with a critical phrase left out, and then produce the phrase. I expended no effort and got little from the experience of generating the answer. Still, I believe the idea behind programmed learning—getting students to be interactive, to provide responses—is a very good one, if done right.

MS: How would you do it right?

HLR: Make the student expend more effort in the process. As I said, the basic idea behind programmed instruction was probably right on target. We need to get students to produce critical responses in the right context to show they understand the concept. However, the method of getting students to produce the responses needs to be changed, in my opinion. There is research on the positive effects of testing, on the generation effect, and on the effect of spaced practice that should be brought to bear on this issue. We know that testing one's knowledge can have quite positive effects on retention and in fact correct performance on a test can have a more positive effect on a second test, given later, than would additional study of the material. Correctly retrieving information on a test has a powerful mnemonic effect. In addition, much research shows the power of generating information under somewhat difficult circumstances relative to simply reading it. So, potentially, programmed learning could benefit from both of these effects. However, in both cases, the positive effects can be short-circuited if the testing or generation occurs under easy conditions. That is, the rememberer must expend some effort during the test (it should not involve retrieval from short-term memory) or during generation. Much research in laboratory studies of human memory is consistent with these conclusions. People need to learn to retrieve material out of context and in the face of interference.

MS: Given what you have said about testing, do the types of tests used in education make a difference?

HLR: I think so. I know multiple-choice tests are favored in some quarters because of problems of reliability in scoring essay tests and the fact that multiple-choice tests can sample the material more widely. And, of course, they can be scored much more quickly and easily. Still, research by educational psychologists shows that students preparing for essay tests report studying more and also looking for general themes and the higher order structure of material they are studying. When preparing for multiple-choice tests, students report studying less time and also spending more time learning local facts rather than general themes. In laboratory experiments in which students are told to expect either a recall test or a recognition test, they generally show superior performance (no matter what kind of test is given) when they have prepared for a recall test. So, encoding strategies in preparing for recall tests (essay, short answer) seem to be more effective than those used in preparing for recognition tests (multiple choice, true/false) even when the amount of study time is equated between the conditions. And if study time is not equated and students do what they naturally do to prepare, they study longer for short answer/essay tests than for objective tests like multiple choice or true/false.

Another reason to favor essay and short answer tests over recognition tests is that the testing effect is greater when people have to produce answers rather than read answers and check them off as correct or incorrect. Looking across the literature, it appears that tests that require people to actively retrieve and generate information have more positive effects.

So, learning is probably greater when preparing for essay/recall tests than for multiple-choice tests for all these reasons. Yet, in many educational circumstances, students mostly receive multiple-choice tests for practical reasons.

MS: How does your own research on transfer appropriate processing impact decisions about testing?

HLR: The idea of transfer appropriate processing is simple, but I believe that the implications for testing are profound. Briefly, the idea is that one can create many different tests of memory and (most broadly) of knowledge. These tests can vary dramatically in their requirements. As I have already pointed out, the idea that all memory tests measure the same type of "memory strength" is fallacious according to this transfer appropriate processing point of view. The argument from the transfer appropriate processing approach is that the type of processing engaged during a study or presentation episode transfers to a greater or lesser extent to a test depending upon the matching of the test and encoding requirements. Some types of strategies

lead to encoding that is quite appropriate for some types of tests but inappropriate for others. Conversely, "inappropriate" strategies that can be used for encoding information for some tests may be perfectly appropriate for other types of tests.

My own research comparing explicit and implicit memory tests has revealed a large body of evidence consistent with these general principles. Most explicit memory tests tap meaning; I have called them conceptual tests. Standard free recall, cued recall, and recognition tests all generally fall into this category. Therefore, study operations that emphasize meaning relative to other dimensions (such as sound or appearance or other surface features) lead to better memory performance, in general, on most explicit memory tests. Variables such as modality of presentation (auditory or visual) often have little or no effect on most conceptual tests. However, on certain implicit memory tests where the interest is in whether priming past experiences affects behavior, modality often matters and meaning does not. Consider the task of completing fragmented words such as e_e_han_. People have a difficult time completing that fragment out of context, but if they have seen the word *elephant* half an hour ago in a list, the completion rate is much higher. Past experience implicitly primed or biased completion of the word fragment. This type of priming test in which a fragmented word is given is called a perceptual implicit memory test and, as the name implies, we have shown this type of priming to be quite sensitive to the type of perceptual processing in which a person is engaged during study. For example, visual presentation of a word leads to about twice as much priming as does auditory presentation of a word on this test. However, on a standard test like free recall there is not much effect for modality except on the last few items presented in a list on an immediate test, and even then auditory presentation leads to better recall than visual presentation. Another interesting case involves presentation of pictures and words (a picture of an elephant or the word *elephant*). On most explicit tests, pictures are remembered better than words. However, on perceptual implicit tests such as word fragment completion, words produce considerable priming and pictures produce little or no priming at all. Once again, perceptual implicit tests are tuned to the surface properties of stimuli whereas conceptual implicit tests tend to be affected by meaning

The huge body of evidence supporting the transfer appropriate processing perspective also shows the relativity of retention. That is, a particular variable (such a standard level of processing manipulation) may have powerful effects on one type of memory test, but no effect or even the opposite effect on other types of tests. From the old strength theory perspective—memory traces vary in strength and tests are simply differentially sensitive in revealing this strength—such complex interactions are

unexplainable. However, from the transfer appropriate processing perspective these sorts of interactions are the rule and not the exception. Tests of memory differing in their requirements produce widely different learning effects.

MS: What does transfer appropriate processing have to do with education?

HLR: Everything. All educational assessment involves particular types of tests. Therefore, following through on our principle, we can certainly say that some types of study will lead to better performance than will other types of study and therefore study strategies should be tailored to the test. More generally, the types of tests required in educational assessment may not be those required later in life where the knowledge must be transferred and expressed in some different way.

Consider graduate education. Many departments have strict qualifying exams, although I would like to think these are becoming a thing of the past. In qualifying exams, students have to take a test that lasts a day or two, usually part written and part oral, in which they are responsible for a gigantic body of literature and have to answer questions about it. The material is often drawn from all of psychology. Students study long and hard for these tests, and surely they learn much from the experience. However, I have always argued against this form of graduate assessment for admission to candidacy to the PhD. After all, the learning experience is completely inappropriate for later requirements in the field. Never again in the student's career will he or she be asked to sit for a test. Why have them spend six months or a year of graduate education preparing for this test? It is simply inappropriate for what they will later be asked to do.

MS: What kind of qualifying exam would you recommend in graduate education?

HLR: In most graduate programs, students are required to take a certain number of courses in their discipline and across neighboring areas, to ensure some breadth. I would let their course work substitute for the usual qualifying exam experience. For qualifying exams, my strong preference would be to ask students to pick a broad area in which he or she wants to become an expert. Suppose, for example, the student says that she would like to become an expert on the literature on memory for prose and what factors create good and poor memory for prose materials. That is a nice broad category. I would then have the student construct a reading list on this topic on which I (and other committee members) would comment and approve. Then the student's task would be to (a) read this literature, (b) summarize and synthesize it in a review paper, and (c) write a set of research proposals (hypotheses to be tested, general methods by which the research would be conducted) as part

of the qualifying exam. Then students could be assessed on their knowledge and defend their writing in an oral exam.

To my way of thinking this would be a wonderful experience for the student and better prepare him or her. If that student is going to become a serious academic, he/she needs to learn to review and synthesize the literature on a problem and to ask cutting-edge questions about that field or topic. Therefore, this type of qualifying exam might help launch the student in the direction of her dissertation and establish an important set of skills for professional life. Spending six months or a year becoming an expert in a broad topic and learning to think creatively about that topic is certainly much more useful to a student than simply studying a large body of literature across all of psychology in order to answer questions on a relatively arbitrary test. In short, this type of qualifying exam would transfer appropriately to the student's later career as an academic.

MS: Who initiated the transfer appropriate processing ideas that you are applying?

HLR: The issue of "appropriate transfer" has a long history within psychology and dates at least to John Dewey and his theory of identical elements of transfer. However, in his treatment, the emphasis was on contents of material and not so much on types of processing. The modern incarnation of transfer appropriate processing theory and research was introduced by John Bransford, Jeff Franks, and their collaborators at Vanderbilt University in the late 1970s. It was their work, really, that inspired me and I applied their ideas to the differences found between explicit and implicit memory tests. Of course, the ideas bear a fundamental similarity to Endel Tulving's ideas on the encoding specificity principle and the study/test interactions that he produced in many experiments in the 1970s. It is interesting to me that shortly after John Bransford developed these ideas, he essentially left his mode of research as an experimental psychologist and went on to do his important work in education and technology. Transfer appropriate processing ideas seem ready-made to apply to education.

MS: Do you see yourself moving into more education-oriented circles as well?

HLR: I am very interested in educational issues, and Bob Bjork and I are discussing the possibility of collaborating on research where we would apply ideas from cognitive psychology to more educationally relevant materials and tests. I have alluded to this kind of work repeatedly during this interview. We need to make sure that what we believe we know from basic research on learning and memory does indeed apply to educationally relevant situations. This may be a possible future avenue of research for me, but I really have not pursued it yet.

MS: Let's end on a broad note. Where do you think cognitive psychology is headed as a field? What lies in the future?

HLR: That's a hard one. Anyone who answers will be telling you more about their own biases than about the field of cognitive psychology. Still, I was persuaded to write a chapter a few years ago on this topic, and called it "The future of cognitive psychology?" (Roediger, 1997). The question mark indicated that probably 80% of the specific predictions in the book—about the future of cognitive psychology—would be wrong.

As one broad trend, I believe the great interest in cognitive neuroscience approaches will continue and even increase. Rather than the metaphorical box-and-arrow models that cognitivists favored from the 1950s to the 1990s (and that some continue to use today), today's models will be neural ones and the theories to explain behavioral phenomena will be tested both in behavioral studies and neuroimaging studies. That is certainly happening in the study of memory. More broadly, thinking about the neural implementation of cognitive processes also sharpens cognitive theory about these mental processes. So, neuroscientists borrow cognitive paradigms and in studying them force cognitive psychologists to confront questions that had not yet occurred to the field. The collaborations are therefore mutually beneficial. I will stop at this one point and send any interested reader to the book chapter for more.

Send interviews or interview ideas to Interviews Editor, Michael Shaughnessy, at Eastern New Mexico University, Station # 25, Portales, NM 88130. e-mail: michael.shaughnessy@enmu.edu

---

REFERENCE

Roediger, H. L., III (1997). The future of cognitive psychology? In: Solso, R. L. (ed.), *Mind and Brain Sciences in the 21st Century*, MIT Press, Cambridge