

Creating Proactive Interference in Immediate Recall: Building a Dog from a Dart, a
Mop and a Fig

Gerald Tehan
University of Southern Queensland

Michael S. Humphreys
University of Queensland

Mailing Address:
Gerry Tehan
Faculty of Sciences
University of Southern Queensland
Toowoomba, Q. 4350
AUSTRALIA.

e-mail: tehan@usq.edu.au

Running Head: Creating Proactive Interference.

Abstract

Phonemic codes are accorded a privileged role in most current models of immediate serial recall, although their effects are apparent in short-term proactive interference (PI) effects as well. The current research looks at how assumptions concerning distributed representation and distributed storage involving both semantic and phonemic codes might be operationalized to produce PI in a short-term cued recall task. The four experiments reported here attempted to generate the phonemic characteristics of a non-rhyming, interfering foil from unrelated filler items in the same list. PI was observed when a rhyme of the foil was studied or when the three phonemes of the foil were distributed across three studied filler items. The results suggest that items in short-term memory are stored in terms of feature bundles and that all items are simultaneously available at retrieval.

Creating Proactive Interference in Immediate Recall: Building a Dog from a
Dart, a Mop and a Fig.

Most current models of immediate memory are in agreement on two things: Firstly, that speech-based codes play a dominant role in short-term tasks. The second follows from the first in that these phonemic codes quickly become degraded if rehearsal is prevented. This latter feature ensures that at recall, a target must be produced from an impoverished trace. Consequently, quite a number of current models specify a redintegration or deblurring process in which an item is derived from a fuzzy approximation of that item (Brown & Hulme, 1995; Henson, Norris, Page & Baddeley, 1996; Lewandowsky & Murdock, 1989; Nairne, 1988; 1990; Neath & Nairne, 1996; Schweickert, 1993). The specifics of the reconstruction process vary from model to model as a function of representational, storage and retrieval assumptions.

The most obvious instance of where coding and reconstructive processes interact is in the phonemic similarity effect. That is, the fact that items that have similar sound characteristics can often be mistaken for another similar item in the list (Baddeley, 1966; Conrad, 1965). Those models that employ localist representations and localist storage, for example Henson, Norris, Page and Baddeley (1996), explain phonemic confusions by arguing that the short-term episodic trace activates the wrong semantic output node. Where distributed representations but localist storage are used, a similar explanation is provided. Nairne (1990), for instance, matches a degraded short-term trace with items in a search set that consist of undegraded, LTM traces of the list items. Phonemic confusions emerge when the degraded trace is more similar to the Long-Term representation of another item in the list than it is to the Long-Term representation of the target. Where both distributed representations and

distributed storage are utilised (e.g. Lewandowsky & Murdock, 1989), as is the case with a number of PDP, connectionist or neural net models, items lose their identity at storage. Retrieval processes produce a noisy trace that needs to be cleaned up. This clean-up process is often accomplished by an autoassociative neural net in which the net attempts to "build" an item from the trace (Chappell & Humphreys, 1994; Lewandowsky & Li, 1994; Murdock, 1995). Autoassociators of the Chappell and Humphreys type converge on the pattern that represents a target item by either activating missing features or by suppressing activated features that do not belong to the target item. Errors can be generated in that in turning on or turning off features, it is possible that the net will converge on an incorrect item rather than on the target item. However, under this system an incorrect item should share a considerable number of features with the target. That is, the incorrect item should be similar to the target.

It is evident that simple similarity effects are readily explained by most current models of short-term recall. However, we think that connectionist models that specify distributed representations have the ability to handle more complex similarity effects that are not readily handled by more traditional models. In the following set of experiments we look at some of the predictions of the Chappell and Humphreys (1994) model with respect to item interactions in the context of short-term proactive interference (PI) effects. We think that PI effects are proving to be a very useful mechanism for constraining assumptions concerning representation, storage and retrieval over brief retention intervals.

The Chappell and Humphreys model is a connectionist model in which representational, storage and retrieval processes are well specified. The model uses distributed representations in which all items are represented by the same limited set of features. Storage in this model involves the formation of context to item or inter-

item associations. Each association is represented by a different pattern of connection weights between the input and output layers in the network. For present purposes this is a key aspect of the model because items, per se, are never stored, only associations are. This feature makes the model very much like other distributed storage models (Murdock, 1982; Pike, 1984) in which the superposition of memories ensures that the identity of each study item is lost.

Episodic memory access in the model involves finding the intersection of two sources of information. This cuing process can most readily be explained with reference to an example, say cued recall of a short list using a taxonomic category cue. The instructions to retrieve the item from the list that was an instance of the category ANIMAL would involve finding the intersection of the sets elicited context and category cues. In modelling this process, the context cue elicits all list members and suppresses the representations of all non-list members. The category cue elicits the representations of all animals and suppresses all non animals. Because the model uses distributed representations, the only elements that survive suppression are those that are common to both cues. The intersection process creates a noisy output because it is possible that there will be overlap in some of the features of animals that were not presented and list words that are not animals, for example, the phonemic features of lion and the list-word light. The noisy output from the intersection process then serves as input for the autoassociator, in which patterns representing items have previously been stored (a semantic/lexical memory has previously been created). The autoassociator can turn off the noise leaving only the prelearned pattern active or, if some of the elements of a prelearned pattern are initially suppressed, the autoassociator can turn the missing elements on. This process hopefully converges on the pattern representing the animal that appeared on the list.

The interesting aspect of this process centers on what happens when two patterns (e.g. there are two animals on the list) plus noise survive intersection. In such instances Chappell and Humphreys assert that the autoassociator can converge to one of three states. It can suppress noise and one of the two patterns converging on the representation of the other item, or it can converge to a state where all units are turned off or all the units are turned on. In either of these latter cases, no item is produced. One area in which this issue becomes important is in proactive interference (PI) experiments, where two similar items can act as competitors at recall. Before exploring how the Chappell and Humphreys model might advance our knowledge of short-term recall, we review some of our recent work looking at PI effects using a short-term cued recall task.

In the experiments that follow we explore target similarity effects within a short-term cued recall task in which PI is manipulated (Tehan & Humphreys, 1995). On each critical trial in this task subjects study two four-item blocks under directed forgetting instructions. Subjects are told that once they find out that a trial is a two-block trial (one block trials are included in the experiment) they are to forget the first block and concentrate on remembering the items in the second block. On each PI trial two members of a category are presented in the list. The to-be-forgotten, interfering foil is always presented in the first block amid three unrelated fillers and the to-be-remembered target item is embedded among unrelated filler items in the second block. Figure 1 presents an example of one such trial. In the no-interference or control version of these trials, the target appears in the second block without any foil in the list at all. The list items are presented at a one second rate and after the final item has been presented a cue is shown either immediately or after two seconds of verbal shadowing. Subjects are requested to recall the instance of the category that was in the most recent block. Primary interest lies in the extent to which the presence

of a to-be-forgotten foil intrudes as a response to the cue or has an impact upon target recall.

Insert Figure 1 about here

Tehan and Humphreys (1995) have found that on delayed test performance is invariably prone to the effects of the interfering foil. Target recall is lower than in control conditions and recall of the foil is a major source of error. On an immediate test the pattern is a little different. In most situations, recall is immune to the effects of PI. Target recall under interference conditions is as good as under control conditions and the foil is rarely recalled in place of the target. These findings replicate other research testing immediate memory of short lists (Halford, Maybery & Bain, 1986; Wickens, Moody & Dow, 1981). Tehan and Humphreys demonstrated, however, that PI can be observed on an immediate test if the foil and the target rhyme. Thus, when target and foils come from the same rhyming category (e.g. wrench and bench) and the cue is an ending cue (_ENCH) or are rhyming instances of taxonomic categories (e.g. cat and rat with ANIMAL as cue), PI is observed on an immediate test as well as a delayed test.

Tehan and Humphreys explained the materials differences on an immediate test by appealing to the discriminative information that transient phonemic codes supply in short-term tasks. They argued that short-term recall involved both transient phonemic information and longer lasting semantic information. They were able to demonstrate that two seconds of distractor activity was sufficient to attenuate the effects of the phonemic codes, such that on a delayed test recall was based primarily upon the semantic features (Tehan & Humphreys, 1995, Experiments 1 & 2). With the semantic attributes of both target and foil were present, they suggested

discriminating one item from the other became problematic. To explain immunity to PI on an immediate test they first assumed that the phonemic features of the target would be maintained but those of the foil would be lost due to lack of rehearsal and subsequent interference. Given that the semantic features of both target and foil were present, the phonemic features of the target would provide information which would generally allow the subject to discriminate between the target and foil. The only instance where these phonemic codes would not discriminate was if the target and foil where rhymes of one another.

Tehan and Fallon (in press) have sketched out an adaptation of the Chappell and Humphreys model to show how the model might encompass the above explanation. They propose two changes. Firstly, given the role of phonemic codes in short-term performance they argue that some of the features used to represent an item should reflect the phonemic characteristics of that item. The second is to assume that the phonemic features of an item are quickly activated and thus support rapid learning, but they are transient in that they are also very quickly suppressed (see Schneider & Detweiler, 1988 for a discussion on fast weight and slow weight learning in connectionist models).

Given the experimental procedure and the retrieval dynamics of the Chappell and Humphreys model, the list cue activates the list items and the category cue elicits all the items in the category. Given that in interference trials, there are two instances of each category, the representations of both the target and the foil should survive the intersection process. The autoassociator should converge on the target some times and on the foil on other occasions. Thus, the model readily produces PI. However, to produce immunity to PI, the autoassociator must converge on the target and not the foil. That is, the features of the foil have to be turned off on almost all occasions. Tehan and Fallon have argued that in the case of non-rhyming items, the active

phonemic features of the target are crucial in producing this outcome. Given the assumption that the phonemic features of the target survive the intersection process, but those of the foil do not, the input to the autoassociator contains more active features of the target (semantic and phonemic) than of the foil (semantic only). As the autoassociator begins the clean up process, the phonemic features of the target will tend to reinforce the semantic features of the target, while suppressing the noise and the active semantic features of the foil. In this way the autoassociator converges on the target. However, if the phonemic characteristics of the target are also common to those of the foil, as is the case with rhyming items, the chances of the foil being converged upon are very much enhanced, in that the phonemic features of the target support the semantic features of both the target and the foil.

The above explanation for PI on an immediate test with rhyming items boils down to an item interaction effect. That is, parts of one item, in the case the phonemic codes of the foil, are being supplied by another item, in this case the target. However, there is nothing in the model that necessitates that only the target and foil interact. The critical determinant is what features survive the intersection process. If features lose their identity at storage, then it is conceivable that the phonemic codes of the foil could be provided by any other item, say one of the filler items in the list. Thus, if our interpretation of PI effects is correct, we should be able to produce PI effects with non-rhyming instances if the phonemic features of the foil can be provided by another word. The following experiments set out to explore this prediction of the Chappell and Humphreys' model. To preview the results of the study we are able to show that items do interact with each other in the predicted fashion and that PI can be produced on an immediate test when the target and foil do not rhyme.

Experiment 1

The first experiment set out to test the notion that the phonemic codes of an interfering foil could be provided by another list item. By having a rhyme of the foil as one of the filler items in the second block we hoped to produce an approximate representation of the foil that contained both phonemic and semantic components. To give a concrete example, suppose that in a trial subjects see dog as the to-be-forgotten foil, cat as the target and one of the filler items late in the second block is log. We would argue that the semantic features of both the target and foil would survive the intersection process, as would the phonemic features of the target. In addition we think that the phonemic features shared by log and dog would also survive. Consequently, given that we have a noisy representation that contains the phonemic and semantic features of both the target and the foil, the autoassociator is somewhat less likely to converge on the target cat and somewhat more likely to converge on the foil dog, than is the case in the standard interference condition where there is no rhyme of the foil in the list. In other words, we expect that PI will be observed.

In the first part of the experiment we attempted to capitalize upon recency by placing a rhyme of the interfering foil as the last filler item on the list. While recency effects in immediate serial recall are not as pronounced with visual presentation as with auditory presentation, there is often a modest improvement in recall for the last serial position with visual presentation (Crowder, 1976). Furthermore, those models that posit rapid retroactive interference with phonemic codes, still tend to maintain some semblance of recency for visually presented items (Nairne, 1990). Consequently, we thought that our best chance for observing PI would be if the rhyme of the foil occurred in the final serial position.

However, our argument implies that phonemic codes are not limited to the terminal item, they should be present across all serial positions in the second block. Thus, in the second part of the experiment, we put the rhyming foil in either the

second or third serial positions, the within-block positions occupied by the target. This meant that on each trial, the target and the rhyme were in adjacent positions in the second block.

The experiment consisted of three conditions, the standard control and interference conditions, plus the interference condition in which a rhyme of the foil was present. The expectation was that we would observe immunity to PI in the case of the standard interference condition where no rhyme was present, but PI would be observed in the case where a rhyme of the foil was present in the list.. Furthermore, while it is possible to make four types of response on this task; target recall, foil recall, extra-list intrusion, or an omission, the strongest test of our assertions concerns differential rates of recalling the interfering foil. Consequently, while measures of correct target recall and omissions are presented, the primary interest in the following experiments is in the number of block-1 intrusions (recall of the foil from block-1 in place of the block-2 target) that are made.

Method

Subjects

The forty people who participated in the experiment were either first-year psychology students from the University of Southern Queensland who participated for course credit, or second and third level undergraduate students who participated for a ticket in a lottery for a small cash prize. Twenty participated in Part A of the experiment and 20 participated in Part B.

Procedure

At the beginning of the experiment subjects were informed that they would be studying a series of one-block and two-block trials in which a block consisted of four words. However, it was also stressed that at any one point in time they only had to remember that most recent block of four items. Consequently, if the trial was a two

block trial, and this was signified by the presence of an exclamation mark (!) as the block separator, they were to forget the first block and concentrate on remembering the second block because it would be on this block that they would be tested. They were told that one-block and two-block trials would be randomly interspersed throughout the experiment and that since they would not know in advance what type of trial it was, it was in their best interests to treat each trial as a one-block trial until they learned otherwise.

Each trial began with a READY sign displayed on the computer monitor for two seconds. The study items were then displayed individually in lower case at a rate of one word per second, and subjects were instructed to remain silent throughout the presentation of the study items. On two block trials, the block separator, (!), was presented for one second after the fourth word in the first block and before the first word in the second block. At recall, a category cue was presented in upper case for two seconds. On an immediate test the cue appeared immediately after the fourth item in the block. With the appearance of the cue, participants were requested to verbally recall the item from the most recent block that was an instance of the category. Subjects had five seconds to make a response before the next trial began. The experimenter recorded the subjects responses (correct recall, intrusion errors, omissions, etc) on a hard copy of the subject's input file.

Materials

The interfering foils and the target items used in this experiment were selected from the taxonomic categories generated by Nelson and his colleagues at the University of South Florida (McEvoy & Nelson, 1982). The filler items were selected from the unused categories from the South Florida norms and from the Shapiro and Palermo (1970) norms, such that there was no overlap between the category membership of filler and critical items. This ensured that filler items were always

unrelated to the critical items although they tended to have similar characteristics.

That is, they tended to be one or two syllable, concrete nouns as was the case with the critical items. As such, we felt that fillers, foils and targets could not be discriminated on the basis of common word characteristics, and that strategies based upon these features would not be employed.

For the critical items, two instances were sampled from each category. The item selected to be an interfering foil in block-one was a high dominant instance of the category (using controlled association procedures, on average the item was produced by 36% of the subjects and was either the most frequently or second most frequently produced item). The block-two target was a relatively weak member of the category (produced by 1.5 % of the subjects, and was, on average, the twelfth most frequently produced instance in the category). The materials used are presented in the Appendix. The targets and foils were matched for word frequency.

In the standard version of this task, category specific proactive interference is manipulated in the two block trials in that on an interference trial an interfering foil from the same category as the target item is presented in the first block, and the target appears in the second block. On control or no interference trials, a target item is presented in the second block with no related item in the first block. Both parts of the current experiment involved a third condition in which a target and interfering foil are present (an interference condition), but one of the filler items in the second block is a rhyme of the foil. The rhyming instance for each of the interfering foils was generated by going through the South Florida rhyme category norms (Walling, McEvoy, Oth & Nelson, unpublished manuscript). The selected rhyme had the same characteristics as the other filler items in that it was not a member of the relevant taxonomic category and it was in nearly all instances a concrete noun.

In creating the trials for each subject, the target and foil were first randomly assigned to the various experimental conditions. On the ten interference trials the target and foil were then embedded amongst randomly selected filler items, which were also all concrete nouns, such that two four-item blocks were created. In the case of the ten control or no-interference trials, the interfering foil was replaced by a filler item, such that the target was the only instance of the category in the list. On the interference trials, foil and target always appeared in the same serial position (on half the trials position 2 and the other half position 3) in their respective blocks. This was done to reduce the influence of primacy and recency upon absolute levels of recall. The remaining ten two-block trials were interference plus rhyme trials in which the foil was present in block-one and the target and a rhyme of the foil were in block-two. In Part A of the current experiment, the rhyme of the foil always appeared as the last item in the list. In Part B, the rhyme was always presented in positions two or three.

All experiments also contain a number of one-block trials that were also tested via immediate cued recall. These trials were included to ensure that subjects attended to the first block in a two-block trial and to allow serial positions one and four to be tested. Generally speaking performance on these trials indicated that they served their purpose and as such they are not considered any further in any of the analyses. The order of the filler and experimental trials were randomized for each subject. This ensured that subjects never knew in advance, whether the trial would be a one-block filler trial, a two-block interference trial or a two-block control trial.

Results

In reporting and analyzing our data for this and all the following experiments, we have first presented the means for the control, standard interference and interference plus rhyme conditions for correct recall, for omissions and for block-1 intrusions. Extra list intrusions are a minor source of error in all experiments and

appear to have no relationship to the different types of interference trial. In short, it appears that our participants, when in doubt, prefer to make an omission rather than to guess. For the correct recall, omissions and block-1 intrusions we provide 95% confidence intervals for the population means, based upon the pooled error term for the three interference conditions (Loftus & Masson, 1994). This information is supplemented by effect size estimates based upon the following comparisons, which remain constant throughout the following experiments. PI effects are primarily determined by looking at performance on the interference conditions compared to performance on the control trials in which no foil is present in the list. Thus, the first comparison involves target recall in the control and standard interference conditions. Here we expect to see immunity to PI which will be reflected in small effect sizes. The second comparison looks at target recall in the control and interference plus rhyme condition. This is the condition under which we expect to find PI, which will be reflected in larger effect sizes. We also do the same comparisons for omission errors. The final comparison involves the strong test of our hypotheses and involves the difference in number of block-1 intrusions between the standard interference and the interference plus rhyme conditions. We expect to observe large effect sizes here. In each of the comparisons, effect size was calculated by dividing the difference between treatment means by a pooled estimate of the population standard deviation. The comparison of block-1 intrusions for the standard interference condition and the interference plus rhyme condition is complicated by the fact that on an immediate test in the standard interference conditions, very few block-1 intrusions occur. This makes parametric tests somewhat dubious. Since the strong test of our theory involves the frequency with which the foil is mistakenly produced as the target, we wanted to adopt a conservative test in the first instance. Thus we have first utilized the non-parametric Wilcoxon Signed Ranks Test as the most appropriate way to analyze these

differences. The statistic we report is a Z statistic where the sum of the signed ranks is divided by the square root of the sum of the signed ranks. We have supplemented the non-parametric test with an effect size estimate based upon a repeated measures t-test, knowing that there is some doubt as to the validity of this measure. It should be said that in all experiments the parametric and non-parametric tests suggested the same conclusion.

Part A

A summary of the different response types are presented in the top panel of Figure 2.

 Insert Figure 2 about here

As far as correct recall goes, the data in Figure 2 suggested that there was little difference between the control and standard interference condition. A small effect size of .08 confirmed this observation. In contrast to this, the difference between the control and interference plus rhyme conditions was more substantial with an effect size of .52. The differences in omission errors was not all that great across the three conditions. The effect size for the control and interference comparison was .22. For the control and interference plus rhyme conditions it was .19.

For the block-1 intrusions, 11 of the subjects produced more intrusions in the interference plus rhyme condition than the standard interference condition, 2 participants showed an advantage for the standard interference condition and 7 subjects produced tied scores. The Wilcoxon test on these data indicated that this pattern was highly reliable, $Z = 2.80$, $p < .005$. An effect size of .73 based on parametric techniques likewise confirmed that intrusions were more likely to occur in the interference plus rhyme condition.

Part B

The means for the different measures when the rhyme appears in the middle serial positions are presented in the bottom panel of Figure 2. Target recall appears to be equivalent in the control and standard interference conditions (effect size is .14), but PI effects appear to be present in the interference plus rhyme condition (effect size is .53). There are more omission errors in the control condition than in either the standard interference (effect size is .61) or the interference plus rhyme (effect size is .64) conditions. Block-1 intrusions are more frequent in the interference plus rhyme condition than in the standard condition with 15 subjects showing an interference plus rhyme advantage, three showing a standard interference advantage, and two producing tied scores. According to the Wilcoxon test, the difference in intrusions is reliable, $Z = 2.37$, $p < .05$. The effect size based upon a repeated measures t-test was .62.

Discussion

There are a number of important features of the current results. Firstly, both parts of the experiment produce the same pattern of effects. The position that the rhyme of the foil occupied did not appear to change performance in any way. Both parts of the experiment have replicated previous research indicating immunity to PI for sub-span lists of items from taxonomic categories (Halford et al, 1988; Humphreys & Tehan, 1992; Tehan & Humphreys, 1995; in press; Wickens et al., 1981) in that there is no reliable difference between performance on the standard interference and control trials. However, when a rhyme of the foil appears in the second block, the likelihood that the target will be recalled substantially diminishes and this is primarily due to the fact that the interfering foil is more likely to be

recalled instead of the target. Moreover, effect sizes are reasonably similar across the two parts of the experiment, at least for target recall and foil intrusions.

Immunity to PI in the absence of a rhyme but its emergence with a rhyme present is precisely the pattern of performance predicted from the Chappell and Humphreys (1994) model. In any event, the results are consistent with the idea that a rhyme can provide some of the phonemic information that has been lost from the representation of the foil making the foil a more potent competitor for the target.

The fact that a filler item that is unrelated to the cue can have an impact on target recall and foil intrusions, is a strong test of distributed representations. An even stronger test would involve the distribution of the phonemic components across several of the filler items. In the next experiment we set out to provide this strong test of the distributed storage assumption, by putting the rhyme phonemes of the foil into separate filler items in the second block. To use the example we have been using to this point, we took the rhyme component of the foil dog, og, decomposed it into its components, o and g, and then put the components in the same within-word position in other filler items, for example, put the o in mop and the g in fig. The expectation would be that we would produce the same results as in the first experiment.

Experiment 2.

Method

Subjects

A further twenty undergraduate students from the University of Southern Queensland participated in the experiment for course credit or a ticket in a cash lottery.

Materials and Procedure

The current experiment is identical to Experiment 1 in all respects save one. In the control and standard interference conditions the filler items were randomly

allocated to condition, block and position within block. In the interference plus components condition, two non random filler items appeared in the second block. The vowel phoneme of one of the fillers shared the vowel phoneme with the foil but initial and terminal phonemes differed. The second filler item shared the terminal phoneme with the foil but the initial and vowel phonemes differed. The filler with the shared vowel always appeared in the list before the filler with the shared terminal phoneme. That is, phoneme order was maintained.

Results and Discussion

The results of the current experiment are presented in Figure 3. The results can be summarized quite simply, there were no reliable differences across any of the measures. Thus, the effect sizes for target recall were .12 and .32 for the comparisons involving the standard interference condition and the interference plus components conditions respectively. With omissions as the dependent measure, the effects sizes for the standard interference and interference plus components conditions were .24 and .62, respectively. More importantly, for Block-1 intrusions only three subjects produced an interference plus component advantage, six produced a standard interference advantage, and there were 11 ties. The Wilcoxon test suggested that there was no reliable difference between the two means, $Z = -.79$, $p = .43$. An effect size of .22 confirmed the lack of intrusion differences.

Insert Figure 3 about here

Clearly, a rhyme is greater than the sum of its parts. In the first experiment a filler item that was a rhyme of the foil influenced the degree to which PI was observed in the experiment. In the present experiment when the rhyme was

decomposed and the parts distributed across other filler items no additional interference was created. Taken at face value the current results suggest that distributed storage has failed the strong test.

Before abandoning distributed storage completely, we decided to repeat Experiment 2 but this time ensure that all three phonemes of the foil were present rather than just the vowel and terminal phoneme. In the Chappell and Humphreys model the more components of a prelearned pattern that are active at the start of the deblurring process the more likely it is that the pattern will be produced. In Experiment 2, the lack of the initial phoneme might have proved crucial. Thus, in Experiment 3 in the interference plus components condition, the three components of the foil were spread across the three filler items: one of the fillers shared the same initial phoneme, one shared the same vowel sound, and the third shared the terminal phoneme as the foil. If the foil was dog then the fillers might have been dart, mop and fig. The materials used are present in the Appendix.

Experiment 3

Method

Subjects

Twenty undergraduate students from the University of Southern Queensland participated in the current experiment. Again course credit or a ticket in a small cash lottery were given for participation.

Materials and Procedure

The materials and procedures were identical to those used in Experiment 2, The only change was that in the interference plus components condition one of the filler items contained an initial phoneme that was shared with the foil. Thus all three filler items in the second block shared one phoneme with the foil. The order of the filler words maintained the order of the phonemes in the foil.

Results and Discussion

The results are summarized in Figure 4.

Insert Figure 4 about here

In looking at the target recall data, it would seem that there was no evidence of any PI. Effect sizes for correct target recall are consistent with this being 0 and .08 for the standard interference and interference plus components conditions, respectively. The error data tell a different story, however. The error patterns for the two interference conditions were very different to one another. For omissions, the effect size for the control/standard interference comparison was .15., but for the control/interference plus components comparison it was .87. For the block-1 intrusions, 15 students produced an interference plus components advantage, one produced a standard interference advantage and there were four tied scores. The Wilcoxon test carried out on this data confirmed that the difference in errors was highly significant, $Z = 3.33$, $p < .001$. Parametric analysis produced an effect size of .89.

While the correct recall data suggest that there is very little difference in recall between the three conditions, the error data suggest that the two interference conditions are very different to one another. In the standard interference condition subjects make relatively more omission errors than block-1 intrusions. The opposite is true of the interference plus components condition. Here it is much more likely for a block-1 intrusion to be made than an omission error.

Performance is becoming more complex in that there now exists an interaction amongst error types. However, it is clear that by putting all three phonemic components of the foil in the second block the likelihood of recalling the foil

increases dramatically. In other words, distributed storage assumptions pass the strong test.

Given that we did not find strong effects when just the rhyme components were utilized but we did when all three components, we thought it important to replicate the findings of Experiment 3. Furthermore, we were interested in determining if the order of the filler items was important. In the previous two experiments, the order of the filler items in block-two reflected the order of the phonemes in the foil, dart appeared as the first filler, mop appeared after dart and fig appeared after mop. In the next experiment we had four experimental conditions: the control and standard interference conditions and two interference plus components conditions. In the first of the interference plus components conditions the order of the filler items reflected the order of the phonemes in the foil. In the second condition the order of the filler items was not maintained, fig might have appeared first, dart after that and mop as the last filler item. Note that the relevant phoneme still maintains its position within the filler word. Our expectation was that the order of the filler items was not an essential part of the effect. Both interference plus components conditions should produce interference.

Experiment 4

Method

Subjects

Twenty undergraduate students from the University of Southern Queensland participated in the current experiment. Again course credit or a ticket in a small cash lottery were given for participation.

Materials and Procedure

The three interference conditions used in Experiment 3 were again used in the current experiment. In addition a fourth condition was also used. This latter condition

again involved the components of the foil being distributed across three filler items in the second block but now the order of the filler items was rearranged. In the case where the target item was in position three, the order was: filler with terminal phoneme in serial position one, filler with initial phoneme in position 2 and filler with vowel in position 4. When the target was in position 2 the order was vowel first, terminal third and initial fourth. As was the case in Experiment 3 the order of the phoneme within word still corresponded to that it had in the foil.

The use of an additional condition necessitated the selection of a further ten categories. These were again selected from the McEvoy and Nelson (1982) norms. To ensure that there was no overlap between critical items and fillers, several words had to be culled from the pool of filler items.

Again separate lists were created for each subject with materials being randomly assigned to conditions and the order of the lists being randomized as well.

Results and Discussion

The results of the experiment are presented in Figure 5.

Insert Figure 5 about here

In looking at the target recall data, it would seem that PI effects are present across all three interference conditions, although they appear to be more pronounced in the conditions where the components of the foil are present. Effect sizes for correct target recall reflect this. For the standard interference the effect size was .48, for the interference plus components in same order condition effect size was .65 and for the rearranged order it was .92. There were no large differences in omission errors with effect sizes of .18, .16 and .25 for the standard interference, components in same order and rearranged components conditions, respectively. In comparing block-1

intrusions for the standard interference and components in same order conditions, 7 students produced an interference plus components advantage, 2 produced a standard interference advantage and there were 11 tied scores. The Wilcoxon test carried out on these data indicated that the difference in errors was marginal, $Z = 1.73$, $p = .08$. Parametric analysis produced an effect size of .41. For the comparison between standard interference and rearranged components, 13 students produced a rearranged components advantage, 1 produced a standard interference advantage and there were 6 tied scores. The Wilcoxon test carried out on these data indicated that the difference in errors was reliable, $Z = 3.08$, $p = .002$. Parametric analysis produced an effect size of .83.

In spite of observing some effects of PI on the standard interference condition for the first time, the principle interests of the study have been realized. When the three phonemic components of an interfering foil are distributed over filler items in the second block subjects are much more likely to recall the foil than if there is only random overlap between fillers and foil. Furthermore, it looks as if the order in which the filler items that share features with the foil is not all that important. If anything, performance in the rearranged condition produces more interference than in the condition where order is maintained. We note that Li and Schweickert (1995) have shown that it is the presence of similar phonemes rather than their order (e.g. disk and skid) that is important for producing phonemic similarity effects. If order of phonemes within a word are not crucial for producing a similarity decrement, the current results are not all that controversial. In sum, in replicating the findings of Experiment 3, it would appear that we have established reasonable support for the role of distributed representations and storage as a mechanism that underlies short-term PI effects.

General Discussion

In the introduction to this paper we reviewed the relevant assumptions of memory models that employ distributed representations and distributed storage. One emergent feature of such models is that items can readily interact with each other. We have looked at the Chappell and Humphreys (1994) connectionist model with respect to how a target and foil might interact with each other in a short-term cued recall task. The current experiments were conceived to explore predictions concerning the presence or absence of immunity to PI on an immediate test, with the basic assumption that phonemic features for one item could be provided by other, unrelated list items. The basic prediction from the Chappell and Humphreys perspective was that if other items in the list could provide phonemic features, then PI would be observed. If the other items in the list were all unrelated to the foil then immunity to PI could be expected.

By and large empirical evidence across a range of short-term tasks employing short lists of taxonomically similar items confirms immunity to PI on an immediate test (Halford et al., 1986; Tehan & Humphreys, 1995; in press; Wickens et al., 1981). This pattern is reflected in the standard interference conditions in the current experiments. In all the experiments, bar the last, the presence of an interfering foil in the first block has had no noticeable effect upon target recall, nor do block-one intrusions serve as a major source of error. In the final experiment target recall is depressed but the effect is not strong and appears to be due to an increase in omission errors rather than increased recall of the foil.

There are boundary conditions to immunity to PI. In previous research we have demonstrated that PI can be observed on an immediate test if the foil and target rhyme (Tehan & Humphreys, 1995). In the current set of experiments we have shown that it is possible to produce PI when target and foil do not rhyme. When a rhyme of the foil is presented in the to-be-remembered second block, or all the phonemic

components of the foil are distributed across filler items then recall of the interfering foil is enhanced. We also have demonstrated that if only some of the components are distributed across fillers then PI is less likely to be observed.

These results are consistent with the assumptions of the Chappell and Humphreys (1994) connectionist model and the Tehan and Fallon (in press) adaptation of it. We have presented our explanations for the presence or absence of PI in the introductions to the various experiments and do not wish to reiterate them here. Instead we wish to consider other possible explanations for the effects, something we have not addressed to this point.

The first issue involves possible ceiling effects with target recall and floor effects with errors. Target recall across all conditions across the four experiments surpasses the 90% accuracy level only once. The task is sufficiently difficult that most subjects make some type of error. With respect to errors, particularly block-1 intrusions, it is not the case that a few subjects were producing all the errors. In fact, the vast majority of subjects (70% - 85% across experiments) made at least one block-1 intrusion. Furthermore, ceiling and floor effects only address the issue of how much is recalled, rather than providing an explanation for accuracy of recall (Koriat & Goldsmith, 1996). We think that we have provided an explanation for why recall is quite accurate under control and standard interference conditions but less accurate under conditions where fillers bear some similarity to the foil.

There are some strategies that subjects might have adopted that could produce the pattern of results that we observe. With regards to the first experiment, the results might simply be an artefact of subjects noticing that a rhyme was present which in turn would remind them of the foil, thereby strengthening its representation. (Tzeng & Cotton, 1980). Although we believe that the manipulation of the rhyming filler item was transparent to subjects, we did not do any post-experiment checks to see if

this was so. Even if the manipulation was transparent, a model of short-term memory that was based upon spreading activation might predict such a result. With the presentation of the rhyme, activation might spread through a rhyme network such that the foil in the interference plus rhyme condition would receive additional activation from the rhyme and thereby become a stronger competitor than it would be in the standard interference condition. These explanations of spreading activation and noticing the rhyme of the foil was present should be limited to the first experiment. Distributing the phonemes across fillers should severely limit the spread of activation to the foil and thereby reducing the chances of observing PI^1 . Distributing the phonemes across fillers should also make it impossible for subject to notice the relationship between block-2 fillers and the foil.

The final strategy that might produce PI effects with the rhyme is a generate recognize access procedure. In the conventional versions of the generate-recognize procedure (Humphreys, Bain & Pike, 1989; Jacoby & Hollingshead, 1990), the subject would start to generate instances of the category and then attempt to match the output of the generate procedure with episodic traces of the list items. As each item is generated, it is matched to each item in memory and if a match exceeds threshold then that item is produced as the response. In long-term recognition studies it has been documented that distractors (new items) that are rhymes of targets are more likely to be falsely recognized than non-rhyming distractors (Byrne & Shea, 1979; Runquist & Blackmore, 1973), although this effect is not large. If we assume that phonemic information in the short-term domain has a privileged role, then such false recognition might occur more frequently in the short-term domain. Given that the foil is usually a high dominant instance of the category, it is likely to be generated. Falsely recognizing the rhyme as the foil may well happen on the odd

occasion. Where there is no rhyme of the foil in the list it is unlikely that matching failures are going to occur.

We believe that the results of the last two experiments lessen the plausibility of the above explanations. The problem posed for generate-recognition by these experiments is the very weak relationship involved. That is, each filler item has only a single phoneme in common with the generated item so that the contribution of a match between the generated item and one of these fillers would be very small. Instead, we would assert that these results support two things: representation at the level of features and the simultaneous activation of all the list items at the time of retrieval. We want to address these assumptions in some detail.

The idea that representations can be conceived of as sets of features is not new (Bower, 1967; Underwood, 1969). Models that have taken a feature approach (Drewnowski, 1980; Nairne, 1990) have tended to stress the phonemic features of the stimuli involved and down play or ignore semantic features. Empirical support for phonemic features is readily observed in studies that have examined errors on immediate serial recall. Intrusion errors tend to share the unforgotten phonemic characteristics of the target item (Wickelgren, 1965a, b, c.), for instance, syllabic stress pattern, stressed vowels and the beginning and terminal phonemes (Drewnowski and Murdock, 1980). These results imply that partial forgetting of an item in short-term memory is possible, and that items in short-term memory are not stored in a unitary fashion but rather as an aggregate of smaller features.

Empirical support for the contribution of non-phonological features is less pronounced. However, the Chappell and Humphreys deblurring procedure asserts that all features contribute to identifying the target item. Nairne's Feature model makes a similar assertion. At the empirical level, word frequency effects in span (Tehan & Humphreys, 1988; Watkins, 1977) could well stem from semantic features facilitating

the deblurring process. Poirier and Saint-Aubin's (1995) demonstration of the facilitative effects of taxonomic similarity on span is a more obvious example of semantic features impacting upon immediate recall. In any event, we think that a systematic examination of semantic and episodic long-term effects in serial recall is long over due.

The more contentious issue is that the results of the latter experiments appear to require assumptions about the simultaneous activation of all list items at retrieval. That is, it would appear that at the time of recall the target, the foil and the three filler items are all participating in the recall process. Such an assumption is not unknown in the literature. In fact simultaneous activation at retrieval is a central tenet of Hintzman's (1986) Minerva II model, which employs local storage, and the distributed storage models such as TODAM (Murdock, 1982) and the Matrix Model (Pike, 1984).

As mentioned earlier, the current results present serious problems for conventional versions of the generate-recognize approach. More sophisticated generate-recognize procedures might be possible, however. For example, in SAM (Gillund & Shiffrin, 1984) all generated items are matched to all the items in memory before a response is made. That is, each match produces some matching strength and the strengths of all the matches are then summed into a single value. The recognition decision is based upon the strength of the composite value. The assumption here is that if a list item is generated, the match to the composite value is going to be stronger than if a list item is not generated. In the case of the phonemic characteristics being distributed across filler items, it is possible that when the foil is generated and matched to the filler items, weak matching strengths are produced on each occasion. However, when the weak strengths are summed prior to decision, it is possible that

the combined matching strength might be above threshold level. If so, the foil would be produced as the response.

Although we cannot rule out a generate-recognize process of the type envisioned in SAM, we have indicated a number of times that our preferred alternative is for distributed storage and direct access of the type that is utilized by Chappell and Humphreys (1994). The current results are certainly consistent with the assumptions of that model. As a number of authors have commented, distributed storage does appear to have several advantages. Retrieval dynamics are very much simplified in that access is direct via cues with no search mechanism needing to be specified, partial losses of memory need not have a devastating effect upon performance, and the models lend themselves to associative effects (Lewandowsky & Murdock, 1989; Murdock, 1982). We would also argue that interactions among list items of the type seen in the current experiments are best understood as being an emergent feature of distributed storage.

Conclusions

We started this paper by observing that there was considerable diversity in assumptions concerning representation, storage, forgetting and memory access in current models of immediate recall. The current results, we believe, suggest a number of constraints. The results strongly support feature based representations in which phonemic features have a privileged role at very short retention intervals. The item interaction effects of the type that characterize the current findings suggest that specific items have to some extent lost their identity. Our preferred explanation for this is to argue that these item interaction effects are an emergent consequence of distributed storage. However, we cannot rule out the idea that local storage with a global matching retrieval process is involved. Given that very few current models of immediate recall involve distributed representation with distributed storage, let alone

consider PI effects, the current results should provide a challenge for existing models and place substantial constraints upon the development of future models.

References

Baddeley, A. D. (1966). Short-term memory for word sequences as a function of acoustic, semantic and formal similarity. Quarterly Journal of Experimental Psychology, 18A, 362-365.

Bourassa, D. C., & Besner, D. (1994). A semantic contribution to serial order recall of subspan lists. Psychonomic Bulletin and Review, 1, 122-125.

Conrad, R. (1965). Order errors in immediate recall of sequences. Journal of Verbal Learning and Verbal Behavior, 4, 161-169.

Bower, G. H. (1967). A multicomponent theory of the memory trace. In K. W. Spence and J. T. Spence (eds.) The Psychology of Learning and Motivation, (Vol. 1. pp. 230-325). New York: Academic Press.

Brown, G. D. A., & Hulme, C. (1995). Modelling item length effects in memory span: No rehearsal needed? Journal of Memory and Language, 34, 594-621.

Byrne, B., & Shea, P. (1979). Semantic and phonemic codes in beginning readers. Memory and Cognition, 7, 333-338.

Chappell, M., & Humphreys, M. S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. Psychological Review, 101, 103-128.

Crowder, R. G. (1976). Principles of Learning and Memory. Hillsdale, N. J.: Erlbaum.

Drewnowski, A. (1980). Attributes and priorities in short-term recall: A new model of memory span. Journal of Experimental Psychology: General, 109, 208-250.

Drewnowski, A., & Murdock, B. B. (1980). The role of auditory features in memory span for words. Journal of Experimental Psychology: Human Learning and Memory, 6, 319-332.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. Psychological Review, 91, 1-67.

Halford, G. S., Maybery, M. T., & Bain, J. D. (1988) Set-size effects in primary memory: An age-related capacity limitation? Memory and Cognition, 16, 480-487.

Henson, R. N. A., Norris, D. G., Page, M. P. A., & Baddeley, A. D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. Quarterly Journal of Experimental Psychology, 49A, 80-115.

Hintzman, D. L. (1986) "Schema abstraction" in a multiple-trace memory model. Psychological Review, 93, 411-428.

Humphreys, M. S., & Tehan, G. (1992). A simultaneous examination of recency and cuing effects. In A. Healy, S. Kosslyn & R. M. Shiffrin (Eds.) From Learning Theory to Cognitive Processes: Essays in Honor of William K. Estes. (Vol. 2., pp. 143-159). Hillsdale N. J.: Erlbaum and Associates.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic and procedural tasks. Psychological Review, 96, 208-233.

Jacoby, J. J., & Hollingshead, A. (1990). Towards a generate/recognize model of performance on direct and indirect tests of memory. Journal of Memory and Language, 29, 433-454.

Koriat, A., & Goldsmith, M. (1996). Memory metaphors and the real-life/laboratory controversy: Correspondence versus storehouse conceptions of memory. Behavioral and Brain Sciences, 19, 167-228.

Lewandowsky, S., & Li, S.-C. (1994). Memory for serial order revisited. Psychological Review, 101, 539-543.

- Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. Psychological Review, 96, 25-57.
- Li, X., & Schweickert, R. (1995). Phonological similarity effect on span: Variable positions of repeated phonemes. Poster presented at 36th Annual Meeting of The Psychonomic Society, Los Angeles, November 10-12.
- Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. Psychonomic Bulletin and Review, 1, 476-490.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. Psychological Review, 88, 375-407.
- McEvoy, C. L., & Nelson, D. L. (1982). Category name and instance norms for 106 categories of various sizes. American Journal of Psychology, 95, 581-634.
- Murdock, B. B., Jr. (1982). A theory for the storage and retrieval of item and associative information. Psychological Review, 89, 609-626.
- Murdock, B. B., Jr. (1995). Developing TODAM: Three models for serial-order information. Memory and Cognition, 23, 631-645.
- Nairne, J. S. (1988). A framework for interpreting recency effects in immediate serial recall. Memory and Cognition, 16, 343-352.
- Nairne, J. S. (1990). A feature model of immediate memory. Memory and Cognition, 18, 251-269.
- Neath, I., & Nairne, J. S. (1995). Word-length effects in immediate memory: Overwriting trace decay theory. Psychonomic Bulletin and Review, 2, 429-441.
- Pike, R. (1984). A comparison of convolution and matrix distributed memory systems. Psychological Review, 91, 281-294.

Poirier, M., & Saint-Aubin, J. (1995). Memory for related and unrelated words: Further evidence on the influence of semantic factors in immediate serial recall. Quarterly Journal of Experimental Psychology, 48A, 384-404.

Runquist, W. N., & Blackmore, M. (1973). Phonemic storage of concrete and abstract words with auditory presentation. Canadian Journal of Psychology, 27, 456-463.

Schneider, W., & Detweiler, M. (1988). A connectionist/control architecture for working memory. In G. Bower (Ed.) The Psychology of Learning and Motivation, (Vol 21, pp 54-119.) New York: Academic Press.

Schweickert, R. (1993). A multinomial processing tree model for degradation and redintegration in immediate recall. Memory and Cognition, 21, 168-175.

Shapiro, S. I., & Palermo, D. S. (1970). Conceptual organization and class membership: Normative data for representatives of 100 categories. Psychonomic Monograph Supplements, Vol. 3, No.11 (Whole of No. 43).

Tehan, G., & Fallon, A. B. (in press). A connectionist model of short-term recall. Proceedings of CogSci'95.

Tehan, G., & Humphreys, M. S. (1988). Articulatory loop explanations of memory span and pronunciation rate correspondences: A cautionary note. Bulletin of the Psychonomic Society, 26, 293-296.

Tehan, G., & Humphreys, M. S. (1995). Transient phonemic codes and immunity to proactive interference. Memory and Cognition, 23, 181-191.

Tehan, G., & Humphreys, M. S. (in press). Cuing effects in short-term recall. Memory and Cognition.

Tzeng, O. J. L., & Cotton, B. (1980). A study-phase retrieval model of temporal coding. Journal of Experimental Psychology: Human Learning and Memory, 6, 705-716.

Underwood, B. J. (1969). Attributes of memory. Psychological Review, 76, 559-573.

Walling, J. R., McEvoy, C. L., Oth, J. E., & Nelson, D. L. The University of South Florida rhyme category norms. Unpublished manuscript, University of South Florida.

Wickelgren, W. A. (1965a). Acoustic similarity and retroactive interference in short-term memory. Journal of Verbal Learning and Verbal Behavior, 4, 53-61.

Wickelgren, W. A. (1965b). Acoustic similarity and intrusion errors in short-term memory. Journal of Experimental Psychology, 70, 102-108.

Wickelgren, W. A. (1965c). Similarity and intrusions in short-term memory for consonant-vowel digraphs. Quarterly Journal of Experimental Psychology, 17, 241-146.

Watkins, M. J. (1977). The intricacy of memory span. Memory and Cognition, 5, 529-534.

Wickens, D. D., Moody, M. J., & Dow, R. (1981). The nature of timing of the retrieval process and of interference effects. Journal of Experimental Psychology: General, 110, 1-20.

Author Notes

This research was supported by Australian Research Council grant No. A79531860. Reprint requests should be addressed to Gerry Tehan, Department of Psychology, University of Southern Queensland, Toowoomba. 4350. Australia. e-mail tehan@usq.edu.au

Footnotes

1. In spreading activation models of word identification (e.g. McClelland & Rumelhart, 1981) there is often assumed to be an interaction between phonemes at the letter level representation and at the word level representation. However, this interaction is seen to involve the linkage between peripheral and central codes. In episodic memory, we know of no evidence indicating that items that share a single phoneme are demonstrably more similar than words that do not have any phonemes in common.

Figure Captions.

Figure 1. A sample trial sequence illustrating proactive interference.

Figure 2. Performance measures as a function of interference condition in Experiment 1.

Figure 3. Performance measures as a function of interference condition in Experiment 2.

Figure 4. Performance measures as a function of interference condition in Experiment 3.

Figure 5. Performance measures as a function of interference condition in Experiment 4.

Fig 1.

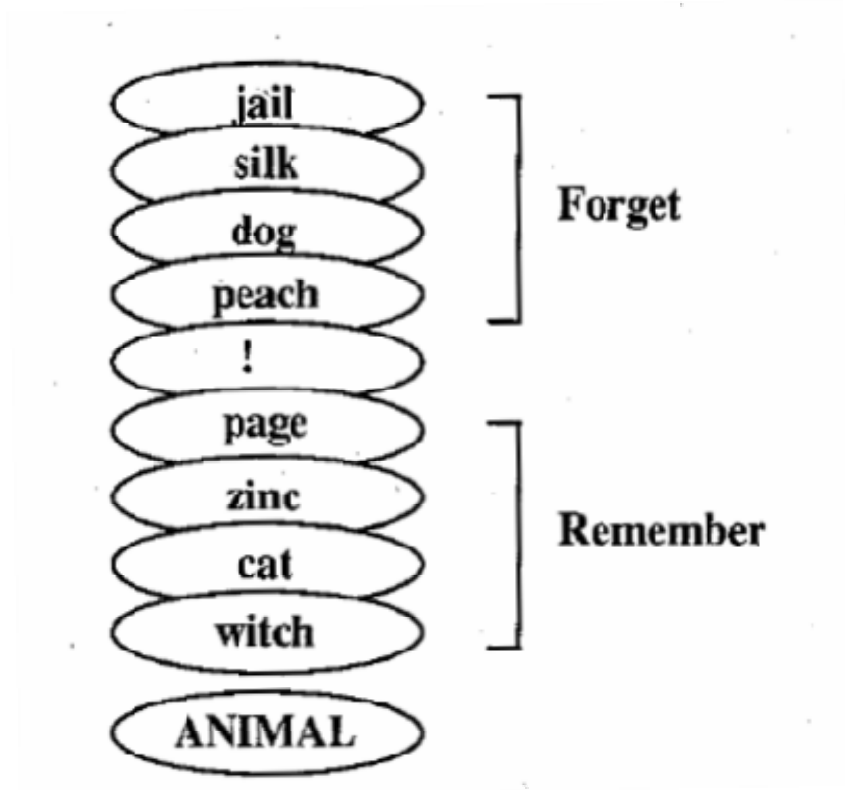


Fig 2.

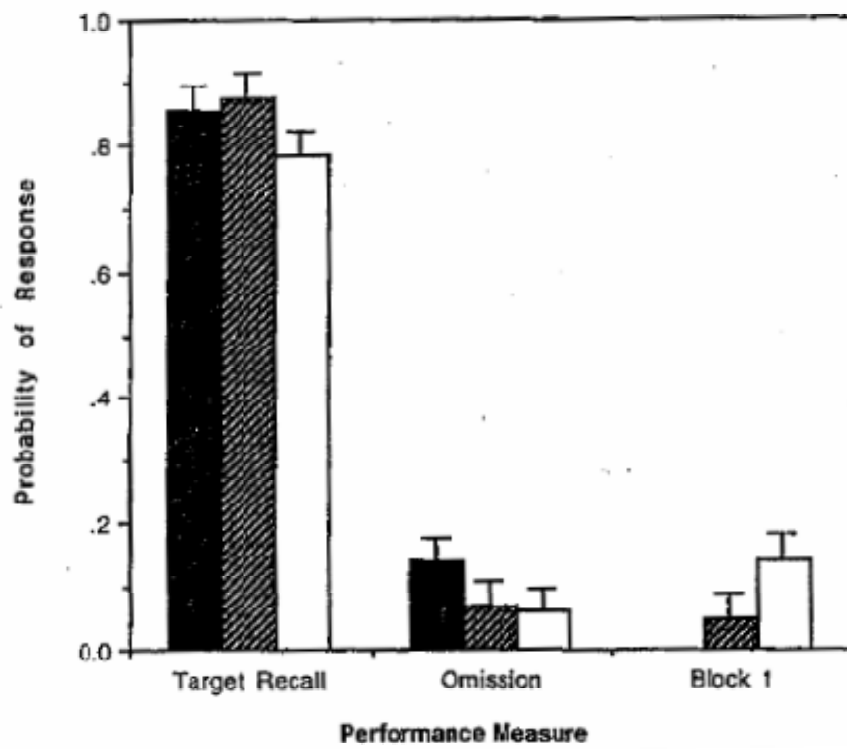
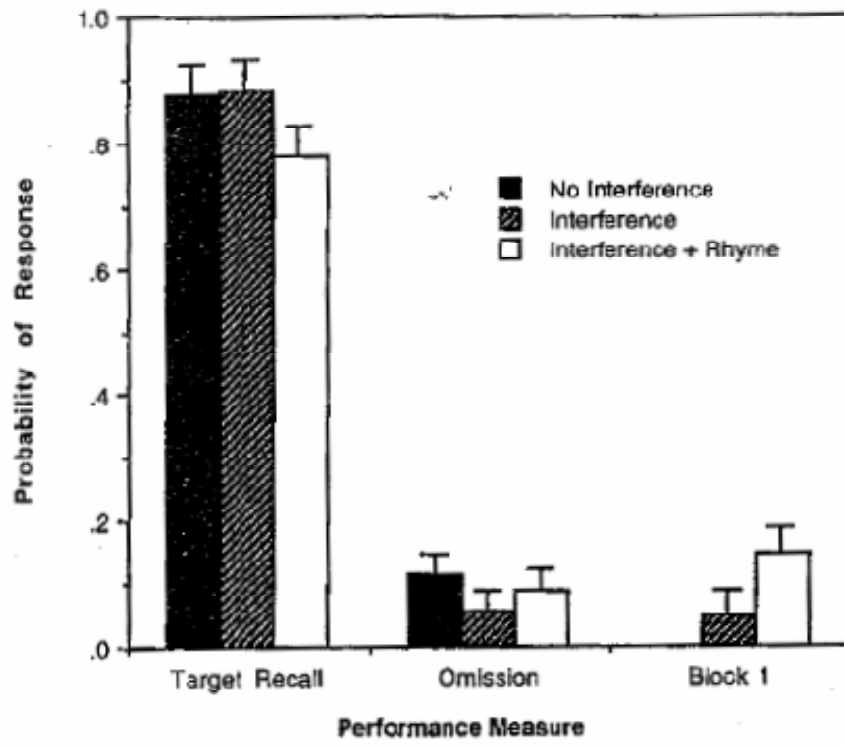


Fig 3.

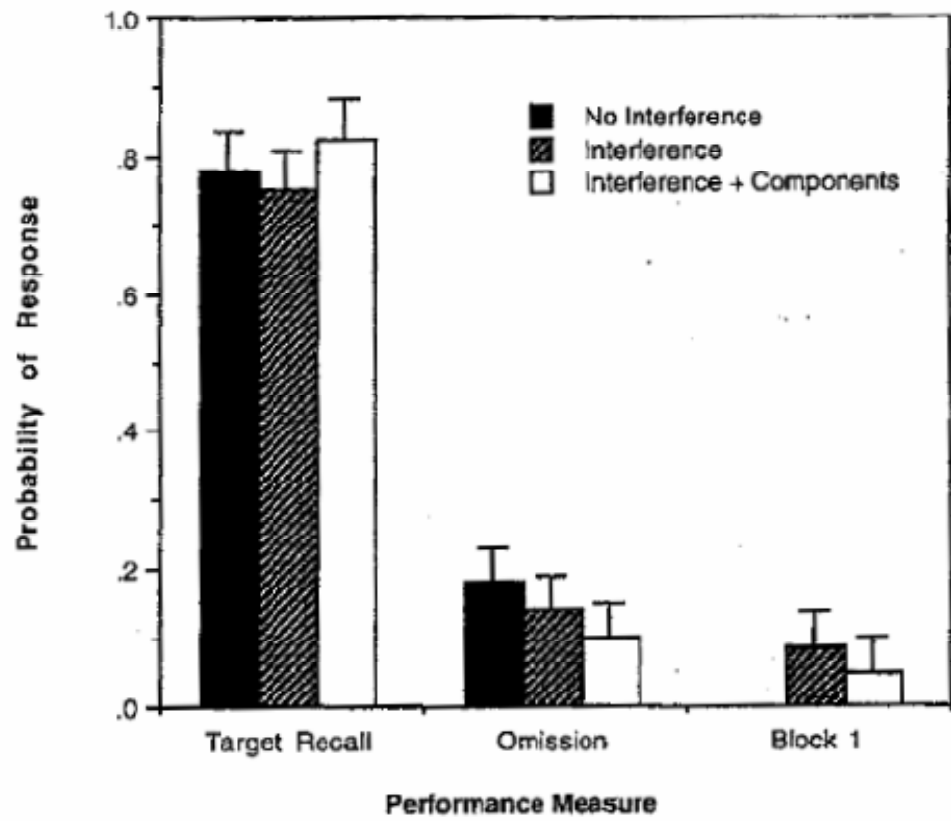


Fig 4.


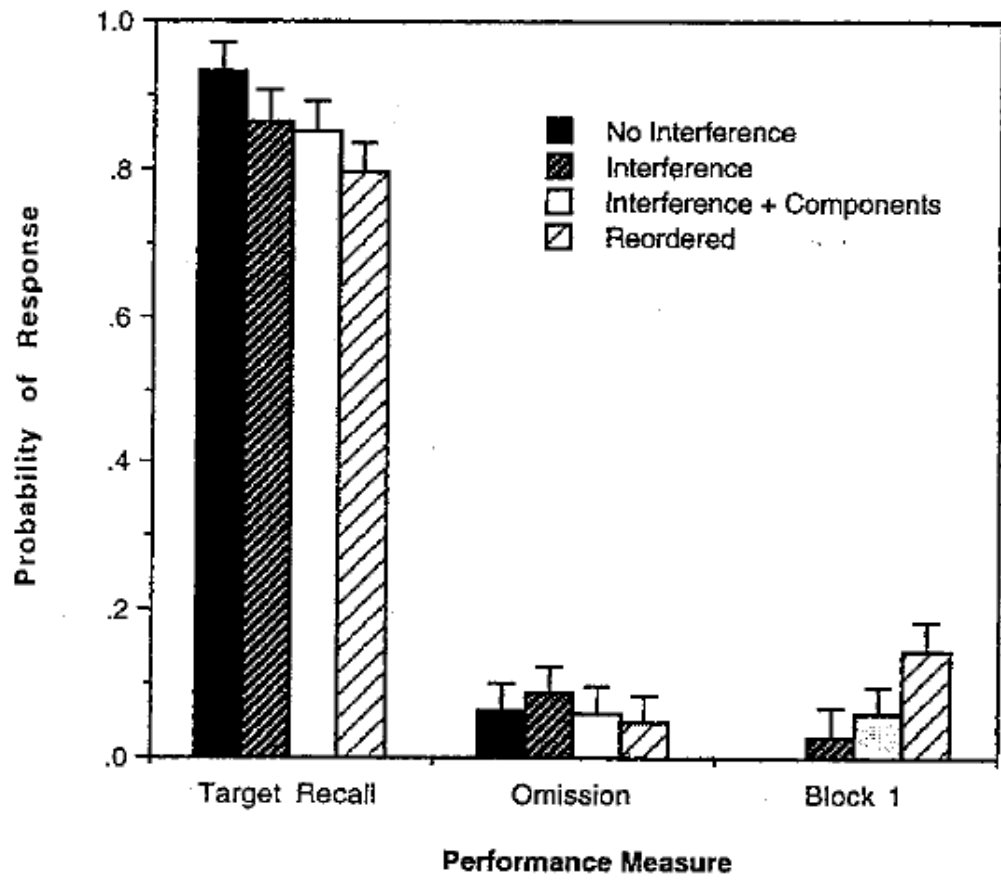


Figure 4. Probability of Response

Fig 5,



Appendix

Cue	Foil	Target	Initial	Vowel	Terminal
WATER-BIRD	gull	swan	ghost	bun	hell
TYPE-OF-WATERWAY	stream	lake	strike	teak	clam
GARDEN-TOOL	rake	hose	room	day	muck
TYPE-OF-MEAT	ham	roast	hole	map	rum
TYPE-OF-VEHICLE	car	wagon	cup	heart	boar
MUSICAL-INSTRUMENT	harp	mandolin	hill	arm	top
TYPE-OF-RODENT	bat	guinea-pig	bill	pack	net
TYPE-OF-FRUIT	peach	mango	page	leap	witch
PART-OF-A-BEDROOM	bed	wardrobe	bug	step	lid
FARM-ANIMAL	hen	turkey	hull	bell	fan
EATING-UTENSIL	knife	chopstick	nut	mine	stiff
MEMBER-OF-ROYALTY	king	duchess	key	silt	slang
AMERICAN-COIN	dime	quarter	dagger	tie	loam
PART-OF-A-TREE	trunk	stem	tram	hum	pink
BIRD-OF-PREY	hawk	falcon	hand	door	tick
PART-OF-A-BOAT	deck	stern	dirt	hen	pack
KIND-OF-LINGERIE	slip	nightie	slot	pit	flap
BUILDING-MATERIAL	brick	studs	broom	whistle	pack
ITEM-OF-CAMPING-EQUIPMENT		tent	cooler	tyre	fender
	paint				
GRAMMATICAL-PART-OF-SPEECH		noun	conjunction		noose
	tower	bun			
ARTICLE-OF-CLOTHING	shirt	cap	shore	burr	pet
COLOUR	white	violet	waist	vine	fit
TYPE-OF-CLOTH	silk	rayon	sum	river	talk
MEANS-OF-COMMUNICATION	phone	letter	fume	moat	can
DAILY-MEAL	lunch	supper	law	hut	pinch
BREED-OF-DOG	hound	terrier	hack	louse	mind

EMOTION	love	shock	lift	gum	hive
ITEM-OF-OFFICE-EQUIPMENT	desk	computer	dart	hem	musk
PART-OF-A-FACE	nose	chin	newt	vote	haze
TYPE-OF-FLOWER	rose	marigold	rubber	note	maze
TYPE-OF-FOOTWEAR	sock	clog	salad	god	pick
GEOMETRIC-SHAPE	cube	diamond	cat	feud	pub
BATHROOM-FIXTURE	tap	shower	toll	hag	pup
TYPE-OF-BREAD	herb	rye	hill	shirt	tub
CAMPING-EQUIPMENT	stove	pan	step	coat	hive
CHEMICAL-ELEMENT	zinc	nickel	zone	pit	honk
COMPASS-DIRECTION	west	south	wall	pen	vast
PART-OF-THE-FACE	cheek	lips	chop	peep	make
GYMNASTIC-EVENT	rings	vault	rat	mint	thugs
HORSE-RIDING-EQUIPMENT	rein	whip	roll	stay	sun

Note: The first thirty sets were used in Experiments 2 (only vowel and terminal consonant) and 3. The bottom ten sets were added for Experiment 4.