CHAPTER 1

Introduction

Memory records our experiences, defining who we are and how we will behave. It is therefore not surprising that psychologists – no matter their research topic – study human memory either directly or indirectly. The layperson is also interested in memory, often because many people feel they have poor memory. The belief often holds even though this is not so; most people get through daily life fairly easily. So why do people think their memory is poor? One hypothesis is that people are usually aware of their memory only when it fails. People forget where they parked their car, forget to take their prescriptions, forget their passwords, and do not always achieve an A+ even when they feel that they are well prepared for an exam. Observing the failures of memory in others can also justify a belief that poor memory is common.

But memory usually serves us quite well, and we might instead be fascinated by the fact that as we go through life, not actively trying to remember anything in particular, we easily remember daily events. We can remember a conversation we just had on the phone with a colleague, even though we were not trying to. The meal you cooked last night was subpar, and you still feel bad, although you would like to forget that. Shown any photo on a phone, we can remember where and when it was taken, even though we were not trying to. Most people's memory is actually quite remarkable, and rather than be too concerned about forgetting, we might be inspired by the mind's covert recordings of daily activities, how we remember distant events, and how we create knowledge. This book is about how we remember, why we forget, and other aspects of human cognition.

Overview of the General Theory

In the 1960s at Stanford University, Richard Atkinson and Richard Shiffrin coauthored a series of technical reports and articles that put forward a theory of memory (Atkinson & Shiffrin, 1965; Atkinson, 1966; Phillips, Shiffrin, &

 https://doi.org/10.1017/9781009440370.003

Atkinson, 1967). They used the term "General Theory" to refer to their proposals, and that is the term to be used here. The key article is titled "Human Memory: A Proposed System and Its Control Processes" (Atkinson & Shiffrin, 1968). It is one of the most-cited publications in the psychological and cognitive sciences; hundreds of experiments have been conducted testing one or another aspect of the theory. This book documents how the theory is able to explain, in quantitative detail, a wide range of memory phenomena. Forgetting and retrieval of memories were familiar areas of research and early topics for investigation. But sometimes the General Theory has been used to understand aspects of human cognition that Atkinson and Shiffrin could not have anticipated, ranging from visual search and implicit memory to the creation of knowledge. The scope of the General Theory expands to this today, and there is no obvious reason why it will not continue to do so in both predictable and unpredictable ways (cf. Goteti, Cybary, Dynes, 2023).

The Memory System. The General Theory is simple, boiling down to two central assumptions. The first is based on a proposal espoused by, among others, William James (1890) and Donald Broadbent (1958): Human memories are classified as those that represent the current contents of thought and those that represent the relatively distant past. These categories were labeled by James and Broadbent as primary and secondary memory, respectively. But Atkinson and Shiffrin referred to them as the *short-term store* and the *long-term store*, highlighting their different functions in human cognition; the short-term store temporarily stores memories and the long-term store permanently stores memories. Atkinson and Shiffrin also assumed a third category that very briefly represents perceptual stimulation, which is referred to as a *sensory register*. The existence of the three types of memory stores is one of two central assumptions of the General Theory.

The memory structures support all thoughts, decisions, and actions of cognition. Following this undeniable assertion, an important question naturally arises: How can only three categories of memory structures support the enormous varieties of human behavior? The General Theory addressed this question from an information-processing perspective that developed in the context of empirical studies indicating that exogenous stimulation is transformed as it proceeds through stages of cognitive processing.[1] By retrieving information from one memory structure and

[1] In this area of research, there was a movement toward formal mathematical models of learning, which provided concrete evidence that there was much to be gained by specifying how the context in which learning occurred altered how information was processed (Estes, 1950, 1955; Estes & Burke, 1953). And advances in computational technologies provided mechanistic analogies and algorithms to compare human information processing, especially those that supported the first models of artificial

storing it in another, information flows between them. In other words, as information flows between memory structures, it must be retrieved from one and stored in the next. The structures and processes for information transfer are sometimes generically referred to as the "memory system."

The memory structures represent different aspects of experience. The short-term store maintains a limited amount of information actively used to perform a task or meet a current goal. With this in mind, Atkinson and Shiffrin referred to the short-term store as *working memory*. The original source of the information in working memory can be the long-term store, a sensory register, or both. For instance, one might view a note indicating one needs to stop at the grocery store on the way home from work, and this visual stimulation that was once briefly represented in a visual sensory register may be transferred and linguistically represented in the short-term store. When entered into the short-term store, the representation of the note can be used to retrieve from the long-term store what is missing from the refrigerator and therefore needs to be purchased. This is an example of how the different memory structures work together to support cognition.

Control Processes. Now, we can establish the second major assumption of the General Theory: information flows through the memory structures to serve specific goals under different circumstances. For instance, the contents of the grocery list can be kept in the short-term store in order to create a new long-term memory for what needs to be purchased. This is an example of how one can control the contents of the short-term store to serve a future goal and how the contents of the short-term store affect the contents of the long-term store.

Cognitive control of information processing gives the General Theory the flexibility to account for behavior in a wide variety of conditions. Since there are a very large number of combinations of tasks and goals, there must also be a large number of models required to describe behavior – that is, to describe how information is transformed and flows among the memory structures. These models are referred to as *control processes* because they manage information in the memory structures to support the performance of cognitive tasks.[2]

intelligence. The interplay between empirical observations and formal theory had been the hallmark of scientific understanding for centuries (Bacon, 1620/1898), but the theoretical and technological developments of the 1950s provided Atkinson and Shiffrin with the opportunity to develop new mathematical models of how information flowed among the human memory structures in response to task demands and to the goals of the human subject.

[2] In the scientific literature, the terms "theory" and "model" are often used interchangeably, almost as synonyms. However, Atkinson and Shiffrin make a sharp distinction between the two. Theory is used to describe the memory system and the variety of controls that determine how information

Returning to our everyday example of going to the grocery store, one might invoke a control process that uses working memory to create a mental image of stopping at the store, walking into the store, and traversing the aisles to create new long-term memories to make this stop as efficient as possible. Indeed, athletes routinely use visual imagery like this to help them prepare for future competitions. The upshot is that what and how we think in the present will affect future behavior when the contents of the short-term store are transferred to the long-term store.

The notions that different aspects of past and present experience are represented in different memory structures and that the flow of information is controlled in different ways in order to achieve a goal are the two essential assumptions of the Atkinson and Shiffrin framework, referred to here as the General Theory because all subsequent developments in the theory are grounded by them. In the scientific literature, some authors refer to the General Theory as the "Modal Model," as a testament to its ubiquity. Why has the General Theory proved so successful over such a long period of time? Two reasons stand out. First, making explicit the distinction between the system itself and the control processes that run the system represented a major change from earlier theories. Second, short-term memory plays a central role in the theory. It is where new memories are created and where they can be combined with older memories to build knowledge structures. It is not just a waystation to storage but a place where information is actively mixed, sorted, and prioritized for long-term storage. For instance, memory does not simply record experiences; we control what is recorded, as one controls what to take a photo of, edits photos, and decides which photos to view in the future. In sum, the theory provides an elegant and dynamic description of the mind's mechanics.

Memory Traces. Although our understanding of human memory is incomplete, one thing is clear: experiences leave a trace of the past within us. Unfortunately, we can only be a little more specific; traces of experience are left in our brains. This might seem obvious, but there are everyday examples that incorrectly suggest something else. For instance, athletes often refer to muscle memory as the ability to perform at a certain level even after taking time off from training, suggesting to them that muscles learn and have memories. This is a common misunderstanding. Memories

flows through the system in a general sense. When the theory is applied to a particular experimental task, a control process must be rigorously specified that is appropriate for the task. The term model is used to refer to a specific instance of a control process applied to a specific task under the conditions defined by the experimental design and procedure employed. Hence, a model is regarded as a special case for implementing a set of assumptions within the framework of the General Theory.

are not stored in muscles but in the brains that control the muscles. Trauma to the brain can lead to paralysis, but losing the ability to move does not affect the ability to learn, think, or remember. Although we cannot point to an individual memory in the brain, we know that various brain traumas and maladies disrupt memory in various ways; so, the memories must be somewhere in the brain.

A view currently popular with neuroscientific memory researchers is that memory leaves traces of past experience widely distributed throughout the brain, an idea proposed by the psychophysicist Gustav Fechner as far back as 1882. However, there is now overwhelming evidence that different parts of the brain are involved in different aspects of memory; disruptions to memory occur when specific parts of the brain responsible for recording or retrieving memory traces are damaged or removed. The disruptions to memory suggest the functional roles of anatomically distinct brain areas. In addition to observing the effects of brain impairments on memory, behavioral research documents how intact memory works or fails under various conditions using memory tests. Changes in memory associated with different tasks or conditions suggest different processes that contribute to a healthy memory. Hence, we can use a variety of methods to investigate memory without knowing exactly where a specific memory resides.

Researchers often refer to a trace of past experience as a *memory trace* or simply a *trace*. For now, I will use the term memory trace to refer to a mental record of past experience, backsliding to the jargon of the memory researcher by using the abbreviated term trace in subsequent chapters. But let us be clear. Everyone knows what a record is, and therefore everyone knows what a memory trace is. Both document what occurred in the past. Just as The Beatles documented their songs on vinyl albums or newspaper box scores document what occurred during a baseball game, memory traces document, albeit in a noisy and incomplete fashion, our personal experiences.

Consider what memory traces record. A memory trace can represent a specific event or general knowledge (e.g., McClelland, McNaughton, & O'Reilly, 1995). Episodic memory traces represent past events that we experience. Semantic memory traces represent general knowledge learned through many experiences. For instance, everyone can remember a childhood event from school, and everyone knows that a school is a place where students go to learn. How experience leads to knowledge is one of the most important questions researchers ask.

One hypothesis derived from the General Theory is that records of similar experiences accumulate in memory traces originally stored at an earlier time (Shiffrin & Nelson, 2013). When similar events are repeatedly experienced, information about each event is acquired and accumulated across events. Hence, these traces no longer represent any single event but represent what several events have in common and may form the basis for inductive inference. For instance, if one encounters a school, one infers that teachers educate students there and the building contains desks and chairs, among the many other attributes or *features* that characterize a school.

Chapter 7 is devoted to discussing knowledge within the framework of the General Theory. For now, we will focus on episodic memory traces because knowledge must arise from experience, and therefore this is a convenient place to begin. It is often assumed that a finite set of information about an event one experiences is recorded in an individual memory trace. That is, there are separate traces representing individual events, like books in a library. If for no other reason than this makes thinking about memory relatively simple, this is the position I will take, recognizing that there is little evidence to support it. And, as you will see, experimental designs used to study memory are amenable to the hypothesis that different events are represented by different memory traces.

Events. An episodic memory trace has a role similar to a journalist, documenting what was experienced and with whom, where an event was experienced, and when it was experienced. What was experienced and with whom is often referred to by researchers as *item information* and where and when an event occurred is often referred to as *context information*. Both item and context information play important roles in the General Theory.

To illustrate the distinction between item and context information, let us assume that you go for a walk as often as you can. Although this greatly simplifies matters, define an event as what occurs between the time that you leave home and the time that you return. The thoughts you had, the weather, and the people you saw would be considered item information, as would other information that you attended to on your walk. However, if asked – How was your walk yesterday? – the combination of the item information (the walk) and the context information (yesterday) specifies a unique event. In fact, researchers have identified neurons in the brain that are referred to as place cells, which fire only when in a particular location (O'Keefe & Nadel, 1979). Hence, to accurately remember an individual event, one requires in a record the basis for inferring when and where the event occurred in addition to what happened.

Laboratory Methods. This example of remembering an event that occurred during everyday life is referred to as an *autobiographical memory* *task*. I will have more to say about autobiographical memory, but research on autobiographical memory has not played a large role in our understanding of memory. The problem with studying autobiographical memory is that it is difficult to scientifically manipulate in a well-controlled fashion, making experimental designs quite expensive, and without such control it is unknown whether one's memories are accurate or not. Therefore, much of the significant research on human memory has utilized what is sometimes called a study-test procedure in laboratory experiments.

The study-test procedure involves asking a subject to try to remember a list of words, pictures, sounds, or novel objects. These stimuli are referred to as *items*, and it is assumed that a new memory trace is created for each item on the list. That is, an event consists of the presentation of an item and the context in which the presentation occurred, and a memory trace records the occurrence of each item in the long-term store.

Most experiments use words as stimuli because they are easy to present and because their occurrence is easy to report in a clear and succinct manner. Knowing what was experienced allows the accuracy of memory to be measured. But use of words in memory experiments presents a challenge. Like daily walks, words have been experienced many times; the context in which a word was encountered allows the different encounters to be told apart. Hence, even in a laboratory setting, the memory trace representing the occurrence of an item on a list must contain both item information (the word) and context information (where and when the word was encountered).

Sensory Registers

Much of the information that comprises thought is obtained from the outside world: conversations we have, sights of our environment, tactile stimulation from a keyboard, the taste of a taco, the smell of a spring day, and so on. One question that immediately arises concerns the number of sensory registers that exist. The obvious answer is that there is a sensory register corresponding to each of the five senses.[3]

Span of Apprehension. Some of the most compelling evidence supporting the existence of a visual sensory register came from research on what was known as the *span of apprehension*, which refers to the capacity of

[3] However intuitively appealing this answer is, Atkinson and Shiffrin (1968) concluded that extant evidence only supported the existence of a visual sensory register, but left open the possibility of additional sensory registers pending new findings.

the visual system to represent objects or stimuli. Perhaps the earliest experiment was conducted by Jevons (1871), who tossed a handful of beans in the air so that some of them would land in a white square outlined in the middle of a black tray. Jevons's task was to estimate the number of beans that had fallen within the white square with only a brief glimpse of its contents. Jevons repeated this process many times. The critical finding was that the accuracy of estimation declined dramatically as the number of beans that fell within the white square increased past nine. Jevons concluded that only about nine objects could be represented by the visual system at a given time.

With a similar experiment, Averbach (1963) controlled the amount of time black-dot stimuli were visually displayed. In doing so, Averbach was able to more rigorously control the availability and exogenous visual stimulation used to make the numerical judgment. The results obtained were remarkably similar to Jevons's. However, Averbach's subjects reported "seeing" all the black dots, even after the stimuli were removed from the display. These introspections suggested that although the subjects could not accurately report the number of stimuli displayed when more than about nine stimuli were presented, the visual system nevertheless represented all of the stimuli, albeit quite briefly.

Contrary results to those of Jevons and Averbach were separately obtained by McDougall (1904) and Whipple (1910), both of whom used a tachistoscope to briefly present a visual array of letters varying in number. The subject's task was to report what letters were displayed. Whereas Jevons and Averbach found that about nine beans or dots could accurately be counted. McDougall and Whipple found that only about four or five letters could be reported.

The different measures of the span of apprehension could be understood following the results of a seminal study conducted by George Sperling (1960). Using tachistoscopic methods similar to McDougall's and Whipple's, Sperling presented a three-row array of letters for just a few milliseconds. In one condition, the subject's task was to report as many of the letters in the display as possible. Not surprisingly, subjects were able to report four or five letters. In the critical condition, an auditory cue presented after the visual array indicated which one of the three rows of letters to report. Again, subjects were able to report four or five letters, which indicated that the span of apprehension was actually twelve to fifteen letters because all elements of each row must have been available upon the presentation of the auditory cue. In other words, all the letters of a three-row array were available for processing. Moreover, as the delay between the presentation of the visual array and the auditory cue increased, the number of letters reported decreased, suggesting that while all visual stimuli were initially represented in the visual register, the contents of the register very rapidly decayed.

Sperling's results suggested a way to reconcile the results of Jevons and McDougall. It is clearly the case that it takes less time to verbally report a single estimate of the number of items in a display than to verbally report the identity of each one, as Jevons and McDougall's task required respectively. Hence, if the contents of visual register decay quite rapidly, then one should be able to verbally report with greater accuracy the number of items displayed than to report each of their identities. Indeed, Averbach's results indicated that subjects were able to make such judgments. Thus, these critical findings indicate that representation in the visual sensory register is like a photograph that fades quite quickly.

Neural Basis for the Visual Short-Term Store. The relationship between the mind and the brain is one of the most interesting and important topics of scientific investigation. One key assumption of the General Theory is that different memory structures support different types of memory representations. The relationship between these structures is important to understand from an information-processing perspective, and the notion that information is transformed via stages of mental processing is consistent with the notion that different brain structures support different stages of processing.

These ideas have been around for a long time, of course. One of the more provocative proposals actually inspired a cultural phenomenon in the early to mid-nineteenth century known as phrenology. Phrenology grew from the work of Franz Gall, who proposed that different parts of the brain support different mental traits, and, in 1820, the Edinburgh Phrenological Society was established by the brothers George and Andrew Combe. Phrenologists, like the Combe brothers, believed that by inspecting the shape of the skull under the human scalp, one could ascertain what traits were more or less present in different individuals, and people naively subjected themselves to phrenological examinations as means of entertainment and even as the basis for social reform. The assumption that one's personality could be derived by a simple physical inspection of the scalp was dubious at best, but the idea that different parts of the brain support the different structural components of thought also gave rise to the enormous field of cognitive neuroscience.

Take, for instance, the visual sensory register. Early visual processing involves several brain regions, each responsible for a different stage of information processing. All visual processing begins in the posterior of the brain in a region known as the visual cortex (or V1). At this stage of processing, the simplest visual features of the stimulus are automatically detected and the results passed to more anterior regions of the brain that develop higher-level representations. There are actually two streams of visual processing (Schneider, 1969). A dorsal stream is responsible for locating the stimulus in the visual field and a ventral stream is responsible for identifying the stimulus (Goodale & Milner, 1992). These two visual representations – the where and the what – are found in the visual sensory register from which the short-term store can selectively process in a controlled fashion.

With respect to the neural basis of the visual short-term store, much of the research began with animal studies. A prominent task that researchers use is known as the *delayed-match-to-sample task*. A good example of how this task is utilized was reported by Goldman-Rakic (1988), who studied the visual short-term store in monkeys. A monkey was placed in a cage and shown two wells placed outside of the cage. Food was then placed in one of the wells before both wells were covered. A screen was then drawn down between the cage and the wells. A short time later, the screen was raised. The goal of the monkey was to remember the location of the food. This was an easy task for monkeys with an intact brain. However, monkeys with a lesion in the prefrontal cortex known as Area 46 did not remember the location of the food.

Groundbreaking research on the human visual short-term store was made possible by the development of noninvasive neuroimaging technology in the late twentieth century. One such technology is positron emission tomography, or PET, which involves the injection of radioactive isotopes into the bloodstream, which are then concentrated in parts of the body, including the brain, where oxygen in the blood is required. Carefully conducted experiments measure blood oxygen-level differences in the brain to ascertain what parts of the brain are associated with the performance of simple cognitive tasks. In one experiment by Jonides et al. (1993), human subjects viewed a series of displays. In one condition, visual fixation was followed by the presentation of a circle, followed by a short delay, and then the presentation of a dot. The subject's task was to determine whether the circle contained the dot. This very easy task is strictly a visual task that relies only on the visual sensory register. In the memory condition, the presentation of the dot was followed by its removal, and after a short delay, the circle was presented. The subject's task was to determine if the area defined by the circle previously contained the dot. But note that this task requires the subject to maintain a short-term memory for the location of the dot. When the brain activity associated with the first, simple task was subtracted from the brain activity associated with the task that also required the use of visual short-term memory, a heightened state of brain activity in Area 46 of the prefrontal cortex was revealed. The finding is consistent with earlier animal studies that suggest Area 46 of the brain supports visual short-term memory.

Additional Sensory Registers. Over the ensuing years, evidence for the existence of additional sensory registers emerged. Broadbent's research suggested the existence of an auditory sensory register where early representations are of the sensory features of acoustic stimuli. For instance, one might selectively process auditory stimuli presented to a given ear. The speculation bore fruit when Darwin, Turvey, and Crowder (1972) used methods inspired by Sperling to investigate the nature of the auditory sensory register. Their results were quite similar to those obtained by Sperling, with one important difference. Whereas the contents of the visual sensory register decayed within a few hundred milliseconds, the contents of the auditory sensory register took much longer to decay, which may reflect the sequential nature of the auditory stimulation and the parallel nature of visual stimulation used in Sperling's experiments. Likewise, similar findings concerning a potential tactile sensory register suggested that tactile stimulation was processed in parallel and did not require attentional resources (Craig, 1968; Shiffrin, Craig, & Cohen, 1973).

Automatic and Controlled Sensory Processing. The topic of parallel versus serial information processing is fundamental to research within the framework of the General Theory. At a high level of analysis, research on serial versus parallel processing of information addresses the extent to which one controls information processing. Parallel processing refers to simultaneous identification, use, or transformation of multiple memory representations. Such processing is usually assumed to take place *automatically* or without the need for attentional recourses. For instance, when scanning the produce aisle for bananas, the visual sensory register automatically maintains a brief representation of the colors of the fruits and vegetables on display, and attention is directed with no effort to those that are yellow. In contrast, serial processing requires the allocation of attention to individual items and is therefore more closely associated with the operations of the short-term store. Once the yellow items of the produce aisle are noticed, attention can be directed toward the processing of lemons, squash, and bananas in a serial fashion to discriminate or classify the yellow items, ensuring that one does not put squash on their cereal in the morning. The distinction between automatic and serial processing is potentially relevant to many other activities. There has been a great deal of research conducted over the past few decades on this topic, and we will review some of the important advances made in the areas of the sensory registers (Chapter 3) and the long-term store (Chapter 4).

Short-Term Memory Stores

One of the more controversial assumptions of the General Theory is the assumption that the short-term store and the long-term store are distinct. The assumption goes back to the earliest writings in psychological science (James, 1890) and later to Broadbent (1958). James's original proposal was based on logic and experience rather than on the results of well-controlled experiments. It seemed obvious to him that the momentary contents of thought were easily accessible, whereas the recollection of distant events was often more difficult. Later, empirical research by Brown (1954) and several others helped revive research on short-term memory, and their results influenced Broadbent's dual-store assumption, which was similar to James's. Likewise, the General Theory assumes that different memory structures serve different purposes: The short-term store temporarily keeps information highly accessible for immediate use, as long as it is attended to, whereas the long-term store retains permanent records of experience and knowledge. However, there is another critical structural distinction: The short-term store temporarily retains a limited amount of information, whereas the long-term store permanently retains memory traces corresponding to the events that we have experienced throughout our lifetime.

Beyond these assumptions, the dual-store approach was greatly expanded upon by Atkinson and Shiffrin (1968), who proposed that the flow of information between the memory structures is under the control of the subject and long-term memories represent what has been processed in the short-term store. Hence, although the short-term store is functionally distinct from the long-term store, they work in concert with each other. The flow of information between the short-term store and the long-term store is revealed by many everyday activities. From a student's point of view, we have all had the experience of studying material for an exam and feeling it is well learned, only to discover later that it was not as well learned as we thought. Within the framework of the General Theory, current thoughts are representations that reside in an active state in the short-term store. While studying, the contents of the short-term store represent what we are actively studying, readily available for taking notes, discussing with peers, integrating with prior knowledge, and so on. For instance, when studying vocabulary, we obviously attend to a new word and the meaning of the new word. In that moment, this short-term memory is consciously experienced.

Because items in the short-term store are easily accessible, we might incorrectly assume that they will also be easily accessible in the future. But when taking an exam, say, the next day, access to the items studied the night before requires successfully retrieving that information from the long-term store, which may not be as easy to do. One reason for this is that the items are in the short-term store when they are being studied but not when taking a test. That is, testing usually requires retrieving information from the long-term store, but this is more difficult than retrieving information from the short-term store. Retrieved information from the long-term store is transferred to the short-term store. The contents of the short-term store are then readily available for answering an exam question.

The tip-of-the-tongue feeling (Brown, 1954) also reflects the distinction between short-term and long-term memory stores. Upon sitting an exam, a topic is tested, which the student distinctly recollects studying the night before, but the answer cannot be put to paper. Instead, it feels like the answer to the test question is "on the tip of the student's tongue." Tip-of-the-tongue feelings reflect an awareness that information is in long-term memory but cannot be successfully retrieved. For instance, although an item cannot be retrieved from the long-term store, when in a "tip-of-the-tongue" state, one is often able to correctly determine the first letter of the item. In this case, only partial information from the long-term store has been retrieved and represented in the short-term store (cf. Schacter & Worling, 1985). Following the exam, the student may consult the textbook or notes and is reminded of the correct answer, or the answer spontaneously pops into the student's mind later during the exam. This confirms the suspicion that the tip-of-the-tongue information was in long-term memory all along. This is a common source of student frustration, and it is consistent with the dual-store model.

Neural Basis for Short-Term Stores. Few topics have received more attention from memory researchers than the dual-store assumption, including neuroscientific case studies, such as those obtained from patient Henry Molaison (H.M.) by Scoville and Milner (1957). These findings heavily influenced the General Theory, and they are worth describing to establish an important empirical basis for the dual-store model. Before doing so, we acknowledge the limited generalizability of findings obtained from case studies, as brain traumas in different individuals may not be the same. But the following findings have subsequently been observed in human and nonhuman patients following H.M., and the interpretation of these findings is generally accepted among memory researchers.

H.M. was a male introduced to Milner and Scoville in the 1950s. As a child, H.M. was in a bicycling accident that precipitated severe epileptic seizures. At the time, medical treatment options for seizures were rudimentary at best, and H.M. had a bilateral resection of parts of his medial temporal lobe including, but not limited to, the hippocampus, an area of the brain now thought to be crucial to the formation of long-term memory traces. The treatment was a success insofar as the seizures were concerned, but there was a significant side effect: H.M. lost the ability to remember what he experienced following surgery.

Larry Squire (2009) reviewed the literature on H.M. shortly after his death, and came to the following conclusion:

> H.M.'s intact intellectual and perceptual functions, and similar findings in other patients with large medial temporal lesions, have been well documented. A key additional finding was that H.M. had a remarkable capacity for sustained attention, including the ability to retain information for a period of time after it was presented. Thus, he could carry on a conversation, and he exhibited an intact digit span (i.e., the ability to repeat back a string of six or seven digits). Indeed, information remained available so long as it could be actively maintained by rehearsal. For example, H.M. could retain a three-digit number for as long as 15 [minutes] by continuous rehearsal, organizing the digits according to an elaborate mnemonic scheme. Yet when his attention was diverted to a new topic, he forgot the whole event. In contrast, when the material was not easy to rehearse (in the case of nonverbal stimuli like faces or designs), information slipped away in less than a minute. These findings supported a fundamental distinction between immediate memory and long-term memory (what William James termed primary memory and secondary memory).
>
> Notably, time is not the key factor that determines how long patients like H.M. can retain information in memory. The relevant factors are the capacity of immediate memory and attention, i.e., the amount of material that can be held in mind and how successfully it can be rehearsed. The work with H.M. demonstrated that the psychological distinction between immediate memory and long-term memory is a prominent feature of how the brain has organized its memory functions. (pp. 3–4)

Hence, H.M.'s short-term memory store was spared, but his long-term memory was not. Such selective impairments are not limited to those caused by brain injury. In fact, impairment of the hippocampus need not

 https://doi.org/10.1017/9781009440370.003

be permanent to affect long-term memory. For instance, midazolam is a benzodiazepine that produces dense but temporary impairment of the encoding of long-term traces, and subjects behave much like H.M. while under its influence (Polster et al., 1993).

More recently, noninvasive technology has allowed brain activity to be measured in healthy subjects. Using PET, functional magnetic resonance imaging (fMRI), and other devices, researchers proposed that memory is supported by a network of brain areas that has become known as the default mode network (Gusnard & Raichle, 2001; Raichle et al., 2001). A common question is what part of the default mode network is responsible for short-term memory. In addition to the medial temporal lobe, prefrontal cortical areas, the posterior cingulate cortex, and the angular gyrus are thought to be involved in human memory, and Frank, Loughery, and O'Reilly (2001) wrote, "It is almost universally accepted that the prefrontal cortex (PFC) plays a critical role in working memory, even though there is little agreement about exactly what working memory is or how *else* the prefrontal cortex contributes to cognition" (p. 137). The prefrontal cortex comprises several distinct subareas, and it seems unlikely that no single part of the prefrontal cortex is responsible for short-term memory.[4]

[4] To illustrate the complexity of the prefrontal cortex, consider the results of just one influential experiment. Nyberg, McIntosh, and Tulving (1998) found increased activation in the left prefrontal cortex during semantic retrieval and increased activation of the right prefrontal cortex during the retrieval of episodic memories. In addition, Nyberg et al. reported heighted activation of the left prefrontal cortex during episodic encoding, indicating that the left prefrontal cortex has a functional role in both encoding episodic memories and retrieving semantic memories. It is possible the hemispheric heterogeneities reflect their different roles in short-term memory. For instance, the left prefrontal cortex might support short-term memory representations to be encoded in long-term memory and used as retrieval cues, whereas the right prefrontal cortex might be involved in controlling these representations. In any case, short-term memory is likely to depend on the interaction of several brain areas, as Frank et al. suggest. It is almost certain that different parts of the prefrontal cortex are involved in attention, representing information, and controlling the processing of information. In addition, there are reports that the support of short-term memory is not limited to the prefrontal cortex (Jonides et al., 1998).

It appears that visual short-term memory is at least as complex as verbal working memory (cf. Goldman-Rakic, 1992). For instance, Smith and Jonides (1995) found that visual short-term memory was supported by Area 46 in the prefrontal cortex, but a more recent review suggests that visual short-term memory is supported by a network including parietal brain areas (Sheremata, Somers, & Shomstein, 2018). It is possible that different parts of the network are involved in attention, representation, and control. Not surprisingly, the situation becomes more complex when autobiographical memory is considered because it involves a mixture of event memory and knowledge (as discussed in Chapter 4). Burianova, McIntosh, and Grady (2010) posited that the overlap in prefrontal activations occurring during the performance of episodic and semantic memory tasks is evidence for a common functional brain network supporting not only episodic and semantic memory but also autobiographical memory. On this basis, one might hypothesize that prefrontal areas play an instrumental role in constructing a mental model to be used as a retrieval cue.

The General Theory's specification of short-term memory greatly oversimplifies its implementation in the brain. Yet, the concept of short-term memory within the framework of the General Theory is still useful. As Marr (1982) famously proposed, there are different levels for understanding information-processing systems, with each level serving a purpose. At the highest level, there are models of findings or the problems to be solved. These are sometimes referred to as task analyses, which are proposals for what the information-processing system must accomplish to achieve a goal. For instance, given the task of remembering a series of digits, one proposal is that these digits need to be kept in a readily accessible state. The middle level of understanding specifies an algorithm to achieve this goal. Hence, one might propose a model for rehearsal specifying the number of items to be rehearsed and in what order. Last, one can specify the physical mechanism that implements a model, and one approach is to study the brain and identify what parts are involved in rehearsal and how they are interconnected. At this time, there is no research that disconfirms the notion that different memory structures exist, and most assume they have the essential properties described by the General Theory. Hence, the identification of brain networks that support short-term memory is likely to continue to be a goal of neuroscientists, and the General Theory can serve as a means for organizing this brain research.

The Short-Term Store and Higher-Level Control. Atkinson and Shiffrin's proposal that information is transferred from the short-term store to the long-term store was a novel contribution to memory research, but the proposal that one may control how information flows between memory structures was revolutionary (Atkinson & Shiffrin, 1971). All conscious mental activity takes place in one or more short-term stores according to the General Theory, and cognitive control processes manage what, when, and how information is used in the short-term store. The items comprising a thought may have been selected from the contents of a sensory register and/or retrieved from the long-term store. While attended to, regardless of their origin, items will remain in a short-term store. When tasks change or one is distracted, our thoughts also change because the contents of short-term memory stores change. Hence, humans can control what they think about.

The theoretical construct of cognitive control is closely aligned with another important assumption of the General Theory: short-term memory stores have a limited capacity. This should be obvious to everyone: it is difficult to think about more than one thing at a time. For instance, the contents of the short-term store would differ depending on what I am thinking about. If I were daydreaming about being on a boat on the ocean versus doing an algebra problem, the contents of the short-term store would reflect these different thoughts. Indeed, daydreaming about being on a boat would distract me from doing the algebra problem because of the limited capacity of the short-term store. Hence, the assumption that the short-term store has a limited capacity indicates a need to control the contents of current thought. And a limited capacity short-term store requires choices about what information it contains and what it does not contain at a given moment.

Cognitive control is intimately linked to capacity limitations. It stands to reason that it is difficult to accomplish any cognitive task without attending to the items necessary to perform the task, and there is abundant research indicating that the ability to devote attention to the processing of information in the short-term store is required to solve novel problems, plan for and imagine the future, and reflect on the past. Regardless of the task, the ability to efficiently utilize a limited capacity memory requires one to control its contents. Indeed, many intelligence and neuropsychological assessments depend on measures of an individual's capacity to maintain items in the short-term store as a means of detecting individual differences in intelligence or cognitive impairments (Conway & Engle, 1996).

Finally, returning to the dual-store framework, managing the contents of the short-term store has consequences for long-term memory. Attending to the contents of the short-term store transfers this information to a long-term episodic trace, which allows this event to be retrieved from the long-term store in the future. When attention is divided or taxed by some other mental activity, the transfer of the contents of the short-term store is disrupted, making it more difficult to remember the event in the future. Hence, managing the contents of the short-term store affects not only current thought but also what will be remembered and thought about in the future.

Goals. The goals one has when performing a task determine the contents of the short-term store. Again, this is quite intuitive: when planning a trip to the ocean, our thoughts may contain items such as a beach towel, swimming suit, sunglasses, and so on, whereas if we are doing algebra problems, our thoughts contain items such as numbers, variables, and mathematical operators. Items define thoughts related to goal-driven behavior. However, there are an infinite number of possible goals, and since we cannot read someone's mind (at least not yet), one challenge for memory researchers interested in short-term memory is to make inferences about what the goals of the subject are.

There are methods for inferring the goals of the typical subject, but for now it is worth briefly considering two types of goals: normative goals and subjective goals (Malmberg, 2008). A normative goal is met when the subject performs whatever task that is assigned in an objectively optimal manner. The optimal level could be determined by a mathematical model. For instance, ideal observer models are commonly used in research (Green & Swets, 1966). Many times, however, one determines the level of performance or behavior that satisfies a personal goal, considering the time spent on a task, motivation from self or others, and environmental stress, among other factors. For instance, a subject who arrives at the lab stressed about an upcoming exam, family situation, or physical condition may perform the free recall task with a different subjective goal than one with fewer stressors. On the assumption, therefore, that goals motivate behavior, inferring the goals of the subject is critical to an understanding of what, how, and when information is processed.

It is often assumed in the laboratory that instructions provided by the experimenter are a proxy for the goal of the subject (cf. Schnorr & Atkinson, 1969). For instance, in a free recall experiment, the subject is instructed to study a list of items and recall as many as possible in any order when prompted. Such instructions are commonly used, but they leave much to the imagination of the subject. How many items does the subject *want* to recall? Presumably, most subjects will try to recall as many items as possible, but what is possible may differ between subjects depending on their subjective goals and can change during the task. When trying to learn the items, should the subject rehearse the items as many times as possible or create a sentence or visual image of the items? How long should the subject attempt to recall words before giving up? These are just some of the questions one might ask, and the answers may depend on the motivation of the subject, the situation, experience with the task, and so on.

In the real world, however, there are usually no instructions provided. Oftentimes, we find ourselves with subjective choices about what is important to attend to and what is not. For instance, when driving to the store in a familiar setting, one might choose to attend to a conversation on the radio. But in a novel setting one may decide that attending to the road or directions is more important. Likewise, when studying for an exam, there are choices to be made. It makes little sense to study material that one has already learned, and one should focus on material that is not well learned. Hence, the goals of the subject direct attention to items or information that are used to meet a goal, and information that is attended to in the short-term store is more strongly transferred to the long-term store.

Experiments in which the goals of the subject are directly manipulated have highlighted how goals impact what is attended to and later remembered. A well-known experiment was conducted by Anderson and Pichert (1978). They had subjects read a story about a house and they were instructed to imagine being either a homebuyer or a burglar. After a delay, long-term memory for the items in the house was tested via free recall. Subjects who imagined being a homebuyer tended to recall more facts about the house relevant to homebuyers (e.g., new exterior siding), whereas those who imagined they were a burglar tended to recall more items relevant to burglars (e.g., the side door is always unlocked).

A straightforward way of interpreting this result is that homebuyers and burglars have different goals, and they attend to different items when analyzing the contents of a house. Hence, attention is directed to those most relevant to their goals while perhaps ignoring items that are irrelevant to one's goals. It is important to note that the story was the same in both conditions; only the subjects' goals were different. And regardless of the specific goal, a mental representation of the passage occupied the short-term store and was constantly being updated as the subject read the passage. The result is that goal-relevant information is more strongly transferred from the short-term store to the long-term store and is therefore easier to recall later.

Rehearsal. Although the contents of the short-term store continuously change as our thoughts change, at times it is important to maintain information in the short-term store. When you go to the kitchen, for instance, it is frustrating when you forget why you went there. Today, many smartphone or computer applications require multistep authentication to maintain the integrity of a computer network. Under ideal conditions, entering your password is easy, as it has been memorized via repeated retrieval from the long-term store. When you are texted a random security code, however, you may need to repeat the elements of the code in your head to enter the code.[5] This mental repetition of the security code is an example of rehearsal, and according to the General Theory, the more often one rehearses an item, like a password, the more effectively it will be transferred to a permanent long-term trace.

The term rehearsal is metaphorically borrowed from the performing arts to emphasize its role in learning or memorizing. What is rehearsed is under

[5] During the writing of this book, smartphone applications began automatically copying the security code so that one does not need to rehearse it. An interesting question is whether such cognitive support by computer devices will impact (positively or negatively) cognitive control processes. For instance, it may be that the control processes invoked fifty years ago to perform a task may differ from those used fifty years from now to perform the same task.

one's control and depends on task demands or the goals of the subject. For a free recall task, subjects are asked to recall as many items from a list as possible in any order, and it is often assumed that subjects attempt to keep as many recent items in the short-term store as possible. Since these are "in mind" and readily available, they are quite likely to be recalled when memory testing begins, especially compared to those that are no longer in the short-term store. However, the longer an item is rehearsed, the more information about the item is transferred to long-term memory. Hence, there are benefits to rehearsal even when one must retrieve a memory from the long-term store.

The Long-Term Store

Perhaps the most commonly misunderstood aspect of human memory is forgetting, a phenomenon associated with the long-term store, and research over the last fifty years has been devoted to developing various models to better understand how long-term memory supports human cognition. Everyone has difficulty remembering the details of prior experience, whether it refers to a particular event or knowledge derived from learning or practice. Difficulty retrieving from long-term memory may lead one to conclude that the memories have been forever lost. However, human memory research indicates otherwise. Failure to successfully retrieve a memory does not mean that the relevant trace has been lost. In fact, the General Theory assumes that long-term memory traces are permanent. An obvious question concerns why retrieval of long-term memories is sometimes difficult if they are permanently retained.

Permanence of Long-Term Memories. Over the years, many models of long-term memory have been developed to describe forgetting in different situations and for different tasks, yet a small number of assumptions underlie those models. As alluded to in the prior paragraph, long-term memories are assumed to be permanent. Forgetting occurs because memories in the long-term store are available but difficult to access. The distinction between the availability or existence of long-term traces and their accessibility was established in the very first experiments on human memory. Ebbinghaus (1964) showed that, although previously well-learned items are sometimes unable to be retrieved from the long-term store, they can nevertheless be relearned much more quickly than new material. The fact that relearning of the item is relatively fast indicates that the original long-term traces are still available, but simply inaccessible. More recently, this finding has been extended to indicate that memories previously acquired almost seventy years prior are still available in the long-term store (Maxcey, Shiffrin, Cousineau, & Atkinson, 2022). Thus, research on human memory indicates that long-term memories are lost like a cell phone is temporarily lost; they are still in the long-term store waiting to be found even decades after they were originally acquired.

Other phenomena that support the distinction between availability and accessibility are commonplace. Almost all students realize that multiple-choice exams are easier than short-answer or essay exams. In fact, the inability to answer a short-answer question does not imply that one would not be able to correctly answer a multiple-choice question (even after taking guessing into account). Thus, forgetting is rather transient; some queries of the long-term store may be successful, even though others may not be.

Retrieval Cues. According to the General Theory, the ability to retrieve from long-term memory depends on the availability of effective retrieval cues. Retrieval cues are information contained in short-term memory; they are what one thinks about while attempting to retrieve information from long-term memory. From a student's perspective, when testing via multiple-choice questions, specific cues are provided in the form of the question and the possible answers. But when memory is tested via essay or short-answer questions, only a question is provided, and the student must retrieve the answers. Multiple-choice tests are easier because the student is presented with a comprehensive set of retrieval cues: the question and the answer. When retrieval fails, "hints" are additional retrieval cues that can help jog a student's memory. This is probably the most obvious example of "cue-dependent memory" and makes clear the assertion that accessibility of long-term memories depends on the cues used to query them.

In the laboratory, different memory tasks are largely defined by the cues provided to the subject by the experimenter. Like different forms of academic tests, not surprisingly, some laboratory procedures are easier than others. In other words, the amount of forgetting observed in the laboratory depends on the manner memory is tested. For *recognition*, an item that may or may not have been studied is provided by the experimenter (i.e., a true/false memory test), and the subject's task is to only endorse those that were previously studied in a specified context. The context is usually assumed to be a specific study list (e.g., the last list). Hence, the retrieval cues used to probe memory for recognition are the item provided and the context. If the subject recognizes the test item as having been studied on the recent list, the answer is "yes" or "true." *Cued recall* follows studying pairs of items. When a similar set of cues is provided, one item from a pair and context, the subject must retrieve from long-term memory the item with which it was paired.

Subjects typically find cued recall more difficult than recognition. *Free recall* is even more difficult. It requires subjects to recall as many items from a study list as they can in any order they choose. It is particularly difficult because only a context cue is used to query memory for items from "the last study list." Once a studied item is retrieved, it can then be used in combination with the context cue provided by the experimenter to retrieve another item in a manner similar to cued recall. Hence, the ease of a memory task reflects the accessibility of a trace in the long-term store given the set of retrieval cues the subject has in their short-term store.

Interference. However, according to the General Theory, there is another source of forgetting. Everyone knows that it becomes more difficult to remember a list of items as the number of items to be remembered increases. The "list-length" effect is also observed in the laboratory (even after potential confounds are controlled) when subjects forget a greater proportion of items from lists containing a greater number of items than from shorter lists. According to the General Theory, the list-length effect is caused by interference from traces competing to be retrieved. During retrieval, many traces in the long-term store are accessed via the set of retrieval cues to query or probe memory. In principle, any one of the traces in the long-term store might be retrieved, but those that are most similar to the set of retrieval cues are most likely to be. Moreover, the more traces there are that are similar to the retrieval cue, the less likely it is that one of them will be successfully retrieved. Fortunately, most traces in the long-term store contain information that does not match the context used as a retrieval cue, and therefore they cause very little interference. However, one can only retrieve one trace at a time, and those traces with context that matches the retrieval compete to be retrieved. This competition increases as the number of competing traces increases. An apt analogy helps visualize interference that results from retrieval competition: It is easier to rebound a basketball against a team with one player than a team with five players because there are fewer competitors to negotiate. As the number of competitors increases, the probability that any one of them will rebound the basketball decreases, but the probability that any one of them will retrieve the ball increases. The details of this probabilistic retrieval process will be described in the next chapter. But for now, the increase in forgetting with an increase in the number of traces in the long-term store is due to competition or interference from traces in response to a query.[6]

[6] Note that as the contents of the retrieval cues change, so do the matches between them and contents of the traces in the long-term store. Hence, interference and forgetting can be reduced by changing the retrieval cues to probe memory, which is consistent with the assumption that forgetting does not imply a permanent loss of traces in the long-term store.

The Interaction between the Short-Term Store and the Long-Term Store

To this point, the different memory stores have been treated as distinct entities. And they do have different properties and functions in human cognition. However, the various components of the memory system work in concert to perform a cognitive task by implementing control processes to manage information flowing among the memory structures (Shiffrin & Atkinson, 1969; Atkinson & Shiffrin, 1971). The processes involved in the transfer of information have been extensively investigated, but before turning to the "nuts and bolts" of the various models (in Chapter 2), let us consider how the General Theory assumes the memory system operates in concert to understand an historically important finding: the free recall serial-position function.

As mentioned earlier, free recall is a laboratory task in which subjects usually study a list of items (e.g., words) one at a time and are subsequently instructed to recall as many items as possible from the study list in any order they choose. A serial-position function plots the probability of correct recall as a function of the order in which the items were studied. For instance, one might compare the ability to recall the first item on the study list to the ability to recall items from other serial positions. Figure 1.1 shows the serial-position functions obtained by Glanzer and Cunitz (1966). Subjects studied fifteen-item lists and there were three delay conditions; memory was tested immediately, 10

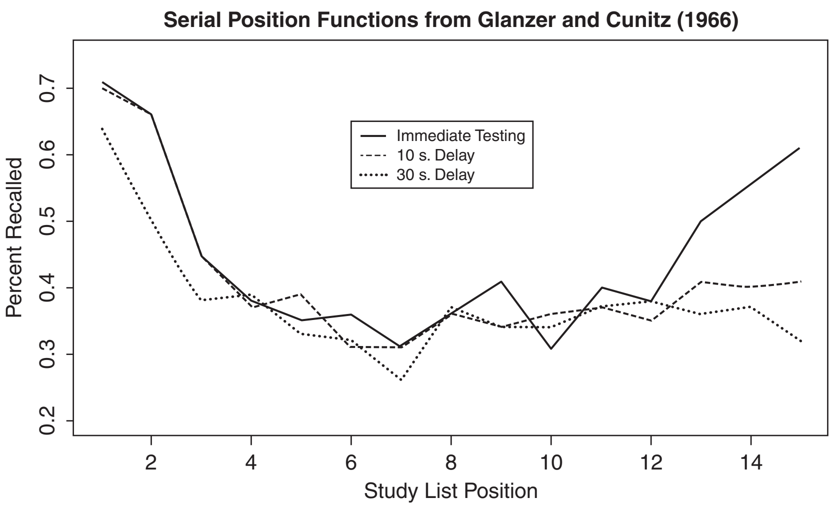

Figure 1.1 Typical serial-position curves obtained from a free recall procedure as a function of the duration of a filled retention interval.

seconds after, or 30 seconds after presentation of the final item. In all conditions, free recall is greatest for the first item and decreases until about the third item on the study list. This is known as the *primacy effect*. Immediate free recall was also better for the last two or three items on the study list. This is known as the *recency effect*. The recency effect was not observed following a 30-second delay in free recall testing.

The General Theory predicts a primacy effect based on the dual-store assumption. The first items on the study list are rehearsed longer in the short-term store than items that appear in the middle serial positions or toward the end of the list. The prediction falls out of the assumption that the short-term store has a limited capacity and that subjects rehearse items until a capacity limit is reached; because the first items were rehearsed longer, they are more easily retrieved from the long-term store.

We can refer to the set of items concurrently rehearsed as a *buffer*. At the point where the buffer reaches full capacity, one of the items currently in the buffer is removed and replaced by the next item on the study list. Different task demands may favor different strategies for controlling the contents of the buffer when the capacity is reached, but a simple model is a "first in, first out" stack algorithm, whereby the first item to enter the rehearsal buffer is the one to be removed. In fact, this is what it appears subjects do. In a classical study, Rundus (1971) asked subjects to rehearse aloud, and the items from early serial positions on the study list were given more overt rehearsals than those that followed. By adopting this strategy, the first items on the study list are transferred more strongly to the long-term store and are easier to retrieve in the future.

Subsequent to a point when capacity is reached, items are rehearsed for about the same amount of time, and therefore they are transferred to the long-term store about equally well. But those items are not as likely to be retrieved from the long-term store as those from the initial serial positions because they are not rehearsed as long. This is because the oldest item is removed on each study trial once the buffer capacity is achieved. Hence, retrieval from the long-term store is about equally likely for items studied in the middle serial positions (e.g., serial positions 4–12 in Figure 1.1), but not as well as those in the primacy portion of the serial position function.[7]

[7] This account of the primacy effect may seem logical or even trivial in the context of the General Theory, but it is very important to consider just how counterintuitive the primacy effect is. Everyone knows that memory for more distant events is worse than memory for recent events. The primacy effect for list learning, therefore, runs counter to our everyday experience. The General Theory account of the primacy effect is a great example of the role scientific models play in our understanding of the world.

Figure 1.1 shows that the recency effect depends on the delay (or retention interval). When free recall is delayed, and the retention interval is spent performing a task that requires use of the short-term store, items from the study list must be removed from the buffer so the limited capacity short-term store can be used to perform tasks assigned during the retention interval. Therefore, free recall following a delay is based solely on retrieval from the long-term store. In this case, no recency effect is observed because all items tend to be rehearsed for the same amount of time once the buffer is full. When there is no retention interval, as in immediate free recall, the contents of the buffer remain, and the recency effect observed for immediate free recall is attributed to their retrieval.

The dual-store account of the free recall serial-position function is supported by many findings. For instance, primacy and recency effects are selectively impaired in patients suffering from brain damage, depending on the location of the damage (Crocket, Hadjistavropoulos, & Hurwitz, 1992). H.M. had an intact recency effect, suggesting that his ability to rehearse items was unimpaired, but his memory for the items from earlier serial positions was greatly impaired (Milner, 1998). A noteworthy caveat is that damage to the brain is rarely identical in different patients. Yet, other human patients and even animals produce a serial-position curve similar to H.M.'s when their hippocampus is damaged (Baddeley & Warrington, 1970; Kesner & Novak, 1982).

Concluding Remarks

This chapter chronicled the fundamental assumptions of the General Theory. The remaining chapters cover the accounts of various memory phenomena in detail. But there are too many models to cover, and therefore the focus is on several models that illustrate the central assumptions of the General Theory and how they can be implemented in a formal structure. Chapter 2 covers the implementation of the basic assumptions of the General Theory and its modeling frameworks. It is understood that this material will be challenging at first. It is, perhaps, a good strategy to initially grasp the gist of these modeling frameworks rather than focusing on the details of each model. Later, Chapter 2 can serve as a reference for the various models and how to modify them to serve a new purpose and to better understand the material in subsequent chapters.

Chapters 3 and 4 illustrate how the framework of the General Theory relates several memory tasks that on the surface may appear to have little in common. For instance, research on visual search and recognition memory shares a goal of understanding how behavior sometimes relies on automatic processing of information and sometimes relies on more deliberate serial processing. Chapter 4 focuses on influences on memory that are largely uncontrolled, reflecting basic structural aspects of memory, which give rise to forgetting, and Chapter 5 presents several models of forgetting.

Chapter 6 covers several findings that were problematic for some models and how those models evolved to account for them. A key assumption that emerged is that long-term memories become differentiated. Differentiation asserts that as more information is accumulated about two randomly similar events, the traces representing those events become easier to distinguish. Differentiated memory traces result from the accumulation of information in existing traces from repetitions of items or events that are highly similar. The tantalizing suggestion is that trace accumulation may be the basis for inductive learning that results in knowledge derived from experience. Chapter 7 presents several models of lexical access, long-term priming, and learning that emerged from research focused on the interaction between experience and knowledge.

Chapter 8 presents how memory is affected by testing memory and some principles of model development. It demonstrates how findings that seem difficult to understand when the consequences of memory testing are not considered make perfect sense when the effects of prior memory tests are acknowledged. Of course, the General Theory has not been without its critics over the years. Although some of these criticisms are the result of misconceptions, some have persisted and entered into conventional wisdom. In Chapter 9, I discuss how several findings thought to challenge the General Theory are actually a priori predictions of the family of buffer models originated by Atkinson and Shiffrin.