

Chapter 1: Introduction

The Problem of Serial Order

How do we store and retrieve a sequence of items in the correct order? For example, how do we recall a six-digit telephone number with apparent indifference to the other 719 permutations of the digits? Moreover, how do we recall that number with apparent indifference to the numerous permutations that correspond to other telephone numbers, bank codes, etc? This is the *Problem of Serial Order* (Lashley, 1951). The present thesis attempts an answer to this problem, at least with respect to short-term memory for temporal order.

More generally, all our actions are sequenced in time. Consider the articulation of a single word: This involves the ordered production of a number of phonemes. Given a typical production vocabulary over 50,000 words in English, this is a huge number of sequences drawn from a set of only 45 or so phonemes. Yet we have little problem articulating words rapidly and correctly. For example, it is extremely rare for the four phonemes in “style” to be misordered to produce articulation of “slight” instead. How is this huge capacity stored in memory, virtually free from interference?

We can normally remember the order of major events in our lives without recourse to external aids such as diaries. Here we are more prone to error, forgetting whether something happened yesterday or the day before, but there is certainly the impression of an autobiographical continuum, along which events recede hazily into the past. This continuum may be punctuated by factual knowledge, such as dates of important occurrences, but these aside, how are events ordered on such a continuum?

These questions illustrate the fundamental nature of the problem of serial order; an unresolved problem that cuts across many traditional psychological distinctions of memory. Clearly, phone numbers heard once may be forgotten completely an hour later, whereas others, such as those of our first home, may stay with us forever. This difference is often thought to reflect temporary storage in short-term memory versus more permanent storage in long-term memory (e.g., Baddeley, 1986). Spelling a word aloud may involve declarative memory,

whereas the sequence of motor commands used to write the word involves procedural memory (e.g., Squire, 1994). Finally, episodic memory for the order of past events can be distinguished from semantic memory for facts (e.g., Tulving, 1983), such as the order of British Monarchs for example. There is no a priori reason why such different memories should use the same means of representing serial order; there is similarly no reason why they should not.

This thesis is concerned with the problem of serial order in short-term memory. More specifically, the domain is short-term, episodic memory for the temporal order of verbal material, a memory tapped by the task of serial recall. With respect to the first example, this is the memory that allows one to retain an unfamiliar telephone number long enough to dial it shortly afterwards, even though the same number might be forgotten minutes later. The advantage of this domain is that many variables that may affect memory for serial order can be readily manipulated and controlled in the laboratory. Though the answers that emerge from this restricted domain may not generalise to all of the examples above, they represent a good starting point. The extent to which the answers do generalise is discussed in Chapter 8.

Theories of Serial Order

There are three basic theories of serial order: chaining theory, positional theory and ordinal theory. Each theory is introduced below, in the general terms of ordering a sequence of elements, where those elements might be digits in a telephone number, movements in an complex action, or events in autobiographical memory.

Chaining Theory

This theory assumes order is stored by the formation or strengthening of associations between successive elements. The order is retrieved by stepping along these associations in a process called *chaining*, where each element cues the recall of its successor.

Chaining is probably the oldest approach to serial order (Ebbinghaus, 1964) and certainly the most intuitive. It is a simple extension of stimulus-response theory, where each response can become the stimulus for the next (Lashley, 1951). In its various guises, it has remained popular in several different models (e.g., Elman, 1990; Jones, Beaman & Macken, 1996; Jordan, 1986; Lewandowsky & Murdock, 1989; Richman & Simon, 1994; Wickelgren,

1965b). However, chaining theory faces several problems, as discussed below.

The simplest chaining models assume only pairwise associations between adjacent elements of a sequence (e.g., Wickelgren, 1965b), and a cue which consists entirely of the immediately preceding response (upper illustration in Figure 1-1). There are several immediate objections to such *simple chaining models*. For example, how do they handle sequences with repeated elements, in which two or more different elements will share the same cue? Or how do they allow recovery from error, because once an error has been made, the cue for all subsequent responses will be incorrect? This should lead to a cascade of further errors (“a chain is only as strong as its weakest link”).




Figure 1-1: Simple and compound chaining models
(upper and lower illustrations respectively).

These questions are less problematic for models that assume remote associations as well as adjacent ones (e.g., Ebbinghaus, 1964; Slamecka, 1985). In these *compound chaining models* (lower illustration in Figure 1-1), the cue consists of a number of preceding elements, an approach which is popular in recurrent neural networks (e.g., Elman, 1990; Jordan, 1986). These compound cues allow disambiguation of elements repeated in a sequence, by virtue of the additional context of elements preceding the repeated elements. They also mean that a single error is less devastating, the additional context allowing recovery from that error.

There are other possible solutions. With respect to the problem of repeated elements, one can appeal to the type/token distinction, so that two occurrences of the same type have nonidentical token representations. For example, representations of each element may be embedded in different temporal or spatial contexts, allowing the same element to function as a different cue at different positions in a sequence, as in the case of Wickelgren's "allophones" (Wickelgren, 1969). This distinction is discussed in Chapter 7.

With respect to the problem of errors in recall, the TODAM model (Murdock, 1983) assumes only pairwise associations¹, but allows recovery from errors by only cuing with the previous response if it is correct. Otherwise, a cue approximating the correct one is used (Lewandowsky & Murdock, 1989). Though this may be appropriate when feedback of the correct response is provided, it is inappropriate for most situations (such as serial recall), where one does not always know whether each response is correct. Nevertheless, this approach illustrates an important distinction in chaining theory: whether the cue consists of the preceding elements recalled, which may be erroneous, or whether the cue consists of the preceding elements stored, irrespective of whether or not they are recalled correctly. This distinguishes *closed-loop* chaining models, where responses are fed back as cues, from *open-loop* models, where there is no feedback and therefore no necessary detrimental effect of errors. This distinction is particularly important in Chapter 2.

Nevertheless, there are many arguments against the sufficiency of chaining theory as a general account of sequential behaviour (e.g., Johnson, 1972; Lashley, 1951). One argument concerns the interference predicted by most chaining models as soon as several sequences of the same elements are stored (at least for those that assume type representations). For example, how could the order of letters in the words *pat*, *apt* and *tap* be retrieved if each letter is associated with almost every other (Houghton & Hartley, 1996)? The scale of this *interference problem* is apparent in the above example of 50,000 words drawn from a set of 45 phonemes: The degeneracy of associations assumed by chaining theory would surely predict much greater interference in speech production.

Another argument concerns the feedback of responses in closed-loop models. Though

1. (though see Murdock, 1993, 1995 for extensions of TODAM that use more than pairwise associations)

“...the only strictly physiological theory that has been explicitly formulated... postulates chains of reflexes, in which performance of each element of the series provides excitation of the next” (Lashley, 1951, p. 114), Lashley argued that many actions are executed so fast that there is simply not enough time for feedback to cue the next response. The finger presses of a skilled typist, for example, are too fast for proprioceptive feedback (Rumelhart & Norman, 1982). This argument has had a profound effect on motor control theory (Bruce, 1994), leading to the idea of motor programs that are independent of feedback, and rejecting the idea of reflex chains (though clearly there is some role for feedback; MacKay, 1982).

Lashley’s most persuasive arguments concern everyday observations that order involves more than the linear structure of chaining theory. Many sequences are ordered hierarchically, such as the order of words in a sentence, the order of syllables in each word, the order of phonemes in each syllable, etc. Moreover, the order of elements typically respect syntactic or schematic structure beyond interitem associations. For example, speech errors show that a noun is far more likely to swap with another noun than a verb, even if the noun is further away in the sentence (Levelt, 1989). Such structure is reflected at all levels of speech production. For example, the structure of a syllable imposes constraints on the order of its constituent phonemes that is clearly beyond chaining theory (Hartley & Houghton, 1996).

Finally, the question remains as to what cues the first element in a sequence, in order to “kick-start” the chaining process. Many chaining models appeal to a separate contextual cue (Murdock, 1995) or plan unit (Jordan, 1986). Thus, chaining theory is clearly not a sufficient answer to the problem of serial order; later chapters will question whether it is even necessary.

Positional Theory

This theory assumes order is stored by associating each element with its position in the sequence. The order is retrieved by using each position to cue its associated element. In other words, rather than using the item-item associations of chaining theory, positional theory uses position-item associations.

The simplest example of a positional theory is Conrad’s “box” model of short-term memory (Conrad, 1965). Conrad assumed that people possess a number of boxes in memory, into which elements of a sequence can be placed (upper illustration in Figure 1-2). The order

of elements can be retrieved by stepping through the boxes according to a predetermined routine. This model does not have a problem with repeated elements, because they are stored in separate boxes, nor with recovery from errors, because the retrieval mechanism can continue to the next box irrespective of whether the contents of the previous box were correct (i.e., no feedback is required). This is the method by which conventional (Von Neumann) computers store and retrieve order, through routines accessing separate addresses in memory.




Figure 1-2: Positional models of Conrad (1965) and Burgess and Hitch (1992) (upper and lower illustrations respectively).

As a psychological model however, such a simple model is inadequate. Firstly, how many boxes do people have in short-term memory: five, six, seven, or more? If a new box is created for each element in a sequence, there would be no limit to the length of sequences people could hold in short-term memory, which is clearly not the case. Secondly, the model provides no immediate rationale for the errors people make when they misremember sequences: People are more likely to confuse elements close together in a sequence than far apart (Chapter 2). There is no reason for this with the perfect coding and retrieval of positions assumed by Conrad (1965).

One way to explain the above errors is that positional codes become confused over time. For example, the Perturbation Model (Lee & Estes, 1977; 1981) assumes the positions of elements are initially coded perfectly, but can get perturbed during storage such that nearby elements swap. Another way to explain such errors is that positional codes are not perfect, but overlap, in the sense that the code for one position is similar to the codes for nearby positions (lower illustration in Figure 1-2). This is the approach taken by the Articulatory Loop Model of Burgess and Hitch (1992). In this model, the circles in Figure 1-2 represent nodes in a connectionist network. The filled nodes are active nodes; the unfilled nodes are inactive nodes. The “window” of active nodes moves from left to right for each position in a sequence, and is associated with other nodes (not shown) representing each element. However, because there is some overlap in the set of active nodes for nearby positions, elements at these positions can sometimes be confused during retrieval.

Positional theory can be extended to a hierarchy of positional codes (e.g., Lee & Estes, 1981). For example, a phoneme can be coded for its position in a syllable, a syllable can be coded for its position in a word, a word can be coded for its position in a sentence, etc. As a general solution to the problem of serial order however, the status of positional theory remains unclear. There is a sense in which the problem is not solved, but circumvented. This sense reflects the question of how the order of the positional codes, rather than the order of the elements, is stored and retrieved from memory. This question cannot be answered without specifying the nature of the positional codes. One suggestion is that the codes are successive states of internal oscillators in the brain (Brown, Preece & Hulme, 1996; Burgess & Hitch, 1996a, 1996b). In these models, elements are associated with different states of the oscillators, and these states can be reconstructed simply by resetting the oscillators and letting them change under their own dynamics. In other words, the oscillators represent a biological clock, which can be rewound in order to retrieve a sequence from memory.

Positional theory does not solve the interference problem however. If a positional model is to store and retrieve the order of letters in *pat*, *apt* and *tap*, then it must employ different positional codes for each word (otherwise each letter will be associated with several positions). If these positional codes were states of internal oscillators for example, a different

set of oscillators might be required for every word we know. Furthermore, the syntactic constraints on serial order, such as those in speech production, involve more than simply positional information (Chapter 8).

Nonetheless, this thesis provides good evidence that people do use positional codes in short-term memory for serial order, and describes a new model that attempts better specification of the codes. The question of whether such a positional model is sufficient for long-term memory for serial order is resumed in Chapter 8.

Ordinal Theory

This theory assumes order is stored along a single dimension, where that order is defined by relative rather than absolute values on that dimension. Order can be retrieved by moving along the dimension in one or other direction. This theory need not assume either the item-item nor position-item associations of the previous theories.

For example, Grossberg (1978) assumed that order is stored in a primacy gradient of strengths in memory, such that each element is stronger than its successor. The order of elements is retrieved by selecting the strongest element, suppressing it, selecting the next strongest, suppressing it, etc. (Figure 1-3; suppression indicated by the broken lines). This idea has been incorporated into the Primacy Model of short-term memory (Page & Norris, 1996b), where the strengths might represent the degree of association to the start-of-sequence context, or even simply activations of item representations in memory.

The original Perturbation Model (Estes, 1972) is also an ordinal theory, where order is inherent in the cyclic reactivation of elements. Perturbations in the timings of reactivations lead to erroneous reorderings of the elements, like shifts in the relative phases of a series of oscillations.² Yet another ordinal model is that of Shiffrin and Cook (1978). This model assumes associations between each element and a “node”, but only the nodes are associated with one another (i.e., unlike chaining models, where it is the elements themselves that are associated with each other). By moving inwards from nodes representing the start and end of the sequence, the associations between nodes allow the order of items to be reconstructed.

2. Later developments of this model (Lee & Estes, 1977, 1981) assumed the perturbations were of positional attributes of items, rather than the timing of the items themselves, making it a positional model.




Figure 1-3: An ordinal model (e.g., Page and Norris, 1996b).

Ordinal models like the Primacy Model require token representations in order to handle repeated elements: The order of repeated elements could not be represented over type representations with a single strength. As regards errors in recall, ordinal models imply that errors will cooccur, in the sense that one error will cause another (because order is defined relationally). For example, if an element becomes stronger than its predecessor in the Primacy Model (owing to random noise), then the elements will transpose, causing two errors. This is an attractive property, because such transpositions are common in people too (Chapter 2).

Ordinal models do not require feedback of responses, and a process like suppression can operate independently of errors occurring at later stages of output. Some ordinal models have also been extended to hierarchical structuring in a manner that allows them to overcome the interference problem (Niggin, 1993; Page, 1994). This is achieved by assigning a new node to each sequence, and associating elements to that node with a strength proportional to their relative order (via a primacy gradient). Because each sequence is associated with a unique node, the order of different sequences can be stored without interference (Chapter 8). This is a simpler solution than assuming separate positional codes for each sequence.

Thus ordinal theory escapes some of the criticisms of chaining and positional theories. Nonetheless, the present thesis will argue that ordinal theory is insufficient as an account of people's short-term memory for serial order.

Theoretical Differences

In spite of the various strengths and weaknesses of specific models discussed above, important differences remain between the three theories of serial order. The difference between chaining and positional theories is obvious: the retrieval cue in the former is the previous element; the retrieval cue in the latter is some (abstract) positional code. The difference between positional and ordinal theories is less obvious, but relates to whether the position of an element in a sequence can be defined independently of its surrounding elements. In positional theories, it can; in ordinal (and chaining) theories, it can not. The consequence is that, in ordinal models, the middle element in a sequence can only be retrieved after retrieval of its predecessors (or successors, in the model of Shiffrin & Cook, 1978). In positional models however, it is possible to retrieve the middle element directly, by reinstating the appropriate positional cue. This is crucial in explaining a class of positional errors found in short-term, serial recall (Chapter 3). More specific differences between these theories are detailed in Chapters 2, 3 and 4.

There are, of course, means of storing serial order that do not fit easily within these theories. For example, order may be stored propositionally, in the form of statements such as “R was before M” and “Q was after M” (though these might be regarded as a form of interitem association). While appropriate perhaps for semantic memory (e.g., the order of British Monarchs), this approach is less appropriate for short-term memory for serial order. Remembering that “R was before M” may supplement short-term memory, but such memories do not appear necessary, and no explicit models have been developed along these lines. Subsequent chapters are therefore restricted to the three theories outlined above.

Empirical Differences

The three theories of serial order can be distinguished empirically, though not always easily. Most of the relevant experiments have been performed within the domain of short-term memory (see Chapter 8 for evidence from other domains). In these experiments, the sequences are typically novel lists of verbal items to be recalled shortly after their presentation. These experiments fall into three main types: serial learning, probed recall and serial recall.

Serial Learning

Chaining and positional theories were first pitted against one another within the serial learning paradigm, under the guises of item-item versus position-item associative theories (e.g., Young, 1968). The serial learning paradigm uses lists of 10 or more items, beyond normal memory span (i.e., the lists are rarely recalled correctly on the first attempt). Subjects are given repeated presentation and recall trials until they learn a list to criterion. The manner in which they have learned the order of items is normally investigated by the amount of transfer to a second task. For example, list learning might be followed by paired associate learning of items that were either adjacent or nonadjacent in the list. If a list *RMQJHV* were learned by forming associations between successive items, as predicted by chaining theory, positive transfer might be expected in paired associate learning of *RM* and *QJ*, but not *HR* or *VM* (e.g., Young, 1962). Alternatively, if serial learning of the list *RMQJHV* were followed by serial learning of a second list *RJQVHM*, or *JHVVRMQ* (e.g., Ebenholtz, 1963; Keppel & Saufleay, 1964; Slamecka, 1964), then positional theory, but not chaining theory, might predict positive transfer in the former case (because three of the items retain their positions from the original list), while chaining theory, but not positional theory, might predict positive transfer in the latter case (because most of the items retain the same predecessor from the original list).

Unfortunately, conclusions from the serial learning paradigm were mixed (Ebenholtz, 1972; Jensen & Rohwer, 1965; Young, 1968). Indeed, they led some researchers to propose that both item-item and position-item associations are utilised in serial learning (Battig, Brown & Schild, 1967; Houston, 1976). Alternatively, the mixed results may reflect a problem with the paradigm, in that the two theories are hard to dissociate fully. For example, in the paired associate task, any transfer for the associates *RM* and *QJ* could equally well reflect mediation of similar positional codes (given that they come from nearby positions), rather than interitem associations. Positive transfer in such cases does not rule out the positional theory therefore. In fact, given the proximity required for interitem associations, the chaining and positional accounts will always be confounded in these situations. Another problem is that transfer between the two tasks presumes that the associations underlying the tasks are identical, and yet there may be quite separate mechanisms underlying serial and paired associate learning.

The second example of transfer between two serial learning tasks is not without problems either. One problem is that, without a more detailed model of how serial learning is achieved (via item-item or position-item associations), the exact predictions of the two theories remains unclear. For example, if new item positions impaired learning at the same time as old positions improved learning, then a lack of transfer from *RMQJHV* to *RJQVHM* could reflect a balance between positive and negative transfer for individual items. Similarly, a lack of transfer from *RMQJHV* to *JHVRMQ* could reflect a large negative effect of unlearning old interitem associations before the learning of new ones is possible, particularly if those associations are remote as well as adjacent. Another problem is that the results from this transfer task are sensitive to the exact method of serial learning, such as whether subjects are aware of the relationship between the original and derived list (Maisto & Ward, 1973).

In summary, the use of transfer tasks has several methodological problems, and the serial learning paradigm has not proved a fruitful means of testing theories of serial order.

Probed Recall

Another means of testing theories of serial order is with probed recall (e.g., Murdock, 1968). For example, given a list followed by a probe item from that list, chaining theory predicts it to be a simple matter to recall the item that followed the probe in the list (via the association between them). If, on the other hand, the probe were a number corresponding to a position in the list, positional theory might predict it to be a simple matter to retrieve the item at that position. Ordinal theory predicts both tasks to be somewhat harder.

The main problem with the probed recall paradigm however is that there are many ways to perform the task. For example, in the first task above of *item-probed successor recall* (e.g., Waugh & Norman, 1965), performance might not be based on a direct item-item association, but rather an indirect series of item-position-item associations, by first retrieving the position of the probe and then cuing with the next position. The fact that people can perform the task of *item-probed position recall* (e.g., Jahnke, Davis & Bower, 1989; McNicol, 1975) is consistent with this hypothesis. The basis of performance on the second task above of *position-probed recall* (e.g., Nairne, Whiteman & Woessner, 1995) is uncertain because there may not be any simple transformation of the probe into a positional code (i.e., no direct

mapping between numbers and internal positional codes).

The uncertainty in the processes underlying probed recall is reinforced by examining response latencies. Sternberg (1967) showed that the latency in item-probed successor recall increases linearly with position of the probe in the list. This suggests serial search from the start of the list, until the probe is encountered. In other words, performance on this probe task is probably based on covert serial recall of the whole list. This is exactly what the ordinal theory would have to predict. Nonetheless, performance is improved if the probe item is accompanied by a spatial, positional probe (Hitch, 1972), which attenuates the effect of probe position on latency (Monsell, 1973). This suggests some role for positional information, consistent with data from experiments using a spatial probe alone. Sanders and Willemssen (1978a), for example, showed that response latency with a spatial probe is a nonmonotonic function of position, with a recency advantage in addition to the primacy advantage of Sternberg's task. A spatial probe may therefore ameliorate the need for serial search, at least for terminal items, supporting positional theory. The problem with this approach however is that it remains unclear how a spatial probe interacts with memory for temporal order, and whether the results would hold in the absence of spatial information (Chapter 5).

In summary, the probed recall task has also faced problems distinguishing theories of serial order, particularly for ordinal and positional theories, mainly owing to uncertainty in the processes underlying performance of the task.

Serial Recall

A better way to test theories of serial order is with serial recall. This task is simply to recall the whole list in a forward order, from the first to the last item. It proves particularly fruitful to study the pattern of errors when recall fails.

One example of errors in serial recall are *associative intrusions* (Wickelgren, 1966). These are transpositions between the items immediately following repeated items, and are more common than corresponding transpositions following nonrepeated items. Wickelgren used such errors to support chaining theory: Repeated items are ambiguous cues for the items that follow them because they are associated with more than one such item. These errors are reexamined in Chapter 7, though it is argued that they do not, in fact, constitute evidence for

chaining theory. Another example are *serial order intrusions* (Conrad, 1960). These are errors of items that occurred at the same position in the previous trial, and are significantly more common than predicted by chance. Such errors clearly support positional theory, and are examined in more detail in Chapter 3.

Thus, as Estes (1972) observed: “When retention is imperfect, the confusion errors that occur are highly systematic” (p. 161). Though a single error may reflect a temporary failure to realise an accurate representation in memory, large numbers of errors show striking patterns in their distribution. These patterns shed light on the mechanisms subserving serial recall and hence the underlying representations of serial order over which the mechanisms operate (Conrad, 1959). Errors in serial recall are seldom random guesses.

There are several close cousins to serial recall, such as backwards recall (where items are recalled in reverse order), or positional recall (where items must be placed in the correct positions, but the order of recall is unconstrained), which may also shed light on theories of serial order. However, serial recall remains the most important task, given that it underlies the basic index of short-term memory, *memory span*, and many of the relevant empirical dissociations in short-term memory research (Baddeley, 1986).

Serial recall is also an everyday cognitive activity (e.g., recalling a telephone number). Indeed, a forward order is the default (and optimal) recall order from short-term memory, at least for up to six or seven items. Serial recall is therefore less likely to be contaminated by specialised strategies than are other unusual, and perhaps artificial, laboratory tasks, like probed recall. Moreover, because people are trying their best not to make errors (as opposed to trying their best to make use of a probe), error analysis is more likely to reveal the underlying representation of serial order, particularly given that people are often unaware of their errors (Chapter 6). Not surprisingly therefore, experiments in this thesis are confined to serial recall.

Measuring Serial Recall

The earliest measurement of serial recall was the proportion of lists recalled correctly (e.g., Crannell & Parrish, 1957). This measure underlies the span index of short-term memory: One’s memory span is usually defined as the length of list that one can recall correctly 50% of

the time. However, this measure ignores differences in the recall of each item in a list. Murdock (1968) pioneered the use of *serial position curves*, which plot the proportion of items recalled correctly at each position of a list. These curves are bowed, with an advantage in recall for the first and last few items. The advantage for early items is termed *primacy*, and the advantage for later items is termed *recency*. In immediate serial recall, primacy is normally more pronounced than recency (Chapter 4).

Later work distinguished two main types of error: *order errors* and *item errors* (e.g., Estes, 1972). Order errors are list items recalled in the wrong position; item errors are list items not recalled anywhere in the report. Estes analysed order errors further by comparing an item's position in the list with its position of recall. This analysis showed that erroneous items are clustered around their correct position, rather than being randomly distributed.

However, few studies actually go beyond measuring the proportion of lists correct or plotting serial position curves, let alone analysing the distribution of order errors. This thesis attempts a more comprehensive classification of errors. Indeed, a major theme behind the thesis is that a great deal of information is available through analysing error patterns in more detail. This may be why Conrad is reputed to have said that error analysis is “the royal road to memory” (A. D. Baddeley, personal communication, 1995).

Classification of Errors

The classification of errors used in the present thesis is described below, with examples given in Table 1-1. This classification distinguishes an item's position in a list, its *input position*, from its position in a subject's report of that list, its *output position*.

When scoring by output position (i.e., taking each response in a subject's report), errors can be broadly categorised into *omissions* and *substitutions*. Omissions arise when no item is given for a position; substitutions arise when an incorrect item is given. Substitutions may be either *transpositions* or *intrusions*. Transpositions are list items in the wrong position; intrusions are items that were not present in the list. Intrusions may be items outside the experimental vocabulary (the set of items from which all lists are constructed), but most often they are items appearing on previous trials. Those intrusions that come from the immediately preceding trial are called *immediate intrusions*.

Error Type	List (Input Positions)	Report (Output Positions)
Omissions	R M Q J H V	R M K J = =
Transpositions	R M Q J H V	R M Q H J V
Intrusions	R M Q J H V	R M E J Y V
Confusions	R M Q J H V	R M Q K H P
Repetitions	R M Q J H V	R M Q R H M
Associates	R M Q J H V	R M J H V Q
Interpositions	R M Q J H V	R M V J H Q
Protrusions	F P Y K Z W R M Q J H V	F P Y K Z W R M Y P Z V

Table 1-1: Example errors in serial recall.
(Errors are in bold; items corresponding to a particular error type are underlined.)

A special class of substitutions are phonological *confusions*. These are incorrect items that are phonologically similar to the correct item, and are common in tests of immediate memory. A special class of transpositions are *repetitions*. Repetitions are items that occur more than once in a report, even though they only occurred once in the list. The distribution of repetitions sheds light on the retrieval processes underlying serial recall (Chapter 4). Another special class of transpositions are *associates*. These are items recalled in the correct order

relative to the previous item recalled, albeit in the wrong position (e.g., Wickelgren's associative intrusions). These errors are predicted by chaining models (Chapter 2).

Two further types of *positional errors* can be identified. *Interpositions* arise when lists are split into groups (e.g., by a pause between presentation of every third item). They are transpositions between groups that maintain their position within a group. *Protrusions* are similar errors, but maintain position between trials rather than between groups (i.e., Conrad's serial order intrusions). Note that the definition of protrusions is orthogonal to that of transpositions and intrusions, in that a protrusion may be either an intrusion or a transposition with respect to the current trial (which is why the term is preferable to Conrad's). Positional errors are predicted by positional models (Chapter 3).

Unless stated otherwise, errors are classified by output position. However, additional information is provided by scoring against input position (under which categorisation of errors is similar, and fairly self-evident; Table 1-1). For example, when omissions are scored against input position, they represent items that are not recalled anywhere in a report. This distribution of omissions can differ from that plotted against output position (Chapter 4). The distinction between item and order errors also generally refers to input position (e.g., Healy, 1974), though this distinction is not used often in the present thesis.

This concludes the majority of errors distinguished in this thesis. Though the complete classification might appear somewhat complex (certainly more comprehensive than conventionally attempted), each type of error plays an important role in constraining models of serial recall, as subsequent chapters will demonstrate.

Additional Terminology

In the present thesis, *error position curves* are plotted instead of conventional serial position curves. These show the percentage of responses that are errors at each position. When plotting any error, these curves correspond to an inversion of serial position curves about the 50% performance line (and the distinction between input and output position is irrelevant). However, errors of different types can also be plotted separately, against either input or output position, so providing more information than conventional serial position curves.

Transpositions can be categorised with respect to both input and output position. This

information can be represented in matrix form, where the entry in the i th row and j th column represents the number of items in Position i of a report that came from Position j in the list. The serial position curve corresponds to the main diagonal of this matrix, while the rows can be plotted to give *transposition gradients*. These gradients are generally peaked, with the peaks representing correct items ($i=j$) and the number of transpositions decreasing with increasing transposition distance $|i-j|$ (e.g., Figure 2-2 in Chapter 2). Immediate intrusions can also be classified in this way, where the j th column represents the position in the previous trial (either the previous list or the previous report). Such *intrusion gradients* show a similar, though flatter, pattern to transposition gradients (e.g., Figure 3-6 in Chapter 3).

Another error of interest is the *first error* to occur in a report. This allows calculation of the conditional probability of an error, given that previous responses are correct (Henson, Norris, Page & Baddeley, 1996). Using survival analysis (Appendix 1), this measure reveals more subtle changes in error probabilities across positions, such as the effects of group boundaries (Chapter 3).

Finally, the terms *Short-term Memory* (STM) and *Long-term Memory* (LTM) used in the present thesis may have different meanings in other contexts. STM and LTM are not meant as theoretical constructs, such as the *Primary* and *Secondary* memories of Waugh and Norman (1965). Though STM (or *Working Memory*, Baddeley, 1986) may correspond to a distinct memory system (Schacter & Tulving, 1994), the only distinction used here is that information in STM is temporary, being forgotten after a matter of minutes, while information in LTM is permanent. The relationship of STM to other aspects of memory is discussed in Chapter 8.

Experimental Design

The eight experiments described in this thesis employed serial recall of between five and nine items. The items were familiar, so that only memory for their order was required, which changed every trial. The items were presented sequentially, in the middle of a computer screen, each one replacing its predecessor. Thus serial order was only defined temporally; there was no spatial information. Presentation rates were between one and two items per second. Subjects were instructed to guess if unsure, but could omit if no item came to mind.

Apart from the main manipulations of interest, the experiments differed in whether the items were digits (Experiment 2), letters (Experiments 1, 6, 7, 8) or words (Experiments 3, 4, 5). Some experiments required vocalisation of the items as they appeared (Experiments 3, 6, 7, 8); the others required items to be read in silence (Experiments 1, 2, 4, 5). Some required recall immediately after the last item had disappeared (Experiments 1, 2, 4, 5); the others delayed recall by three, vocalised distractor digits (Experiments 3, 6, 7, 8). Some required written recall (Experiments 1, 2, 4, 5, 6, 8); the others required spoken recall (Experiments 3, 7). These procedural differences were not of primary importance and their implications are only mentioned in passing.

The main manipulations were the phonological similarity of items (Experiment 1), the length and temporal grouping of lists (Experiment 2), the length of the intertrial interval (Experiment 3), the relative size of different groups (Experiment 4), the relative length of different lists (Experiment 5), the presence of repeated items (Experiments 6, 7) and the effect of guessing instructions (Experiments 4, 5, 8). These manipulations arise from predictions of different theories of serial order and different models of serial recall.

Statistical Tests

Many analyses in the present thesis concern proportions or probabilities (e.g., the proportion of responses that are errors). Such proportional scores are suspect to floor and ceiling effects when they are close to 0 or 1, sometimes producing skewed distributions. To make some allowance for this, an empirical, log-odds transform is used (Appendix 1). This transform is particularly useful for proportions of errors of a certain type (e.g., the proportion of errors that are transpositions), given that some subjects make more errors than others: The log-odds can be weighted, giving more weight to scores from subjects who make more errors. In other words, the mean proportion can be weighted by sample size.³

The means and standard deviations of proportions are given in tables (the latter in brackets). When the sample size is fixed across subjects, these statistics normally represent

3. This weighting is restricted to pairwise comparisons of log-odds in the present thesis. Weighted ANOVAs can be performed on log-odds, but they lack orthogonality, making interpretation of interactions ambiguous. ANOVAs in the present thesis were therefore restricted to unweighted log-odds.

untransformed data. When sample sizes vary across subjects however, the statistics are calculated by retransforming weighted log-odds, in order to give a truer indication of population statistics. This retransformation is always indicated in the table captions. (One caveat with this approach is that retransformed proportions will not always sum to exactly 1.) When summary statistics refer to approximate proportions, they are expressed as percentages.

Finally, a comparison between two mean proportions is deemed statistically significant for alpha levels below .05. When making multiple pairwise comparisons, significance levels are based on Holm's method for adjusting alpha, which is essentially an iterative application of the Bonferroni correction (Appendix 1).

Computational Modelling

This thesis combines both experimental and modelling approaches. In particular, many of the analyses performed in Chapters 2, 3 and 4 were used to develop the computational model of short-term memory described in Chapter 5. This model makes predictions which were used to guide further experiments in Chapters 6 and 7. This interaction between empirical and computational approaches illustrates the important role of modelling in psychological research.

There are many advantages of computational modelling over more traditional verbal theorising (e.g., Hintzman, 1991). The main advantage is that computational models can be specified unambiguously. This reduces misinterpretation and makes their predictions clearer, improving their testability (indeed, provision of the relevant computer program allows anyone to validate a model's predictions). Computational models also allow direct quantitative fits to data, rather than the qualitative hand-waving made by many verbal theories (e.g., they can not only predict an interaction, but also the size of the interaction). In particular, the ability of computational models to simulate nonlinear, probabilistic processes means that complex interactions can arise from relatively simple mechanisms; interactions that are hard to predict a priori (or analytically, in the case of mathematical models). Given the sequential dependency between responses in serial recall (Henson et al., 1996), this ability proves particularly important in explaining the complex interactions between different error types (Chapter 5).

Overview of Thesis

In Chapter 2, the predictions of chaining theory are tested in immediate serial recall of lists of phonologically confusable and nonconfusable items. The data from Experiment 1, together with those in Henson et al. (1996), provide no support for chaining theory.

In Chapter 3, the effects of list length, grouping (Experiment 2) and proactive interference (Experiment 3) are examined. The pattern of errors found between groups and between trials is explicable by a positional theory of serial order, but not an ordinal theory.

In Chapter 4, meta-analyses of error data are performed on a number of experiments performed recently at the Applied Psychology Unit. These analyses produce a set of empirical constraints that any model of short-term, serial recall must meet. No previous model can.

In Chapter 5, a new, computational model of serial recall is developed, the Start-End Model (SEM), which meets the empirical constraints of Chapter 4, and provides quantitative fits to the data from Experiments 1, 2 and 3. This model is an example of a positional theory. Extension of SEM to other phenomena in STM is discussed, as is its relationship to previous models. Most importantly, SEM predicts an new property of positional errors in serial recall.

In Chapter 6, Experiments 4 and 5 confirm the predictions of SEM, and pose a serious challenge to other positional models, particularly those that assume positional codes are generated by internal oscillators in memory.

In Chapter 7, Experiments 6, 7 and 8 examine the effects of repeated items in serial recall. The results, together with those in Henson (1996b), are consistent with the basic assumptions of the SEM, but suggest that several additional processes are involved in memory for repeated items.

In Chapter 8, the more general assumptions of a positional solution to the problem of serial order are discussed. It is concluded that, while not denying other representations of serial order, particularly in procedural memory, positional theory appears a promising approach to the problem of serial order, particularly in episodic memory.