7 Is there a common mechanism underlying word-form learning and the Hebb repetition effect? Experimental data and a modelling framework

Mike P. A. Page and Dennis Norris

Overview

The Hebb repetition effect (Hebb, 1961) is a phenomenon whereby performance on the immediate serial recall of a list of familiar items is seen to improve over unannounced repetitions of a given list. One possible real-world counterpart of this effect is the learning of phonological word-forms that are themselves sequences of familiar items, in this case phonemes or syllables. We discuss this hypothesis with reference to a variety of recent data, and propose a modelling framework, based on the primacy model of immediate serial recall (Page & Norris, 1998), that seeks to identify common underlying mechanisms.

Introduction

In this chapter, we will be exploring the possibility that the mechanism responsible for the Hebb repetition effect (Hebb, 1961) is essentially the same as that underlying the learning of phonological word-forms. In his influential paper, Hebb demonstrated that the immediate serial recall (ISR) of a given digit-list improved to the extent that that same list had been repeatedly presented (and recall attempted) on every third trial of a set of ISR trials. List repetition was unannounced, and Hebb showed that learning was seen regardless of whether subjects reported having noticed this manipulation. This latter finding has subsequently been supported by a number of other studies (e.g., McKelvie, 1987; Stadler, 1993). To this extent, it can be said that the Hebb repetition effect is sometimes implicit, even though it is manifested in the performance of an explicit serial recall task.

To those who are disposed to make the distinction, the Hebb repetition effect is a paradigmatic example of the transfer of information from short- to long-term memory. To be specific, the effect appears to involve the gradual development of a durable representation of the item and order information in a given list that can enable that list's subsequent recognition and enhanced recall. In recent years, a number of quantitative models of the ISR task have been developed (e.g., Brown, Preece, & Hulme, 2000; Brown, Neath, & Chater, 2007; Burgess & Hitch, 1992, 1999; Farrell & Lewandowsky, 2002; Henson, 1998; Neath, 2000; Page & Norris, 1998) and we will discuss the Hebb effect here in relation to our own primacy model (Page & Norris, 1998). The primacy model can be thought of as an implementation of the phonological loop component of the working memory model (Baddeley, 1986; Baddeley & Hitch, 1974), inasmuch as it comprises a mechanism specialized in the ordered recall of speech-based material retained over the short term. Indeed, in recent work (Page, Madge, Cumming, & Norris, 2007), we have characterized both the phonological loop, and our model of it, as a system that permits the reproduction of a short portion of speech that has just been heard or, alternatively, that has just been recoded from the corresponding visual stimulus. In this conception, the phonological store is identified with a high-level utterance plan that, when executed, will result in the repetition of a stimulus list. In our model, the primacy gradient of activations that is instantiated across localist representations of list-items constitutes just such an utterance plan.

It is the identification of the phonological store/primacy gradient with a speech reproduction system that raises a question regarding the "ecological" counterpart of the Hebb repetition effect. Just as we presume that the phonological store is not available solely to enable laboratory performance of the ISR task, we can also presume that the learning processes seen in the Hebb repetition effect are not similarly confined to the laboratory conditions under which they have been demonstrated. Baddeley, Gathercole, and Papagno (1998) have drawn together data supportive of their view that the phonological store is a crucial component of the learning of phonological word-forms. They review evidence from a considerable variety of studies: neuropsychological work with the so-called "short-term memory patients", whose catastrophic performance on auditory ISR is invariably accompanied by a specific deficit in the learning of novel phonological forms; developmental research that shows that, at the early stages of word-learning, vocabulary size can be predicted by ability in ISR and in nonword repetition (NWR); research with gifted adults, which shows an association between language learning ability and immediate memory span, independent of IQ; and research with people with learning disabilities that suggests the same association. Given this proposed dependence of word-form learning on a mechanism that is able to retain sequence information in the short term, it is natural to go further and to suggest that the Hebb repetition effect is a laboratory analogue of word-learning. In naturalistic word-learning, repeated hearings of a novel word, that is a novel sequence of phonemes, can lead to the establishment of a new long-term representation of that sequence in memory. The result of this learning process would comprise the entry of that novel word-form into the mental lexicon.

A relationship between the Hebb repetition effect and the learning of phonological word-forms has, we think, surface plausibility. If the Hebb repetition effect were completely unrelated to word-form learning this would imply that there were two quite distinct mechanisms both of which used phonological short-term memory to learn new phonological sequences. This is particularly the case given the known association between short-term memory and word-form acquisition noted above. To give an example, in a reasonably typical Hebb-type experiment, using auditory letters rather than digits as list-items, one might be asked to recall the list "B J F M L". This would involve repeatedly hearing the sequence, perhaps some rehearsal of it, and repeated recall attempts. Is it likely that this process engages quite different mechanisms from those engaged when one is asked to learn the novel word "beejayeffemmelle", a task that is also facilitated by repeated presentation and (see later) repeated recall? We think not, though some might disagree. Of course, the learning of a novel word is not normally going to occur in isolation, being more likely to be encountered in, say, the learning of an object-name. However, as Baddeley et al. (1998) make clear with reference to such name-learning, it is the learning of the phonological word-form that correlates with short-term serial memory ability, rather than the secondary learning of the association between name and object.

In what follows, we will first supplement this intuition with data from a number of experiments that demonstrate that properties of the Hebb repetition effect align with a number of properties that one would expect to obtain for a word-form learning mechanism. While no amount of such evidence can *guarantee* the existence of a common mechanism, we will maintain that application of Occam's razor places the onus of proof on those who deny it. Then, in the latter part of the chapter, we will offer a qualitative framework, based around our model of short-term ordered memory, within which both Hebb repetition effects and the learning of phonological word-forms can be simulated.

Properties of the Hebb repetition effect

The Hebb repetition effect is not critically dependent on spacing

Following Hebb's (1961) finding of learning for repetitions spaced at every three trials, Melton (1963) claimed that the Hebb repetition effect was significantly reduced at spacings of every six trials and was not evident at all at spacings greater than this. Clearly this result is less than encouraging for our working hypothesis, as it would seem to imply that a novel word would need to be repeated at very short intervals for it to be incorporated into the lexicon. While it is possible that such massed presentation might well assist in the learning of a novel word, it would be preferable to our position if this were not a necessary requirement. In Cumming, Page, Norris, Hitch, & McNeil (2005), we decided to explore the issue further. In three experiments involving the recall of lists of seven or eight single-syllable words, we found an interaction between the spacing of list repetitions and the item-set from which the intervening, nonrepeating, "filler" lists were constructed. To be specific, when all lists (repeating and filler) were constructed by rearrangements of the same fixed set of items (in this case, words), repetition learning was present weakly, if at all, for spacings of either three or six lists. However, when filler lists were constructed from a different set of items from those used in the repeating lists, robust repetition learning was observed at spacings of three, six, nine and twelve lists, with no diminution of the learning effect with increasing spacing. Indeed, if anything, the repetition learning was stronger for 12-apart spacing than for the other spacing conditions.

How does this finding relate back to our working hypothesis regarding the learning of phonological word-forms? Certainly, the confirmation of repetition learning at 12-apart spacing, approaching the highest spacing that can reasonably be achieved in the Hebb (1961) paradigm while maintaining the good humour of participants, is encouraging. Had we found that 12-apart learning was never possible, then we would have had to consider rejecting the word-learning hypothesis. At the same time, the finding with respect to the composition of filler lists had the benefit of demonstrating consistency between our results and those of Melton (1963): in his experiments, Melton had used a closed set of list-items, namely, the digits 1–9, thus inadvertently reducing the extent to which learning could be seen at large spacings. With regard to word-learning, therefore, we predict that it will be difficult to learn, say, a novel word if repeated presentations of that word are spaced by "phonemic anagrams" of that word, that is, by other novel words comprising reorderings of the repeated word's constituent sounds. Given that the average word-length, in English at least, is in the region of 7 phonemes (this is the average over word types, the frequency-weighted token average is about 3.5 phonemes, Cutler, Norris, & Sebastián-Gallés, 2004) each drawn from a pool of around 45, it seems unlikely that the learning of any given novel word will be plagued by such a concentration of phonemic anagrams as are seen in the single-pool variant of the Hebb experiment. Notwithstanding Melton's work, therefore, the working hypothesis lives on.

Multiple lists can be learned simultaneously

In the traditional Hebb (1961) paradigm, learning is restricted to that of a single repeating list in a background of nonrepeating fillers. One would not be surprised, though, to find that learning could subsequently be found for a new repeating list in a fresh learning block, and this is indeed what is found (e.g., Cumming, Page, & Norris, 2003). Is it possible, though, to learn a number of interleaved but repeating lists at the same time? This would certainly be a desirable property if our word-learning hypothesis is to be maintained. The construction of a lexicon would be a rather pedestrian affair if each novel word had to be learned completely before one could begin to learn the next one. Fortunately, the third experiment in Cumming et al. (2005) permitted a clear response to that question. As noted above, repetition learning was demonstrated in that experiment for spacings of 6-apart, 9-apart and 12-apart. Moreover, the experimental design was such that at any given point, the learning of lists at the three different spacings was simultaneously under way. In other words, even though repetitions of the lists at the three different spacings were interleaved with each other and with nonrepeating fillers, robust learning was still observed. This result is self-evidently a good thing regarding our working hypothesis, but it is certainly not self-evident from all perspectives. For example, this finding is difficult to explain from the point of view of a position–item association account of Hebb repetition learning. In such an account (e.g., Burgess & Hitch, 1999), improved performance is explained by the strengthening of associations that are presumed to exist between representations of list-items and representations of the list positions in which they are repeatedly presented. This account has already fared badly in a number of subsequent tests (e.g., Cumming, et al. 2003; Hitch, Fastame, & Flude, 2005a) but its plausibility is weakened still further by simultaneous learning of several different lists. Indeed, in unpublished experiments, we have observed that five interleaved repeating lists can be learned simultaneously. This clearly poses major problems for any account based on learning position–item associations. Each position would have to be associated with multiple items, with no way of knowing which of the associated items should go with each list. This weakness is acknowledged in the latest version of the Burgess and Hitch model (Burgess & Hitch, 2006). In this modified model, multiple banks of positional context units compete for the representation of given lists, in much the same way as multiple, localist, chunk units compete for the same privilege in earlier models by Cohen and Grossberg (1987), Nigrin (1993) and Page (1994), as well as in the model described below. It is hard to see, though, how crosstalk between similar lists can be avoided unless parameters are set such that each list becomes associated with a different set of context units. In that case, though, one would wonder what was "positional" about the model at all, since each bank of context-units might just as well be represented by a localist unit exactly as in the earlier models. Note that one of the main pieces of data that drove the development of positional models was the finding that intrusion errors often involved the recall of an item that had appeared in the same position in a previous list (e.g., Henson, 1999). In a model in which every new list is associated to a novel set of positional-context units, one would not expect to see this form of positional error at all.

There is also a question, in any positional model that resorts to multiple banks of context units, of the account that it might offer of the relationship between immediate serial recall, nonword repetition and the learning of phonological word-forms. Burgess and Hitch (2006) are clear that their modified model only operates at the lexical level, that is, it represents sequences of words. And yet the relationship between ISR and word-form learning seems to call for a common ordering mechanism for words in lists and for sounds in words (see also Gupta, this volume). In our opinion, this calls for a class of model that can operate across the levels of a hierarchy, able to represent order across, say, phonemes, syllables, and words, as the need arises. Gupta's (this volume) model is certainly of this type, having a (single) bank of context units that connects to each of the model's hierarchical levels, driving serial recall at each. (Of course, this property of Gupta's model makes it less appropriate for modelling the Hebb effect.) For Burgess and Hitch, it is rather difficult to see how, say, phonemic associations with one of multiple banks of positional codes could ever constitute a content addressable entry in a mental lexicon, at least not in a manner consistent with the representation of lists of words that they themselves espouse. That is not to say that such a "hierarchically enabled" representation is impossible, but in our opinion it would need such modification (e.g., one context bank per word) as to make it functionally indistinguishable from prior localist, ordinal (i.e., completely nonpositional) models (e.g., Nigrin, 1993; Page, 1994) in which such a hierarchy has already been implemented.

Learning in the Hebb repetition effect is long term

Our word-learning hypothesis would certainly be corroborated if the benefits that accrue from repetition learning were seen to be longer lasting than is traditionally demonstrated. The Hebb repetition effect is usually seen in the context of a series of 20 or more ISR trials, perhaps with a final test at the end of the experiment. Less research has been conducted to investigate the longevity of the learning involved. Fendrich, Healy, and Bourne (1991) showed preserved sequence learning a month after the repeated typing of digit sequences, but conditions were rather different from those seen in the Hebb repetition effect. In the fourth experiment of Cumming et al. (2005), we looked at this directly. A large subset of the participants who had taken part in the experiment, in which they learned several interleaved lists, were unexpectedly asked back into the laboratory approximately 3 months later. They were tested on the recognition and recall of the lists they had previously seen repeatedly, compared with those of previous filler lists (presented once previously) and new recombinations of both Hebb lists and fillers (to control for item, as opposed to order, memory). The key finding was that lists that had been presented and recalled eight times in the prior experiment, were both recognized as more familiar and better recalled than any of the controls. This benefit was not just a result of better item-memory, but was specific to the order in which the words had originally been presented. The learning that underlies the Hebb repetition effect is therefore very long lasting indeed. Learning is of sufficient longevity to make it a suitable mechanism for word-form learning, and therefore consistent with our working hypothesis.

Learning in the Hebb repetition effect is relatively fast

Related to the longevity of Hebb-effect learning is the speed with which learning can proceed. Researchers such as Dollaghan (1985, 1987) have offered convincing evidence that word-learning among children is fast, with children learning both the form and meaning of novel words within a couple of presentations. Clearly, it is difficult to make quantitative comparisons between the learning of quite short, novel words, and the learning of word-lists of somewhat above span-length. Nonetheless, a number of observations can be made. First, it is a matter of logical necessity that the Hebb repetition effect must involve an element of learning on the very first exposure to a given list. If this were not the case, then there would be nothing to distinguish the second presentation of that list from the first presentation, and so on, with the result that learning could never "get started". In an unpublished experiment, we have found even this single-presentation learning to be relatively long-lasting. In a public event focusing on human memory, we exposed casual participants to a number of lists of letters for immediate serial recall. After this brief exposure to 10 or so lists, participants were given an envelope and asked to open it later in the event, after about 20 minutes had elapsed. When they opened the envelope, they were shown a number of lists and asked to rate them for familiarity. Some lists had been seen once previously, and others were entirely new. After they had completed their ratings, they were asked to hand them in to a collection point, together with an estimate of the time that had elapsed since the first part of the experiment. With a mean elapsed duration of over 20 minutes, participants rated lists that they had seen once previously as being more familiar than reordered lists of the same items. While such evidence for fast learning is more illustrative than definitive, it is consistent with observations of sustained improvements in mean recall performance of around 3–4% per repetition in traditional Hebb-effect experiments.

Similar results hold for the nonword repetition paradigm, although once again the relevant data have not often been collected. For example, in an experiment with children's nonword repetitions, we found that the 4/5-year-old participants could recognize, some 4 weeks later, a nonword that had been presented to them on a single occasion. Naturally, if our working hypothesis is correct, then we would expect more traditional Hebb repetition effects to be exhibited by such children too. It is to this issue that we now turn.

The Hebb effect is exhibited by young children

The learning of novel word-forms is clearly something that children do well. If word-form learning and the Hebb repetition effect share an underlying mechanism, we would therefore expect children to show Hebb repetition effects in the context of an ISR task. In a series of experiments, Hitch, McNeil, Page, Cumming, & Norris (2005b) showed that 5-,7- and 10-year old children do indeed show a Hebb repetition effect, although it was somewhat more difficult to find than had been anticipated. This weakness of the observed Hebb effects sometimes necessitated a reduction in the number of intervening filler lists, from Hebb's (1961) two, to one or even zero. Stronger effects were seen, too, when filler lists were taken from different categories (e.g., digits versus letters) from that of the repeating list, somewhat reminiscent of the experiments described above that showed stronger repetition effects for nonoverlapping filler-sets.

Finding Hebb effects in young children, even those young enough to be assumed not to be using strategies such as cumulative rehearsal, is a relief from our word-learning perspective. Nonetheless, the unexpected weakness of the effect might also serve to highlight an important issue in a framework such as ours (and, say, Gupta's, this volume) that assumes a common mechanism for word-learning and list-learning. The issue concerns the hierarchical level at which learning is presumed to take place. In word-learning, the sequence that is being learned is one of sublexical representations, be they phonemes or, for longer words, possibly syllables. These sequences are continuous, in the sense that they do not incorporate significant pauses in between sequence-items, and the sequence elements do not themselves have semantic or syntactic properties (assuming, of course, that they are not morphemes). Word-list learning is, however, somewhat different. The lists are typically punctuated by pauses during which strategic activity, such as rehearsal, might well take place, particularly for adult participants. Moreover, the list elements carry both meaning and syntactic information. Thus, while the data reviewed by Baddeley et al. (1998) strongly suggest a link between mechanisms underlying immediate serial recall and those engaged in word-form learning, we should not be blind to the different contexts in which this mechanism operates in the two tasks. The difference might be particularly important when considering word-learning in young children, for whom a list of words for immediate serial recall might lack the coherence of a novel word for repetition and/or learning. It is certainly possible that adult strategies such as list-rehearsal serve both to establish coherence across inter-word pauses and, of course, to provide additional list-repetitions, both of which would be either unavailable, or less available, to children. This would tend to weaken Hebb effects in children. It might also be that children are more adversely affected than are adults by the composition of the nonrepeating filler lists: this is certainly suggested by the manipulations of spacing and of filler-set. This factor might well interact with a lack of list-coherence, or a lack of focus at the whole-list level, to render the Hebb repetition effect still weaker in child participants. The fact that the Hebb repetition effect is seen at all is, however, consistent (though not exclusively so), with the idea of a mechanism shared with the word-learning process.

Partial lists can be learned

The issue of list-coherence also arises in relation to the learning of parts of lists. Schwartz and Bryden (1971) were the first to show that list-learning was prevented when the first two items of a list were changed on each "repetition". Hitch et al. (2005a) confirmed and extended these results. Together, these results have been used to falsify early positional models of the Hebb effect (e.g., Burgess & Hitch, 1999), and they are also uncomfortable for those that might want to look for an account of the effect based on associative chaining from one list-item to the next. From our perspective, these important results, together with others (e.g., Cumming et al., 2003; Page, Cumming, Norris, Hitch, & McNeil, 2006), suggest a chunking account of the Hebb repetition effect (see below), in which chunk access is conditional on the correctly ordered arrival of all list-items, with the consequence that it fails definitively if early list-items do not match. Nigrin (1993) and Page (1994) both published models of this kind. Moreover, this property is very reminiscent of (and, we will claim, functionally identical to) the well-known cohort effects seen in lexical access (Marslen-Wilson & Tyler, 1980).

Naturally, if list-learning is essentially chunk-learning, then it should be possible to learn partial lists provided the repeating portion is represented as somewhat distinct from the remainder of the list. Using Bregman's (1990) terminology, we should be able to find learning of partial lists that are *streamed* separately from other list-items. A number of preliminary results suggest that this is indeed the case. In two so-far unpublished sets of experiments, we investigated this issue in rather different ways. In the first, we showed that a sequence of 6 digits that appeared repeatedly within otherwise random longer 10-digit sequences, was not well learned when the point at which the repeating portion began in the longer list varied between positions 2, 3 and 4: performance on the repeating portion never differed from that on the equivalent positions of nonrepeating lists. When, however, the repeating portion was presented in a different voice from the remainder of the list (i.e., it was streamed separately by voice), then learning of the repeating portion was evident, even when the entire 10-item list had to be recalled on each occasion. This is reminiscent of results from Hughes and Jones (2004), who showed that the Hebb repetition effect is abolished by having items within a given repeating list presented in different voices. In the second series of experiments, we showed that individual groups could be learned in an experiment in which word-lists were grouped by pauses. In these experiments, the repetition was that of individual groups rather than that of whole lists, and the relative position of repeating groups could vary from trial to trial. This is, perhaps, analogous to the appearance of a novel word in various different contexts over the course of learning. Overall, the results suggested that learning did indeed take place at the level of sublist chunks, where the chunk boundaries were circumscribed by the grouping structure of the overarching list. The learning of these sublist chunks did not seem to comprise the establishment of associations between list-items and either within-list or within-group positions. Nor did they appear to correspond to the learning of chains of associations, since the recall advantage accrued to all items in learned chunks including the first item, no matter where the chunk appeared in the overall structure of the list. We concluded that learning comprised the establishment of new chunks in longer-term memory, with the subsequent recognition of these chunks serving to assist in their recall. The idea of chunk boundaries aligning with group boundaries and, by extension, stream boundaries, is completely consistent with the early findings of Winzenz (1972), who showed, among other things, that list-learning by repetition did not extend to "repetitions" in which the ordering of items remained the same, but the grouping structure changed from trial to trial.

The Hebb effect does not require recall, but is strongest when recall is attempted

Past research (Cohen & Johansson (1967a, 1967b; Cunningham, Healy, & Williams, 1984) has appeared to show that a Hebb repetition effect only accrues when recall is attempted on each relevant trial of the learning phase. In Page, Cumming, Norris, Hitch, and McNeil (2005), we looked at this issue a little more closely, using a somewhat stronger manipulation of the Hebb effect, with more repetitions and a nonoverlapping filler-set. Our conclusion was somewhat different from that of prior work: while Hebb repetition effects were certainly stronger when recall was attempted, we did find evidence for both increased familiarity and better recall of lists that had been presented repeatedly but never recalled. In these experiments, covert recall of the auditorily presented stimulus lists was prevented (as far as was possible) by requiring the reading of another word-list from the screen during what would have otherwise been the recall phase of the trial. Using this design, we were additionally able to show that recall, and to a lesser extent recognition, was enhanced for a list that had been repeatedly read from the screen, but never heard as a stimulus list and, hence, never recalled.

Finding significant Hebb repetition effects under presentation-only conditions is positive from the point of view of our working hypothesis. Word-form learning is clearly able to proceed in children too young to be able to produce the words themselves. If list-production were absolutely necessary in the Hebb repetition effect then we might be forced to explain the different pattern of results. Having said that, our experiments did clearly demonstrate that repetition effects were stronger when recall was attempted and that this was apparently not solely to do with the number of times a list, or a recalled approximation of it, were recalled overall. Do we see the corresponding effect in word-form learning? That is, are word-forms better learned if participants attempt to recall them overtly during learning? To us, the intuitive answer is yes. Casual observation suggests that infants often repeat a portion of what has been said to them, where that portion might correspond to the last word or, later, the last few words. But "data" is not the plural of "anecdote", and we have found it difficult to find any literature that answers the question directly. In a collaboration with Tania Zamuner, therefore, we have carried out some preliminary experiments to look at the matter. Over three such experiments with 4- to 5-year-old children, we did find that repetition of a novel word enhanced its subsequent recognition, even when overall frequency of occurrence (hearings plus recalls) was controlled.

A unified framework for modelling the Hebb repetition effect and word-form learning

In what remains of this chapter, we will outline a framework for modelling both the Hebb repetition effect and the learning of other sequences such as those corresponding to phonological word-forms. The framework is based on models going back to Grossberg (1978), via Cohen and Grossberg (1987), Nigrin (1993) and Page (1994), but departs from each of these in some important ways. Space forbids a detailed quantitative description of the mechanics of the model; rather we will outline the general principles underlying the operation of the model, relating them back to some of the issues covered above.

The model is a localist, connectionist model: that is to say, the presence of any item that can be said to have been learned by the model will be indicated in the model by the maximal activation of at least one connectionist unit, such that that unit does not respond maximally in the presence of any other item (see Page, 2000, for much more on the issue of localist representation in connectionist models). In the context of immediate serial recall of lists of letters, digits or familiar words, it will be assumed that the model already contains at least one localist representation of each list-item. That is what it means, in the model, to say that an item is familiar. In addition, it will be assumed that there are also individual units representing known sublexical items, such as phonemes (or, more accurately, allophones), and possibly syllables. Because their activation signals the occurrence of the corresponding item, these units will be known as "occurrence units", and the total collection of such units will be known as the "occurrence layer". The topology of the model is illustrated in Figure 7.1. Importantly, at any one time, there will be units available in the occurrence layer that have yet to be committed to the representation of any item; in line with the prior models mentioned above, these will be referred to as "uncommitted" occurrence units. There is only one layer of occurrence units, and that can contain within it, say, some units representing phonemes, some units representing syllables, some units representing words,

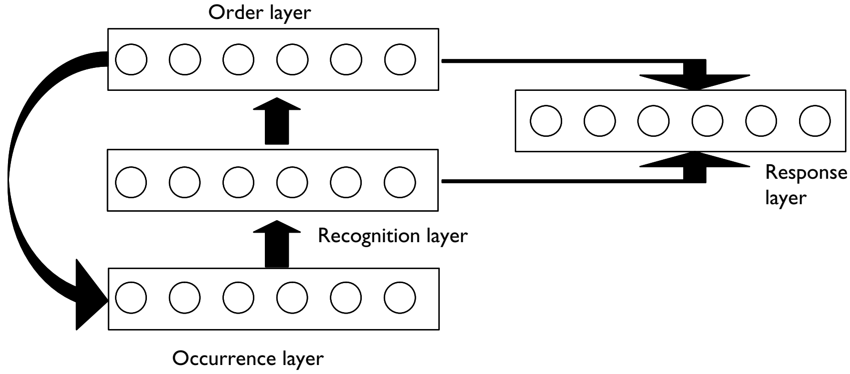

Figure 7.1 The topology of the model. Black arrows indicate one-to-one connections between layers. The order layer has excitatory within-layer connections; the recognition layer has inhibitory within-layer connections.

and some units representing familiar sequences of words (e.g., "The cat sat on the mat"). This single-layer structure is in contrast to the hierarchical structure seen in the predecessor models (e.g., Nigrin, 1993; Page, 1994). For our purposes, we will assume that the phoneme units are primary, in the sense that their activation is driven directly by pattern recognition processes exterior to the model; this assumption is founded on the idea that the recognition of phonemes does not constitute sequence recognition. By contrast, all the other occurrence units are secondary, in the sense that their activation is driven by the sequential activation of other units in the occurrence layer.

Immediately above the occurrence layer, is another layer of units called the recognition units, each of which is in a one-to-one relationship with a corresponding occurrence unit. Strong activation of a recognition unit signals the recognition of the associated pattern within the model. Recognition units compete with each other to signal recognition in the same way as do masking field units in the model of Cohen and Grossberg (1987) and the later variants (Nigrin, 1993; Page, 1994), and in a manner that is commonly described as lexical competition in models such as the shortlist model of word recognition from continuous speech (Norris, 1994). In order briefly to characterize the difference between the occurrence layer and the recognition layer, consider presentation of the stimulus "CAT". In the occurrence layer, five units will activate: those corresponding to the three phonemes (actually probably allophones) in the word and those corresponding to the familiar words CAT and AT, all of which are in some sense present in the stimulus signal. By contrast, at the recognition layer, competitive processes will ensure that only a single unit activates strongly, namely that corresponding to the word CAT. The competitive processes at the recognition layer therefore indicate the most likely lexical parsing, based on the entirety of the stimulus signal. Such processes have been described in great detail elsewhere, and will not be our focus here. In relation to this simple example, though, it is important to note that had the words CAT and AT not been familiar, then there would have been no such lexical entries to suppress activation of units corresponding to the word's constituent phonemes, and it would have been these three units that would have activated at the recognition layer. Suppression of "lower-level" representations (e.g., phonemes) by higher-level ones (e.g., words) is what constitutes chunking in this model and those from which it was developed.

There exists a third layer of units driven once again in a one-to-one and unidirectional fashion by the units in the recognition layer. This layer is called the "order" layer and it stores item and order information corresponding to the sequence of recognized words. The representation of serial order in our model is, unsurprisingly, the same as that in our primacy model of immediate serial recall (Page & Norris, 1998). (Indeed, the primacy model itself is based on elements of the earlier work mentioned above and, in particular, by Grossberg, 1987.) To be specific, the recognition of a sequence of familiar items at the recognition layer, be they phonemes/syllables in a previously unfamiliar word (as in the nonword repetition task) or familiar words (as in the ISR task), will result in a primacy gradient of activation at the order layer, such that units corresponding to items presented earlier in the sequence will have the higher activation. Once an item is represented at the order layer, its corresponding representation is reset at the recognition layer. Again, the details of the primacy model are described in sufficient detail elsewhere to preclude a reprisal here. Finally, each order-unit projects activation back down to its corresponding occurrence unit, and this information is used in the establishment of new occurrence units that come to represent sequences of "lower-level" items.

Before reflecting on some of the properties of the Hebb repetition effect and of word-form learning described above, we shall briefly outline the manner in which higher-order sequence representations can be learned. Again, space precludes a detailed quantitative account, but the qualitative account we offer here should be sufficient to elucidate the framework in which the model operates. As noted previously, the occurrence layer always contains a number of units that are uncommitted to any pattern. These uncommitted units are weakly connected to quite a large proportion of other occurrence units, but the "synaptic weights" on these connections are small and relatively homogeneous. These weights do not, therefore, strongly constrain the circumstances in which a given uncommitted unit activates (see below). By contrast, an occurrence unit that has learned to recognize a given sequence will have a pattern of weights on its input connections that is derived from the primacy gradient of activations representing that sequence. To give an abstract example, an occurrence unit that comes to encode the sequence ABCDE will have connections from the occurrence units for A, B, C, D and E, such that the weight on the connection from the A-unit to the

ABCDE-unit will exceed that from the B-unit, which in turn will be larger than that from the C-unit, and so on. There will thus be a primacy gradient in connection weights, which is essentially a scaled copy of the primacy gradient in activations across order-units that is repeatedly instated during repeated presentation of a novel word-form (where A, B, C, D and E represent, say, phonemes) or a new word-list (where the letters represent list-items).

The way in which occurrence units activate exclusively in response to their learned sequence is also different from that outlined in previous models of this class. In early work (e.g., Cohen & Grossberg, 1987), activation of a higher-level occurrence unit was a function of the dot product between the vector representing primacy-gradient activations and the vector representing that unit's incoming synaptic weights. So if the primacy gradient of activations across representations of the letters A, B and C was 5, 4.5 and 4 respectively, and the corresponding weights to the ABC-unit were 10, 9 and 8, the dot product would be equal to $(5 \times 10) + (4.5 \times 9) + (4 \times 8) = 122.5$, and the ABC-unit would activate as a function of this input. Later work (e.g., Nigrin, 1993; Page 1994) identified problems with the use of the dot product and more complex calculations were developed. All these calculations depended, however, on the combined presence of a set of synaptic weights and a primacy gradient of activations corresponding to the stimulus list. We now believe any such calculation to be unsustainable, at least in the framework we are proposing.

There are two main reasons. First, in the framework proposed above, for there to be a primacy gradient corresponding to each of the items A, B and C in that order, then each of them in turn must have won a competition for recognition at the recognition layer (since activation at the order-layer is contingent on such a sequence of recognitions having occurred). But if each of the constituent items (A, B and C) must already have been recognized before the (dot product) input to the ABC-unit can be calculated, there would no longer be any opportunity, at the recognition layer, for activation of the ABC-unit to suppress activation in the units corresponding to its constituent items. The only reason that previous versions of the model were able to avoid this problem was that they incorporated a hierarchical structure with units corresponding to, say, phonemes at one level, words at another, and word sequences at another. This causes its own problems, which are beyond the purview of this chapter.

Second, it will be a fundamental part of the framework that we are suggesting that so-called short-term memory patients, whose immediate serial recall, nonword recall and word-form learning are very poor, suffer from not being able to form a primacy gradient of activations in response to a stimulus sequence. However, it is well known that such patients are perfectly able to recognize sequences that they learned prior to their deficit. Ipso facto, such recognition cannot depend on the presence of a primacy gradient, or any dot product calculated using it. In short, in our framework, a primacy gradient in activations is necessary during sequence (word-form) learning, since the primacy gradient in synaptic weights is a learned version of the primacy gradient in activations, but is not necessary for subsequent sequence (word-form) recognition.

Our new mechanism for the activation of sequence-representing occurrence units does not, therefore, require the presence of a primacy gradient. It does, however, require the occurrence of each of the constituent items in the correct rank order. Again, a detailed mathematical exposition would be out of place here, but the essential idea can be summarized quite succinctly. Suppose that an ABC-unit (i.e., an occurrence unit representing the sequence ABC), has incoming weights equal to 10, 9 and 8 from the A, B and C units respectively. Assume also that when each of the units A, B and C is activated, each emits a pulse of activation of unit magnitude to all connected units, with that pulse's effect being modulated (multiplied) by the weight of the connection. Further suppose that the ABC-unit has a threshold, and that only incoming pulses above that threshold value can cause the ABC-unit itself to fire. Finally, assume that each time the ABC-unit fires, its threshold is lowered by a given value, say 1. If we assume that the resting threshold for the ABC-unit is, say, 9.5, then it is easy to see that the ABC-unit will only fire its maximum three times, if the items A, B and C, arrive in the rank order encoded in the primacy gradient of weighted connections. When A occurs, a unit pulse is sent from the A-unit to the ABC-unit, modulated by the connection weight 10 between the two. This signal 10 exceeds the threshold of 9.5, so the ABC-unit fires once and, as a consequence, its threshold is reduced to 8.5. When the B-unit fires, it causes a signal of strength 9 to arrive at the ABC-unit. This exceeds the new threshold, so the ABC-unit fires again and the threshold is accordingly lowered to 7.5. The arrival from a signal of strength 8 from the C-unit causes a third firing of the ABC-unit. In this way, the arrival of the items A, B and C in the correct order as specified by the ranking of incoming weights to the ABC-unit causes that unit to fire maximally. By contrast, applying the same procedure, the list ACB will only fire the ABC-unit once (in response to the A), and so on for other approximations to the learned list. In the full model, the firing is probabilistic, and there are multiple pulses per event, but the essential mechanism is the same.

Using this activation mechanism, it is clear that occurrence units that have learned a given sequence will activate whenever that sequence occurs in the input. Whether the sequences are fully recognized depends on the competitive processes at the recognition layer. These processes in turn influence what is stored at the order layer. For example, if the network has occurrence units for ABC, ABCD and DE, then it is perfectly possible that the stimulus ABCDE will be parsed as a recognition of ABC followed by a recognition of DE, which will be represented more compactly as a list of two familiar items (ABC followed by DE) at the order layer than the occurrence of the five items A, B, C, D and E would otherwise suggest. To reiterate, this is what is meant by chunking in this model. The competitive relationships that implement such a parsing at the recognition level are the subject of a good deal of previous work (e.g., Nigrin, 1993; Norris, 1994; Page 1994) and will not be discussed further here.

In outlining our framework, it only remains to specify the conditions under which learning of a new sequence will occur and of what that learning comprises. It will be simplest to use as an example the Hebb repetition-learning of a list of familiar items, say words. In this case, each ISR trial involves the presentation of a sequence of stimuli that is, in some sense, coherent. Sequence-items almost always come from a single spatial source and are presented in a common voice (for auditory) or format (for visual). In the absence of grouping cues, the boundaries of this coherent sequence are clearly delimited. In other words, if a sequence is to be learned at all, then it is very clear what that sequence comprises. Any such coherent sequence will be learned, to some extent, after a single presentation. This is because the system has no way of knowing whether the sequence is going to become frequent in the future (see above regarding the necessity of first-trial learning). It is for this reason that any novel coherent sequence in memory will get partially learned by the most active uncommitted unit that is available when the sequence ends. When partial learning has occurred, but before full commitment is achieved, a unit is said to be provisionally committed.

Because of the way uncommitted units respond to incoming sequences (i.e., sequences of firings of other occurrence units), there is always very likely to be at least one that is active at the end of any given sequence. As noted above, uncommitted units have low incoming weights, but remain responsive to incoming stimuli because they also have a low threshold of activation. Suppose, for example, that an uncommitted occurrence unit is widely (or even fully) connected to other occurrence units in the occurrence layer, with weights that are distributed randomly, but fairly tightly, around the value 1. If we assume also that the threshold is set to around 95% of the maximum weight (as it was for the committed unit discussed above) and that as the unit fires, the threshold reduces by appropriately scaled increments but does not drop below a minimum value of, say, 0.5. It should be clear that such an uncommitted unit will respond to a large variety of possible sequences. For this reason, at the end of any given coherent sequence, there will always be activated uncommitted occurrence units ready to learn the sequence. If the sequence is already (somewhat) familiar, however, there will also be an activated committed (or provisionally committed) unit. In order to prevent proliferation of units committed to the same sequence, a competitive mechanism is needed that will prevent learning at uncommitted units for sequences for which a committed unit already exists. Fortunately, there exists sufficient mechanism at the recognition layer, which is inherently geared towards competition. Essentially, for any coherent sequence for which there is no occurrence unit that spans the sequence, learning will accrue to the most active uncommitted/provisional unit. On the first presentation of a novel list, that will be an uncommitted unit; on the second and subsequent presentations it will be the same unit, which is now provisionally committed, until eventually the unit becomes fully committed to representing the given sequence. It is only at that point, when the unit becomes fully committed, that the unit will be enabled to suppress recognition of its constituent parts by competition at the recognition layer, and parse the sequence by its activation alone.

All this functionality can be achieved by ensuring that: uncommitted/provisional recognition units (i.e., those recognition units corresponding to uncommitted/provisional occurrence cells) can compete with each other; that their activation can be suppressed by competitive activation of other committed recognition units; but that they cannot suppress activation of those committed recognition units until they have accrued sufficient learning to become fully committed themselves. This asymmetric competitive arrangement between committed and uncommitted/provisional recognition units allows uncommitted units to activate when no other existing units better parse the sequence, but ensures that such activation does not prevent recognition of the sequence elements and the generation of the primacy gradient that the uncommitted unit needs to learn. When such an uncommitted unit sustains high activation at the recognition layer, its occurrence unit learns the primacy gradient that is projected down from the order layer to the occurrence units of the list-items that drove its activation. Over the course of several such learning trials, during which a primacy gradient builds up in the provisional unit's incoming weights, the previously uncommitted unit becomes specialized in the recognition and processing of the learned sequence.

How does the learned activation of a new occurrence unit enable progressively better recall of, say, a repeating list in a Hebb-type experiment? Each recognition unit is assumed to be connected to a production unit located in a fourth layer called the production layer. Each unit in the production layer is connected to the corresponding units in both the recognition layer and the order layer. For this reason, when a primacy gradient of activations is in place at the order layer, then the equivalent gradient is copied into the production layer (perhaps when a signal to recall is issued). Moreover, if a provisionally committed unit is active at the recognition layer (which it will be if the input stimulus is consistent with its provisionally learned pattern) then its corresponding production unit will also be active. Under these circumstances, this provisionally active production unit will learn the activation pattern across the other active production units. That is, it will learn a primacy gradient across its outgoing weighted connections that will enable it to produce the correct activation gradient across its constituent items, when it is activated as an isolated chunk. During this learning period, activation of the provisional production unit for an entire list, say ABCDE, will project its partially learned primacy gradient onto the production units of its list-constituents (i.e., the production units for A, B, C, D and E). This small learned gradient will add to that projected from the order layer, to produce a primacy gradient that is higher in activation, and steeper (in terms of activation step) than the primacy gradient corresponding to a list for which no provisional production unit exists. Given that a steeper primacy gradient will result in fewer order errors (when noise disrupts the correct rank order of gradient items) and a higher-activation gradient will result in fewer omission errors, the gradual learning of a primacy gradient in the outgoing weights of a provisionally committed production unit will result in the gradual recall improvement characteristic of the Hebb repetition effect. Once the unit is fully committed, activation of it alone will be sufficient to permit recall of its encoded list.

Finally, while we have described the model in terms of list recall, the same mechanism will apply to the learning of word-forms. Word-forms, in the conception presented here, are lists much like any other. The application to novel word-forms presented in isolation is straightforward. Slightly more involved is the application to word-forms presented in continuous speech. The principal issue here is what constitutes a coherent sequence (i.e., a possible word-form). Remember that in the context of the Hebb repetition effect, it is relatively clear where the boundaries of the to-be-learned list are located. In continuous speech the boundaries are less clear. That having been said, a good deal of research has shown that even quite young infants are sensitive to a variety of cues as to possible/likely word boundaries. These include pauses, lexical stress, lengthening, bigram probabilities, allophonic distinctions, etc. We suggest that the learning of such probabilistic cues, which has been shown both to precede and to influence word-form segmentation and learning, can permit good hypotheses to be made as to where coherent sequences begin and end. Regrettably, space restrictions forbid a more detailed exposition – this would probably require a chapter to itself.

Conclusion

To conclude, we will now briefly consider each of the properties of the Hebb effect and of word-form learning that were described in the first half of this chapter, and relate them to the model developed in the second half. The Hebb repetition effect was found not critically to depend on spacing; in our model, the learning carried out by a provisionally committed list-unit would be easily able to span any reasonable spacing. The use of filler lists all derived from a single-item pool shared with that of repeating lists will lead to a proliferation of units all provisionally committed to lists that are anagrams of each other. Simulations of the threshold-based activation mechanism have shown that in the early stages of learning, anagrams can cause inadvertent (though nonmaximal) activation of a given list-unit. Activation of a large number of provisional recognition units will slow down learning via the competitive relationships that exist between such units at the recognition layer. For this reason, Hebb repetition learning will be faster when filler lists are drawn from a separate item-pool. Multiple lists will be able to be learned simultaneously because, in the model, different units will come to represent those different lists. Given the properties of localist models, such learning can be either fast or slow, and can persist over long periods of time, without any risk of interference with previously stored knowledge. Hebb repetition effects will be found with young children, though the strength of these might be compromised to the extent that the children fail to perceive, say, word-lists as coherent streams. Moreover, partial lists will be learnable if factors relating to perceptual organization (e.g., grouping, streaming) permit somewhat distinct representation of the repeating parts. Finally, an attempt to recall a list will promote learning, not least because the recall attempt will constitute a second (and hopefully veridical) hearing.

In this chapter, we hope to have given some cause to believe that the Hebb repetition effect and the learning of phonological word-forms are related, and we have offered a qualitative description of a localist connectionist framework within which both might be modelled. The long-term viability of this framework will, of course, depend on its ability to furnish quantitative simulations of some of the data discussed above, and more. We are optimistic that it will be able to do so.

References

Baddeley, A. (1986). *Working memory*. New York: Clarendon Press/Oxford University Press.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. H. Bower (Ed.), *Recent advances in learning and motivation* (Vol. 8, pp. 47–90). New York: Academic Press.

Baddeley, A., Gatherede, S., & Papagno, C. (1998). The phonological loop as a language learning device. *Psychological Review*, *105*(1), 158–173.

Bregman, A. S. (1990). *Auditory scene analysis: The perceptual organization of sound*. Cambridge, MA: The MIT Press.

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*(3), 539–576.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*(1), 127–181.

Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, *31*(4), 429–460.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*(3), 551–581.

Burgess, N., & Hitch, G. J. (2006). A revised model of short-term memory and long-term learning of verbal sequences. *Journal of Memory and Language*, *55*(4), 627–652.

Cohen, M. A., & Grossberg, S. (1987). Masking fields: A massively parallel neural architecture for learning, recognizing, and predicting multiple groupings of patterned data. *Applied Optics*, *26*(10), 1866–1891.

Cohen, R. L., & Johansson, B. S. (1967a). Some relevant factors in the transfer of material from short-term to long-term memory. *Quarterly Journal of Experimental Psychology*, *19*(4), 300–308.

Cohen, R. L., & Johansson, B. S. (1967b). The activity trace in immediate memory: a re-evaluation. *Journal of Verbal Learning and Verbal Behavior*, *6*(1), 139–143.

Cumming, N., Page, M. P. A., & Norris, D. (2003). Testing a positional model of the Hebb effect. *Memory*, *11*(1), 43–63.

Cumming, N., Page, M. P. A., Norris, D., Hitch, G., & McNeil, A. M. (2005). Repetition spacing and order competition effects in the Hebb repetition task. Paper submitted to the *Journal of Memory and Language*.

Cunningham, T. F., Healy, A. F., & Williams, D. M. (1984). Effects of repetition on short-term retention of order information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(4), 575–597.

Cutler, A., Norris, D., & Sebastián-Gallés, N. (2004). *Phonemic repertoire and similarity within the vocabulary*. Paper presented at the 8th International Conference on Spoken Language Processing.

Dollaghan, C. (1985). Child meets word: "Fast mapping" in preschool children. *Journal of Speech and Hearing Research*, *28*(3), 449–454.

Dollaghan, C. A. (1987). Fast mapping in normal and language-impaired children. *Journal of Speech and Hearing Disorders*, *52*(3), 218–222.

Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, *9*(1), 59–79.

Fendrich, D. W., Healy, A. F., & Bourne, L. E. (1991). Long-term repetition effects for motoric and perceptual procedures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(1), 137–151.

Grossberg, S. (1978). Behavioral contrast in short-term memory: Serial binary memory models or parallel continuous memory models. *Journal of Mathematical Psychology*, *17*, 199–219.

Hebb, D. O. (1961). Distinctive features of learning in the higher animal. In J. F. Delafresnaye (Ed.), *Brain mechanisms and learning* (pp. 37–46). Oxford: Blackwell.

Henson, R. N. A. (1998). Short-term memory for serial order: The Start-End Model. *Cognitive Psychology*, *36*(2), 73–137.

Henson, R. N. A. (1999). Positional information in short-term memory: Relative or absolute? *Memory and Cognition*, *27*(5), 915–927.

Hitch, G. J., Fastame, M. C., & Flude, B. (2005a). How is the serial order of a verbal sequence coded? Some comparisons between models. *Memory*, *13*(3–4), 247–258.

Hitch, G. J., McNeil, A. M., Page, M. P. A., Cumming, N., & Norris, D. (2005b). Do children show a Hebb effect? Paper submitted to the *Journal of Child Language*.

Hughes, R. W., & Jones, D. M. (2004) *Hebbian auditory sequence learning: The role of passive perceptual processes*. Poster presented at the 45th Annual Meeting of the Psychonomic Society. Minneapolis, USA.

Marslen-Wilson, W., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. *Cognition*, *8*(1), 1–71.

McKelvie, S. J. (1987). Learning and awareness in the Hebb digits task. *Journal of General Psychology*, *114*(1), 75–88.

Melton, A. W. (1963). Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *2*, 1–21.

Neath, I. (2000). Modeling the effects of irrelevant speech on memory. *Psychonomic Bulletin and Review*, *7*(3), 403–423.

Nigrin, A. (1993). *Neural Networks for Pattern Recognition*: Cambridge, MA: MIT Press.

Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition*, *52*(3), 189–234.

Page, M. P. A. (1994). Modelling the perception of musical sequences with self-organizing neural networks. *Connection Science*, *6*(2–3), 223–246.

Page, M. P. A. (2000). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, *23*(4), 443–512.

Page, M. P. A., Cumming, N., Norris, D., Hitch, G. J., & McNeil, A. M. (2005). Exploring the importance of list recall in the Hebb repetition effect. Paper submitted to *Memory and Cognition*.

Page, M. P. A., Cumming, N., Norris, D., Hitch, G. J., & McNeil, A. M. (2006). Repetition learning in the immediate serial recall of visual and auditory materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(4), 716–733.

Page, M. P. A., Madge, A., Cumming, N., & Norris, D. G. (2007). Speech errors and the phonological similarity effect in short-term memory: Evidence suggesting a common locus. *Journal of Memory and Language*, *56*(1), 49–64.

Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review*, *105*(4), 761–781.

Schwartz, M., & Bryden, M. P. (1971). Coding factors in the learning of repeated digit sequences. *Journal of Experimental Psychology*, *87*(3), 331–334.

Stadler, M. A. (1993). Implicit serial learning: Questions inspired by Hebb (1961). *Memory and Cognition*, *21*(6), 819–827.

Winzenz, D. (1972). Group structure and coding in serial learning. *Journal of Experimental Psychology*, *92*(1), 8–19.