UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Computational Explorations of the Irrelevant Sound Effect in Serial Short-Term Memory

Permalink

https://escholarship.org/uc/item/0pq2k1dv

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 22(22)

Author

Beaman, C. Philip

Publication Date

2000

Peer reviewed

Computational Explorations of the Irrelevant Sound Effect in Serial Short-Term Memory.

C. Philip Beaman (c.p.beaman@reading.ac.uk)
Department of Psychology
University of Reading
Earley Gate, Whiteknights
Reading RG6 6AL

Abstract

Although a number of current models of immediate serial recall exist, only one published model (Neath, 1999, 2000) incorporates simulations of the disruption of immediate serial recall by irrelevant background sound. This paper explores a possible model of irrelevant sound effects derived from Neath (1999) and applies the results of the model to previously unconsidered data sets. Studies by Neath (1999, 2000) apply the feature model, a mathematical model of short-term memory (Nairne, 1990), to some basic data regarding the irrelevant sound effect but this approach is ultimately limited by implicit assumptions regarding the nature of interference in short-term memory. Relaxing these assumptions allows for a wider application of a model of the irrelevant sound effect derived from that of Neath but not tied to the implementational detail of the feature model. The new model fits not only the original data considered by Neath (1999, 2000) but also empirical results concerning the effects of word-dose (Bridges & Jones, 1996) and token set size (Tremblay & Jones, 1998). It is concluded that the principles underlying the model provide a promising basis for further theoretical work.

Introduction

Theories of verbal short-term memory have frequently made use of the irrelevant sound effect, the disruption to serial recall of visually-presented verbal lists by background noise, to inform the proposed architecture of short-term memory (e.g., Salamé & Baddeley, 1982; Jones & Macken, 1993). Briefly, sound played to a participant during or immediately following the visual presentation of a to-be-recalled list impairs recall of the list even though the participant was explicitly told to ignore anything they might hear and participants are never tested on the contents of the "irrelevant" or "unattended" sound. It is well established that, although there are individual differences in the level of susceptibility to irrelevant sound disruption (Ellermeier & Zimmer, 1997), most participants show the effect.

There has been a paucity of formal simulations of irrelevant sound disruption, even though it has been claimed that a number of models of immediate serial recall can, in principle, account for the effect (Burgess & Hitch, 1999; Norris, Page & Baddeley, 1995). One model that has been applied to the effect is the feature model of Nairne (Nairne, 1990; Neath & Nairne, 1995), a mathematical model of

short-term memory based around the idea that the items most likely to be recalled from a list are those items which are most distinctive (Nairne, 1988). The model has been applied to a number of short-term memory phenomena including modality differences, interference from concurrent articulation and from post-list stimulus suffixes (Nairne, 1990), the word-length effect (Neath & Nairne, 1995), and latterly the irrelevant sound effect (Neath, 1999, 2000).

The Feature Model

The feature model assumes that interference rather than decay accounts for loss from short-term or *primary* memory. Representations of items in the feature model are vectors that code for the "features" of an item using a binary system allowing features to assume the values of +1 or –1. Features may be modality dependent, coding information available only in a specific sensory modality, or modality independent, coding information that can be conveyed equally by two or more modalities. Interference occurs in the model through overwriting. If a feature takes the same value as its counterpart in the immediately preceding vector, the earlier feature value is overwritten. This is implemented by setting the value of the feature to 0 so that it is informationally uninformative. To give an example, if feature *x* of item *n+ 1* is the same as feature *x* of item *n*, then feature *x* of item *n* is overwritten and can play no part in determining whether or not item *n* is accurately recalled (see Figure 1). Feature values are generated randomly and independently for each feature vector.
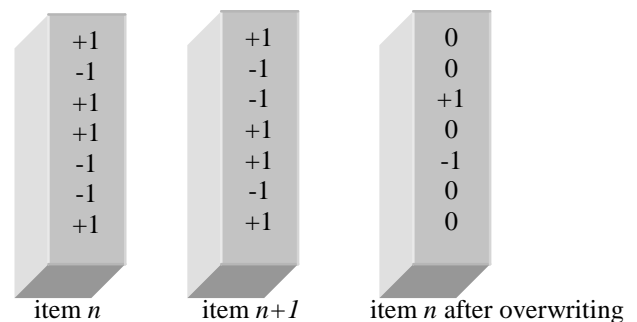

item *n*	item *n+1*	item *n* after overwriting
+1	+1	0
-1	-1	0
+1	-1	+1
+1	+1	0
-1	+1	-1
-1	-1	0
+1	+1	0

Figure 1. Degradation of the representation of a list item in primary memory when a successive item shares some of the same feature values.

Retrieval consists of finding the best match to a degraded cue amongst a set of undegraded feature vectors assumed to reside in what is termed *secondary memory*. Two memory systems coding the same information is in many ways an unsatisfactory situation if retrieval depends on the degraded representation. Nevertheless, it is useful from the point of view of modelling the irrelevant sound effect since accurate recall of the degraded memory trace can be viewed as recall of the correct item in the correct order. In immediate serial recall the to-be-recalled stimuli are typically overlearned, lists of digits or consonants for example, so the task is essentially one of identifying which (known) item appeared in which serial position. As argued elsewhere (Beaman & Jones, 1997, 1998) the irrelevant sound effect consists primarily of a disruption of order information. The distance between the degraded item and its undegraded secondary memory representation is calculated by summing the number of mismatched features, M, and dividing by the total number of compared features, N, as described in Equation 1.

$$d_{ij} = a \sum \frac{b_k M_k}{N} \qquad (1)$$

The value M_k is the number of times feature position x_{ik} does not equal feature position x_{jk}. The parameter a is a scaling parameter that is assumed to correspond to the overall level of attention, and b_k is used to weight particular comparisons if the task makes them more important than other comparisons. Distance, d, is then used to calculate the similarity between the degraded vector and the undegraded secondary memory representation according to Equation 2.

$$s(i, j) = e^{-d_{ij}} \qquad (2)$$

The probability that a particular secondary memory trace, SM_j will be sampled as a potential recall response for a particular degraded memory vector PM_i is then given by Equation 3, where w_{ij} and w_{ik} are possible response bias weights.

$$P_s(SM_j \mid PM_i) = \frac{w_{ij} s(i, j)}{\sum_{k=1}^{N} w_{ik} s(i,k)} \qquad (3)$$

This basic overwriting model was supplemented by Neath (1999, 2000) with two additional assumptions to account for the irrelevant sound effect[1]. The first assumption was that

---

[1] In fact, the full version of the feature model also includes a further recovery equation that produces the characteristic bow-shaped serial position curve. However here we are specifically concerned with the results of overwriting. Since there has been never been any suggestion that interactions between irrelevant sound and serial position might be of theoretical significance the recovery equation has been omitted here and performance averaged over serial position, a procedure also followed by Neath (Neath, 1999, 2000).

irrelevant sound will act like a concurrent articulation task, already accounted for by the model (Nairne, 1990) and overwrite a certain proportion of the modality independent features. The second assumption was that irrelevant sound and concurrent articulation manipulations differ in that effort is required to actively produce irrelevant noises in the concurrent articulation manipulation, which is not true of the irrelevant sound manipulation. Neath therefore proposed varying the attentional parameter a by a greater amount in simulations of concurrent articulation than in simulations of irrelevant sound.

With these amendments the feature model shows the correct qualitative pattern of results across a number of experimental studies altering only those parameters associated with the particular psychological process implicated (Neath, 1999, 2000). For example, the model shows correctly that the irrelevant sound manipulation impairs memory for lists of words, but less so than concurrent articulation. However, as with all simulation studies, there is experimental evidence not addressed by the model. Some of this evidence is directly relevant to the way in which irrelevant sound interferes with memory representations and cannot be accounted for by the feature model as it is currently formulated.

Two inconsistencies exist in the feature model account. Firstly, as described earlier, within-list interference results only in a lack of information about the overwritten item, not misinformation. Equation 1 ensures that these effects will be functionally identical, since only mismatches between the degraded vector and the undegraded secondary memory representation influence the similarity calculation (Equation 2) and both lack of information and misinformation are counted as mismatches. Nevertheless, the theory would be more parsimonious if all overwriting was implemented in an identical manner. The second inconsistency is more serious and concerns the difference between overwriting by irrelevant sound and overwriting by concurrent articulation. Concurrent articulation is implemented as setting half of the modality independent features to a constant value because participants are required to repeated the same utterance over and over "so the same information will overwrite the to-be-remembered items" (Neath, 2000). However in a simulation showing how varied speech (referred to in the literature as "changing-state" irrelevant sound) impairs recall performance more than repeated speech ("steady-state" irrelevant sound) (Jones, Madden & Miles, 1992) this logic was not used. Instead a variation in the attentional parameter is invoked, with variable speech assumed to attract more attentional resources.

The alteration in the attention parameter a is necessary as demonstrated by Figure 2, which shows the average sampling probabilities of a 9-item list in steady state, changing state and quiet control conditions across 200 simulations. The steady state condition comprised of setting half of the modality independent features to a constant value as described in previous simulation studies (Nairne, 1990). The changing-state condition comprised of overwriting half the modality dependent features with different random combinations of +1 and −1. The attentional parameter, a, was set to an identical value for all conditions. All other

weights were set to 1.0. Note that, provided all the other parameters remain unaltered, the same patterns of performance can be obtained at different overall recall levels by simple manipulation of the attentional parameter, *a*. However, this would simply be an exercise in data-fitting and not of psychological interest. The important point to note is that without the adjustment of the attentional parameter no changing-state effect is observed. Parameter adjustment of this type is also perilously close to data-fitting.
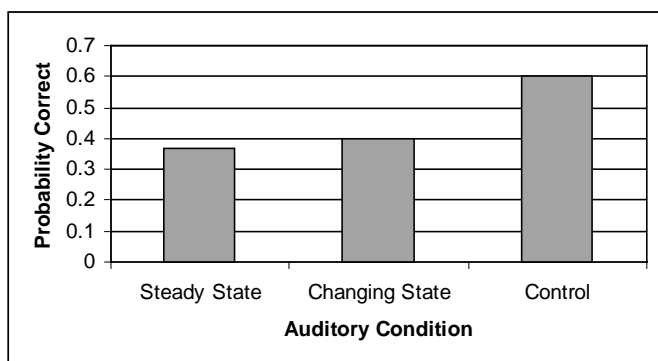

Figure 2. Sampling probabilities of items in the feature model under steady state and changing state irrelevant sound conditions when the attentional parameter is kept constant.

A Revised Model: The Changing-State & Token Set Size Effects

The problem of the changing-state effect can be by viewing it as a special case of what Tremblay and Jones (1998) termed "token set size". These authors argued that the essential cause of disruption by irrelevant sound was the presence of change in the irrelevant speech stream (Tremblay & Jones, 1998). The number of different changes, they argued, was irrelevant: disruption should markedly increase from one token (steady-state) to two tokens (changing-state) and there should be little or no further disruption beyond this token set size.

To give a concrete example, repetition of the utterance "A" in the irrelevant sound stream constitutes steady-state irrelevant sound and a token set size of one. According to Tremblay and Jones this should not cause discernible disruption to immediate recall. Repetition of the utterance "A-B" however, has a token set size of two and introduces change into the irrelevant sound stream and should therefore disrupt immediate recall. Repetition of the utterance "A-B-C" is also a changing-state stimulus (with a token set size of three) and should therefore also disrupt recall, but not necessarily to any greater degree than a token-set size of two since it is the number of changes, not the nature of the changes, which is important. Thus, changes from A to B then back to A are functionally equivalent to changes from A to B to C. Jones and Tremblay (Jones & Tremblay, 2000) argued that the increment in the attention parameter necessary to account for the changing-state effect did not

have a principled basis. If the increase in the attention parameter was necessary to account for increased attentional demands of changing-state stimuli, they argued, it should be increased in linearly as token set size increases, which would result in a linear increase in disruption not present in the experimental data.

A more realistic simulation than that attempted by Neath (Neath, 1999, 2000), and one that is not subject to these criticisms can be attempted by dropping the overwriting inconsistencies within the feature model. Closer examination of the experimental procedure employed in the Tremblay and Jones study reveals that over a 19 second presentation and retention interval Tremblay and Jones presented large numbers of repetitions of the same tokens. 38 separate occurrences of the same token in set size 1 condition, 19 repetitions each of 2 tokens in set size 2 condition, 13 repetitions of 3 tokens in set size 3, 8 repetitions of 5 tokens in set size 5 and 5 repetitions of 7 tokens in set size 7. If a conservative estimate of a covert rehearsal rate no faster than the slowest overt rehearsal rate of 2 items per second is assumed there could have been 38 rehearsals of a single item in this time period. The feature model has to assume that interference with the representations can occur at rehearsal as well as encoding since the experimental data demonstrate that the irrelevant sound effect can occur in an unfilled retention interval, after list presentation but before recall (Beaman & Jones, 1998). Therefore there will have been multiple opportunities for interference in this time period and the feature model's assumptions that changing-state irrelevant sound randomly overwrites half of each item's feature values once begins to look implausible.

Instead assume that each item was rehearsed once as it is presented- this is a standard assumption common to many models of immediate serial recall (e.g., Page & Norris, 1998). This leaves a 10 second retention interval which, with a slow rehearsal rate of 2 items per second and a 9 item list to rehearse gives time for only 2 complete rehearsals of the entire list. Thus altogether there is sufficient time for at least 3 rehearsals of the whole to-be-recalled list. During this time overwriting can occur. If, instead of the rather arbitrary random overwriting in Neath's version of the feature model, feature vectors are generated to represent the irrelevant sound utterances overwriting can then proceed according to the within-list overwriting principles specified by Nairne (1990). To simulate the token set size, the number of feature vectors representing the irrelevant sound was varied. Unlike the previous reported simulation, there was no adjustment of the attention parameter between set size 1 (steady-state) and set size 2 or above (changing-state).

The data regarding the token set size effect are shown in Figure 3, together with a simulation study using the same procedure described here. The number of overwritings was set to 3 per item, and the items chosen to overwrite were randomly sampled from a set size of 2, 3, 5, or 7 randomly generated feature vectors. As Figure 3 clearly shows, this procedure produces a very good match between the performance of the model and the data from the experiment. Notably, the model actually provides a closer fit to the data than the predictions of Tremblay and Jones (1998).
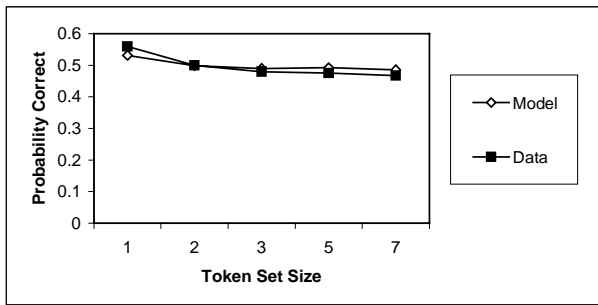
Figure 3. The effects of token set size of irrelevant sound on sampling probabilities in the revised model. The crucial difference between steady state and changing state conditions is represented by the difference between set sizes 1 and 2, and unlike in the feature model, is here reproduced accurately.

The Word-Dose Effect

Further evidence not specifically considered by Neath (2000) which is impossible to fit into his account without amendment includes the so-called "dose" effect (Bridges & Jones, 1996). This refers to the finding that increasing the absolute number of words in the irrelevant sound stream increases the size of the effect. Dose differs from token set size in that, for example, "A-B-A-B" has a set size of two but a dose of four. The word dose manipulation introduced by Bridges and Jones (1996, Experiment 1) shows strongly linear effects (see Figure 4) when recall performance is collapsed across presentation position of the to-be-recalled lists. An attempt to fit a linear trend line to these data yielded an R^2 value of .9978. The original feature model cannot account for these data because there is no mechanism within the model for relating probability of overwriting to number of irrelevant items presented. In the absence of this the model simply implements irrelevant sound interference of any type, regardless of the number of times each irrelevant item is presented as a single overwriting of each to-be-recalled item by a random combination of +1s and –1s. The problem presented by token set size effects is thus repeated, and the model cannot produce word dose effects.
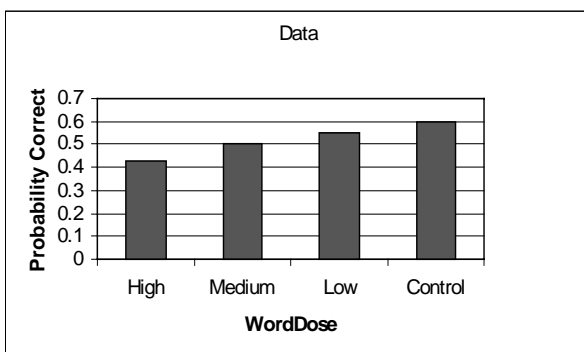

Figure 4. Moderating effects of word dose on disruption of serial recall by irrelevant sound.

However, as before, reconsideration of the word dose data suggests an alternative modelling formulation. Bridges and Jones presented 5 different speech items repeatedly over the 9 second period of presentation of the to-be-recalled lists, a 10 second retention interval and a 15 second response time (Bridges & Jones, 1996). If the simplifying assumption is made to exclude the response time from the analysis then in the high dose condition participants heard 57 separate utterances, in the medium dose, 29 utterances, and in the low dose 19 utterances.

It is not clear how the timing of the rehearsal coincides with the presentation of the irrelevant sound material, however the data indicate that a linear relationship exists across high, medium and low word "dose". In the next simulation therefore I assume one overwrite per item for the low dose condition, and increment the number of overwritings by one for the medium-dose and two for the high-dose conditions. The item chosen to overwrite each time will be chosen at random from a set of 5 vectors representing the 5 irrelevant sound items generated in the same manner as the vectors representing the to-be-recalled list. Overwriting will then proceed in the same manner as within-list overwriting. It is clear that this procedure ensures not only that overwriting becomes internally more consistent but also allows for simulation studies of such effects as word dose that are more directly motivated by the experimental procedure and do not resort to altering free parameters. The results of the word dose simulation are shown in Figure 5. Comparison of this figure to the data displayed in Figure 4 reveals that a reasonable qualitative fit to the data has been obtained. There is a discernible effect of word dose, to which a linear trend line can be fit with R^2=.9198, mirroring the linear trend observed in the experimental data.
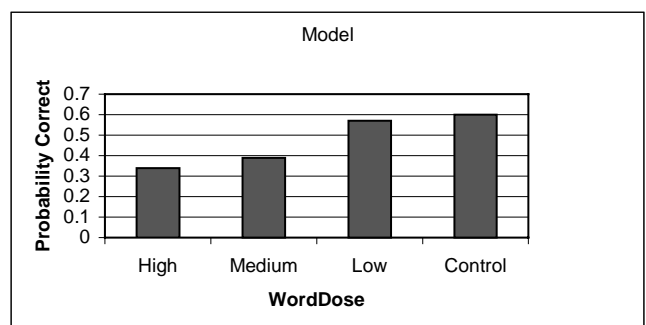

Figure 5. Effects of word-dose on sampling probabilities of the revised model.

General Discussion

Although the basic structure of the feature model was appropriated for this series of demonstrations, the intention was not to produce a simulation of irrelevant sound effects specific to the feature model. Instead, the intention was to investigate how some of the basic data regarding the irrelevant sound effect can emerge from an architecture in which items are represented in a distributed fashion and

presentation of irrelevant sound reduces the signal-to-noise ratio when recall of the order of those items is necessary. This investigation has succeeded in showing that increasing the noise in a distributed representation will reproduce many of the main findings in the irrelevant sound effect literature with relatively few assumptions. As such, there are three important points to note about this exercise.

The first point is to note that many of the feature model's assumptions, although implemented here, did not play any role in determining the outcome of the simulations. For example, although the assumption that overwriting occurs across features sharing the same value was implemented here, it is not necessary to make this assumption in order to obtain these results. Since each vector was constructed using random selections of binary values, the same result would be expected even if overwriting occurred across features with different values. It is possible to state with some confidence that reducing the signal-to-noise ratio by addition of noise to a distributed representation of the to-be-recalled item will therefore reproduce at least some of the key phenomena of irrelevant sound. The second point of note is that the simulations presented here reproduce many of the key characteristics predicted by Jones' changing-state hypothesis (Jones, Madden & Miles, 1992). These include: the changing-state effect itself, the specific disruption of order information, the word dose effect, and the lack of any great effect of token set size above 2 tokens. The simulations produce these effects, however, without the explicit representation of order cues assumed to be necessary by Jones.

The final point in favour of the current set of simulations is their relative parsimony and close correspondence to experimental procedure. Neath (2000) was criticized by Baddeley (2000) and Jones and Tremblay (2000) for the number of free parameters required in his simulations of irrelevant sound effects. The current set of simulations show that incrementing the attentional parameter is not necessary if the original (within-list) overwriting principles of the feature model are followed. This procedure provides a better fit to the data than the addition of the extra parameter. By explicitly matching the possibilities of overwriting to the rehearsal process it also proves possible to account for the word dose effect, which cannot otherwise be accounted for by the feature model. What is envisaged is an interference effect of discrete irrelevant sound elements on a continuous, serial, mental rehearsal process.

Acknowledgements

References

Baddeley, A. D. (2000). The phonological loop and the irrelevant speech effect: Some comments on Neath (2000). *Psychonomic Bulletin and Review*. In press.

Beaman, C. P., & Jones, D. M. (1997). Role of serial order in the irrelevant speech effect: Tests of the changing-state hypothesis. *Journal of Experimental Psychology: Learning, Memory and Cognition 23*, 459-471.

Beaman, C. P., & Jones, D. M. (1998). Irrelevant sound disrupts order information in free recall as in serial recall. *Quarterly Journal of Experimental Psychology 51A*, 615-636.

Bridges, A. M., & Jones, D. M. (1996). Word dose in the disruption of serial recall by irrelevant speech: Phonological confusions or changing state? *Quarterly Journal of Experimental Psychology 49A*, 919-939.

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

Ellermeier, W., & Zimmer, K. (1997). Individual differences in susceptibility to the "irrelevant speech effect". *Journal of the Acoustical Society of America*, *102*, 2191-2199.

Jones, D. M., & Macken, W. J. (1993). Irrelevant tones produce an irrelevant speech effect: Implications for phonological coding in working memory. *Journal of Experimental Psychology: Learning, Memory and Cognition, 19,* 369-381.

Jones, D. M., Madden, C., & Miles, C. (1992). Privileged access by irrelevant speech to short-term memory: The role of changing state. *Quarterly Journal of Experimental Psychology 44A*, 645-669.

Jones, D. M., & Tremblay, S. (2000). Interference in memory by process or content? A reply to Neath (2000). *Psychonomic Bulletin and Review*. In press.

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. (1999). Modelling the disruptive effects of irrelevant speech on order information. In: I. Neath, G. D. A. Brown, Poirier, M. & Fortin, C. (Ed.s). *Short-term/working memory*. Hove: Psychology Press.

Neath, I. (2000). Modeling the effects of irrelevant speech on memory. *Psychonomic Bulletin and Review*. In press.

Neath, I., & Nairne, J. S. (1995). Word-length effects in immediate memory: Overwriting trace decay theory. *Psychonomic Bulletin and Review 2*, 429-441.

Norris, D., Page, M., & Baddeley, A. D. (1995). Connectionist modeling of short-term memory. *Language and Cognitive Processes, 10*, 407-409.

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

Salamé, P., & Baddeley, A. D. (1982). Disruptions of short-term memory by unattended speech: Implications for the structure of working memory. *Journal of Verbal Learning and Verbal Behavior*, *21*, 150-164.

Tremblay, S., & Jones, D. M. (1998). The role of habituation in the irrelevant sound effect: Evidence from the role of token set size and rate of habituation. *Journal of Experimental Psychology: Learning, Memory and Cognition 24*, 659-671.