1 Overview of how to analyze memory tasks

As a starting point in thinking about human memory, we introduce a phenomenon referred to by Johansson, Hall, Sikström, and Olsson (2005) as Choice Blindness. In the original study participants were shown two cards, each containing a picture of a female face, and were asked to choose the more attractive of the two faces. On some trials, immediately after making their choice, they were shown the picture that they had chosen and were asked to describe their reasons for choosing that picture. Unbeknown to the participants, on a few of these trials, a card trick was used to show them the picture of the face that they had not chosen. There were several variations on this basic condition. Across all conditions, only 13% of the deception trials (when the non-chosen face was presented as the chosen face) were identified, though some participants indicated on a subsequent questionnaire that they had some suspicions. These results were published in the journal *Science*, so they were clearly judged to have met exacting standards for scientific importance and general interest. Although many people were surprised by the results, surprise is not necessarily an indicator of scientific importance or good science. Instead it might simply reflect the fact that the reviewers of this article did not understand how to perform a preliminary analysis of an unusual memory task.

The Choice Blindness task can be used to illustrate five components involved in analyzing a memory task: (1) What are the *goals* of the task? (2) What *cues* are used? (3) What *information* is needed in order to solve the task? (4) What were the *opportunities* to learn the required information? and (5) What are the sources of *noise* in the memorial process?

The first step in the analysis of the Choice Blindness task is to identify the *goal* of the task. From the researchers' perspective the goal was to identify the deception. That is, to recognize that the experimenter-supplied face was not the face that the participant had chosen. However, this was probably not the goal of the participants, who presumably thought that they were reporting on their reasons for choosing the presented face.

The *cues* available on the deception trials include the experimenter-supplied face that was not chosen; the concept "chosen"; and temporal information. By temporal information we mean the information that indicates that the experimenter's question refers to the last trial. An implicit assumption here is that this information plays a functional role in the ability to recall details about the last trial. In Chapter 3 we will argue that it is useful to conceptualize the use of temporal information as a temporal context cue. In the Choice Blindness task an attempt to recall the last trial would presumably involve recalling memories for both the chosen and non-chosen face and a memory for having chosen one of the two faces. Information about which face had been chosen might also be remembered, though this might involve a second retrieval process using a memory of one of the retrieved faces as a cue.

With respect to the *information* required to achieve the experimenter's goal, one possibility is an association between the non-chosen face (or some of its features) and the concept "not chosen." Because the non-chosen face is physically present we might think of this as a cued-recall task where the participant uses the presented face as the cue to recall the concept "not chosen." In principle, the experimenter's goal can also be achieved if there is a sufficiently strong association between the chosen face and the concept "chosen." However, to achieve the goal by this route, the participant would have to retrieve an image of the chosen face, or at least sufficient detail that it could be discriminated from the presented, non-chosen face. Temporal information and the concept "chosen" are both potential cues for retrieving an image of the chosen face. In addition, the paired presentation of the two faces during the choice task along with the participant looking back and forth and comparing them on specific features could create an association between them. Thus the presented non-chosen face could conceivably also be used as a cue to recall the non-presented chosen face. However, recalling sufficient detail about the chosen face so that it can be differentiated from the presented non-chosen face may be difficult because the former had only been observed for a few seconds. An alternative explanation for the failure to recall the chosen, but not presented, face is that because of the different goals involved (due to the deception), participants may simply not attempt to recall it.

The next step in our analysis asks how much *opportunity* is there to learn an association between the non-chosen face and the concept "not chosen." Presumably, during the choice task, the participant looks back and forth between the two faces either trying to make a holistic judgment or by comparing them on the basis of individual features. Regardless, it doesn't seem possible to form an association between the non-chosen face and the concept "not chosen" until after the choice is made. At that point, attention is likely to be directed toward the chosen face, and there may be relatively little opportunity to learn the association between the non-chosen face and the concept "not chosen."

The final step in our analysis is to identify the sources of *noise* that might cause a participant to falsely believe that they had chosen the non-chosen face.

By noise we refer to the discrimination problem faced by the memory system and whatever makes that discrimination difficult. Johansson et al. (2005) report on a comment made by one of their participants during the test phase as to why he had chosen a face that he had in fact not chosen during the initial choice task. The participant commented that he "preferred blondes" even though in the choice between a blonde and a brunette he had chosen the brunette. This comment seems irrational but it may be an indicator of a source of noise in the task. It seems likely that when participants look at an individual face they notice features that they like such as the eyes, hair color, hairstyle, facial symmetry, etc. They may even comment to themselves about liking that feature. For example, when choosing between a blonde and a brunette they might comment to themselves that they prefer blond hair. However, it still may be the case that they end up choosing the brunette because they judged the brunette to be more attractive in more of the feature comparisons. Now when presented with the picture of the non-chosen blonde the participant may recollect that during the choice task they had commented to themselves that they liked blond hair. However, they do not recollect any comments about liking the eyes, the hairstyle, or the facial symmetry of the other (chosen) face. A possible reason for this is that the face they are looking at cues the comments they made about that face but not the comments they made about the other face. Given all the work on the role of recollection in reality monitoring (did you see this stimulus or did you imagine it?) and source monitoring (did you see this as a picture or as a word?), it seems likely that such a recollection would be taken as evidence by the participant that they had indeed chosen the blonde (Johnson, Hashtroudi, & Lindsay, 1993).

Gold CIONs: goals, cues, information, opportunity, and noise

To summarize what we think are the implications of our analysis of the Choice Blindness task, we suggest that the best starting point for thinking about human memory is a task analysis. Such an analysis begins with an attempt to identify the goals of the task. Note that the goals may not be the same for the experimenter and the participant. The analysis continues with an attempt to identify the cues that are available. These cues may be inherent in the situation (e.g., temporal cues when the test shortly follows the study episode), they may be supplied by the experimenter (e.g., the cue provided in a cued-recall experiment), or they may be supplied by the participant (e.g., a mental reconstruction of the learning context). Cues supplied by a recurrent environment (e.g., some events reoccur on most Sundays) are especially important because they appear to play a crucial role in learning to plan for, or think about, the future. The information that would be required to achieve a particular goal also needs to be specified. Note that the required information can include associations with memories, not just associations between stimuli/events and responses/outcomes. Once the required information is specified one can start thinking about the opportunities to learn that information and the presentation conditions that might promote as well as retard its learning. These aspects of a task analysis are relatively atheoretical in the sense that they do not depend on specific assumptions about representation (e.g., whether memories are stored in a network, separately, or in a composite memory) and process (e.g., whether the memory access process is similar to cued recall or whether it is an operation where the similarity of the incoming information is compared to the information in memory).

The final component of the task analysis is to consider the discrimination problem (the sources of noise) faced by the memory system. Ideas about the sources of noise are sometimes, but not always, dependent on specific ideas about representation and process. Our thinking in this regard is influenced by ideas about parallel processing and composite memories (Chappell & Humphreys, 1994; Dennis & Humphreys, 2001; Humphreys, Bain, & Pike, 1989; Humphreys, Pike, Bain, & Tehan, 1989), though in most situations the ideas that we invoke as part of our explanations can also be derived from other assumptions. By a composite memory we refer to the idea that memories are represented as distributed patterns (e.g., a matrix or a neural network) and memories are stored by superimposing a new pattern onto the existing pattern (e.g., by adding to an existing matrix). In particular we think that it is important to consider the possibility that there is a difference between the memory access operations that have been described as involving recall and matching. With recall, the memory system must converge to a particular output (e.g., a word) whereas a matching operation only involves a computation of similarity between an input pattern and memory.

A simple way to remember the components of the analysis is to use the acronym Gold CIONs, for goal, cue, information, opportunity, and noise.[1] Note that *coins* has been deliberately misspelled in a blatant attempt to attract attention and hopefully increase the memorability of the acronym. The misspelling is also a reminder of the order in which the different analyses would normally be employed.

Additional practical and theoretical considerations

In starting to think about memory we like to differentiate between knowledge, habit, and memory. Examples of knowledge would include knowing that Paris is the capital of France, or knowing how to ride a bike. These two examples might be further classified as involving semantic knowledge and procedural or

[1] We thank Karin Humphreys for suggesting this acronym.

perceptual-motor knowledge, respectively. We like to separate knowledge from memory because it allows us to ask somewhat different questions. For example, if you have a memory for a conversation with your father and have the subjective experience of hearing his voice, is your memory for his voice episodic or are you augmenting your episodic memory for the event with knowledge of your father's voice that you have acquired over many different episodes? Note that although the distinction between episodic and semantic memory (Tulving, 1972, 1983) has become very popular, it is easier to distinguish between knowledge and memory than between episodic and semantic memory. We postpone addressing the episodic/semantic distinction until the penultimate chapter of this book.

We use the term "habit" for behaviors that occur relatively automatically (habitually) under specified conditions. In the Skinnerian tradition these behaviors would be said to be under stimulus control (Skinner, 1938). For example, when seatbelt laws were first introduced in Queensland, Australia, many motorists complained that they could not reliably remember to fasten their seatbelts when they got into their car. Many of them had been buckling up for long drives but not for short ones and presumably it was this inconsistency that prevented them from consistently fastening their seatbelts once the law had changed and they were required to do so for all trips. With enough consistent practice, the behavior of buckling up could occur reliably even under a considerable amount of distraction (e.g., the kids are screaming in the back seat and the ice cream in the shopping bag is starting to melt). We are unsure how these habits are acquired. Two possibilities are the kind of stimulus-response (S-R) learning that was found in the work of Hull (1943) or the more recent and more sophisticated learning theories such as Rescorla and Wagner (1972). The problem is that with humans the memory system(s) contribute(s) to performance during the acquisition of these habits, so it is very difficult to determine just what processes are contributing to the acquisition of habits.

As we have indicated, we do not have a definition that distinguishes between episodic and semantic memory. However, we do have a process whereby we can explore issues related to this popular distinction. Our process starts with some ideas for identifying the presence of memories in nonverbal and low-verbal organisms. We then ask how well these simple memories will work in solving important problems if they are supported by a recurrent, or partially recurrent, environment. Next, we examine some of the characteristics that appear to be associated with memory performance in adult humans. These include forming an association between a stimulus and an existing memory, or between one memory and another (Yates, 1992); rapid binding (McClelland, McNaughton, & O'Reilly, 1995); controlling the memory access process (Humphreys, Murray, & Maguire, 2009); and a very large capacity (Brady, Konkle, Alvarez, & Oliva, 2008).

We talk about associations, bindings, and information almost interchangeably. Association is the traditional term and probably has too much surplus meaning associated with it to be generally useful. The term "binding" has been borrowed from computer science, where it basically refers to information that allows the program to determine that items or features are related (occurred together, is a part of, is an attribute of, etc.). In this sense, a binding is not an entity because the ability to extract information about A being related to B can involve both the contents of data structures and the processes that operate on those structures. This is illustrated in the Humphreys, Pike, et al. (1989) exposition of how the different global matching models can differentiate between intact and rearranged pairs (see Chapter 4). In the memory literature, binding is a useful term because the term association has such a long and checkered history, and because a discussion in terms of information can be awkward. In Chapter 4 we discuss the complexity of bindings in episodic memory. In Chapter 5 we raise the possibility that there are learnable connections (bindings) between modality-specific memory codes and more central memory codes (see Humphreys, Bain, et al., 1989). Finally, in Chapter 8 we discuss bindings between stimuli and memories. Henson, Eckstein, Wasszak, Frings, and Horner (2014) have also discussed multiple types of bindings that seem to be required to explain priming results. However, they use the traditional terminology of S-R, stimulus–stimulus, and feature-response bindings. A closer comparison of the distinctions that we make about different kinds of bindings and the distinctions that Henson et al. make will certainly reveal a considerable amount of overlap.

It is also important to question the idea that if you can access one component of an episode you will have relatively easy access to the other components. This assumption seems to lie behind some of the surprise at the Choice Blindness findings. That is, it seems surprising that when cued with the picture that was not chosen, the participant did not have access to information about the picture that was chosen. This idea about ready access to the other components of an episode seems to stem from two related theoretical assumptions. First, that it is possible to retrieve a memory trace and then inspect the contents of that trace. Such a possibility requires separate memory traces because this is simply not possible with a composite memory, where a new cue would be required in order to access another component of the event. Second, that it is possible to rapidly search memories, discarding those components that do not satisfy the memory query. A form of search is possible with a composite memory but it is a search through cues, not a search through memories. That is, one has to generate a cue and use that cue to access memory. The memory contents that are recovered can then be used as either an additional cue or by themselves as a cue. As a consequence, a search through cues is likely to be slower and more effortful than a search through memories.

Additional generic themes also emerge from our discussions of particular issues. In our treatment of several memory tasks we explore the possibility that the procedures used to enhance storage do not simply strengthen an association but become part of the memory for the list. In the memory literature there is also a bias to emphasize explanations based on storage processes over explanations based on retrieval processes. We tend toward the opposite bias because retrieval explanations are more testable than storage explanations, which, we feel, means that they should be considered first. This difficulty in studying storage is illustrated by a rhetorical question a leading animal-learning theorist asked the first author, Michael Humphreys, in the early 1990s. He was asked, "How can you study learning when you cannot control the learning history of your subjects?" The answer that came to him a few days later was that he primarily studied retrieval, and with humans the study of retrieval was easy because all you had to do was to hold learning or storage constant and change the instructions about what to retrieve. However, with nonhuman animals it was difficult to study retrieval because you had to train the animal to retrieve in a specific way, which means that the different retrieval groups have different learning histories (for an exception, see the discussion in Chapter 8 of cueing dolphins to flexibly retrieve).

Because adult humans would learn a single item or the association between a single pair of items in one trial, researchers have investigated human memory using lists of items or pairs. We are concerned with how the use of a list to study memory changes the memory task, and of the role of the other items in the list when it comes to the recognition or recall of one of the list items. We are also concerned with the differences between recognition and recall paradigms and with the possibility that enhanced performance with one type of item may come at the cost of reduced performance with another type of item. For example, in Chapter 4 we will be looking at many examples where two types of items (e.g., pronounced and unpronounced) are mixed together in a study list. When one of these item types is better recalled or recognized than the other, we turn to comparisons with between-subject designs to determine whether there has been facilitation or inhibition or both.

We have also chosen to use illustrative examples from the marketing/sponsorship literature and from the recent literature on applications of memory research to educational settings and issues. Sponsorship effectiveness is often evaluated by assessing memory for the brand-event link. This is particularly true when a brand is sponsoring a noncommercial entity such as a charity or an artistic company (Cornwell & Humphreys, 2013). In addition, while sponsorship and marketing materials are definitely more complex than the materials that we have used in our traditional laboratory paradigms, they are still relatively simple. Thus they provide a useful platform for seeing whether ideas derived from our laboratory paradigms can be applied. In using memory research in any applied setting, it is frequently necessary to consider a wide range of memory research. With respect to educational research, in some ways the paradigms that we have used over the last 40 years (a single study trial followed by a test) are less relevant than the multi-trial learning, transfer, and mediation paradigms that were used during the verbal learning era. However, as some of our examples show, there has been a tendency on the part of researchers to work in a silo. These silos are formed when researchers in an area almost exclusively reference the other researchers that are working in that area.

Overview of the book

The next five chapters of the book review the five components of our Gold CIONs analysis. In addition to reviewing some of both the historical and contemporary literature on each of these components, these chapters also contain some more or less speculative explanations for a variety of findings. For example, Chapter 2 examines the distinction between the goals of a memory task and the broader goals that the participant may have at either the study or test. Chapter 3 looks at the importance of cues and Chapter 4 looks at how a consideration of the information used can help in understanding a variety of memory tasks. Chapters 3 and 4 are reasonably straightforward reviews of existing results, though some of the relationships between contemporary concepts and earlier concepts, and even some of the earlier results, have been forgotten by contemporary researchers. The next two chapters complete our Gold CIONs analysis. Chapter 5 looks at those factors that appear to enhance memory storage. A major component of this chapter is an attempt to explain how the use of discriminative information might enhance memory. Chapter 6 looks at the role of noise in human memory with an emphasis on forgetting. It also looks at how an identification of the discrimination problem can help to understand when discriminative information helps.

Following on from our review of the Gold CIONs analysis are three chapters that address broader and more theoretical issues. Chapter 7 looks at how memory is controlled. This includes control by the environment, including the control permitted by a recurrent environment, control by others, and how we control our own memories. A major theme of this chapter is that a relatively simple memory system can solve complex problems if control is maintained by a recurrent environment and if stimuli can be associated with existing memories. Chapter 8 is concerned with how we understand episodic memory. We first review the historical development of the concept. Then we consider how three developments make it so difficult to determine how memory mechanisms have changed through evolution and development. These are: (a) a very large increase in our memory capacity, (b) an increased ability to control our memories, which is developed and nurtured by deliberate instruction during our long childhood, and (c) the use of words, which seems well suited to building up complex sequences in which memories are associated with other memories. The final chapter compares and contrasts three organizing frameworks that can be used to introduce students to the study of memory. First, there is the memory systems approach, which differentiates between episodic and semantic memory and/or between procedural and declarative memory. Second, there is the information processing approach, which differentiates between encoding/storage, retention, and retrieval. Third, there is our Gold CIONs approach. Although we favor the Gold CIONs approach, we also recognize the need to think in terms of systems, especially if we are to understand how memory is used to solve important problems. In discussing this issue we introduce the episodic problem-solving system.

2 Analyzing the goals of a task

In analyzing the Choice Blindness task in Chapter 1 we focused on the memory task, which involved the presentation of either the face that had been chosen as the more attractive face in the previous trial, or the face that had not been chosen, accompanied by the instructions to describe why the participant had chosen that face. This type of memory retrieval task has been formally analyzed by Humphreys, Wiles, and Dennis (1994). Their analysis was based on Marr's (1982) idea about specifying the goal of a task. Marr started by identifying an abstracted version of a task that was solved by the human cognitive system. For example, deriving the shape of an object from the pattern of shading on the retina. The pattern on the retina was derivable from the physics of light and the optics of the human eye. Marr was then able to show how a mathematical analysis (he referred to this as a computational-level theory) of how this input could be used to derive the outcome (a representation of the 3-dimensional shape of the object) could constrain algorithmic level theories of the process. Note that Marr's computational-level theory did not predict such details as how long the process would take or the kind of errors that might be made. These details were left to the algorithmic level theories. In their analysis, Humphreys et al. (1994) started with well-known laboratory paradigms where they could identify the inputs to a memory retrieval task. These included the instructions, cues, and memory structures. The memory structures are similar to our concept of information except that they are formal structures. In the Humphreys et al. article the memory structures were sets of bindings between words, relations, contexts, etc. The instructions were used to specify what would constitute a correct response which was also referred to as the goal of the task. For example, in cued recall with an extralist cue, a non-studied cue that elicits one of the study words in a free association task is provided as a cue. The instructions tell the participant to use the cue to recall a meaningfully related word from the study list. Thus a correct response is a word that was in the list and is meaningfully related to the cue.

Like Marr (1982), Humphreys et al. (1994) hoped that identifying the inputs to a task and the task's goal would help to constrain algorithmic-level theories. However, with memory, unlike vision, we do not know how the inputs to a task