

Psychological Science

<http://pss.sagepub.com/>

Recognition and Source Memory as Multivariate Decision Processes

William P. Banks

Psychological Science 2000 11: 267

DOI: 10.1111/1467-9280.00254

The online version of this article can be found at:

<http://pss.sagepub.com/content/11/4/267>

Published by:



<http://www.sagepublications.com>

On behalf of:



[Association for Psychological Science](http://www.sagepublications.com)

Additional services and information for *Psychological Science* can be found at:

Email Alerts: <http://pss.sagepub.com/cgi/alerts>

Subscriptions: <http://pss.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Research Article

RECOGNITION AND SOURCE MEMORY AS MULTIVARIATE DECISION PROCESSES

William P. Banks

Pomona College

Abstract—*Recognition memory, source memory, and exclusion performance are three important domains of study in memory, each with its own findings, its specific theoretical developments, and its separate research literature. It is proposed here that results from all three domains can be treated with a single analytic model. This article shows how to generate a comprehensive memory representation based on multidimensional signal detection theory and how to make predictions for each of these paradigms using decision axes drawn through the space. The detection model is simpler than the comparable multinomial model, it is more easily generalizable, and it does not make threshold assumptions. An experiment using the same memory set for all three tasks demonstrates the analysis and tests the model. The results show that some seemingly complex relations between the paradigms derive from an underlying simplicity of structure.*

Over its long history, memory research has produced, and it shows every sign of continuing to produce, an abundance of experimental paradigms. In most cases, paradigms have been created to answer specific questions about one area of study, but the results obtained often lead to further experiments, and eventually to new fields of memory research. Although we have learned much about memory in this process, the effect is a tendency for the field to fragment, and to have major areas of research defined by paradigms whose results are difficult to relate to one another. This article presents an analysis of memory-based discrimination based on a single multidimensional representation that may bridge at least some gaps between domains of memory research.

The motivating interest here is modeling the relationship between item memory (sometimes termed old/new recognition), exclusion performance (used in measuring effects of familiarity; cf. Jacoby, 1991), and source memory, or memory for the origin of known material (Batchelder & Riefer, 1990; Johnson, Hashtroudi, & Lindsay, 1993). Although performance on these tasks often seems related in complex ways, the present analysis shows an underlying simplicity in the relations, and a principle that can be applied to other domains.

The memory representation assumed here is based on multivariate signal detection theory (SDT). The decision mechanism is a straightforward multivariate extension of the criterial process used in unidimensional detection models (see Banks, 1970; Green & Swets, 1974; Macmillan & Creelman, 1991; Swets, Tanner, & Birdsall, 1961). The current approach could be viewed also as an application, with simplifying assumptions, of the general recognition theory (GRT) of Ashby and colleagues (cf. Ashby, 1988, 1992; Ashby & Townsend, 1986) to memory, an application that in fact seems overdue.

This approach is offered as an alternative to multinomial analysis, which has been separately applied with some success to source

memory, recognition memory, and exclusion performance (Batchelder & Riefer, 1990; Buchner, Erdfelder, & Vaterrodt-Plünnecke, 1995; Wainwright & Reingold, 1996). However, it is not clear how a single multinomial model could simultaneously cover these three areas without generating a great number of parameters, as well as a very complex decision tree. In addition, the signal detection approach has many advantages over the multinomial approach, among them the lack of strong threshold assumptions (see Kinchla, 1994), and an intuitive spatial representation of the decision process.

MULTIDIMENSIONAL REPRESENTATIONS OF MEMORY INFORMATION

Learning confers discriminability between classes of previously homogeneous items. The object here is to show how the discriminations tested in different tasks—recognition, source memory, and exclusion—can be put into a single representation. In recognition memory, the discrimination is between items that were presented (old) and items that were not (new). In source memory, the discrimination is between Source A and Source B, rather than between old and new items.

In exclusion (Jacoby, 1991), there are two types of old items. Some old items, the included subset, are designated as targets, and are to receive a “yes” response. Other old items are to be excluded; that is, participants are to say “no” when these items are presented, just as they respond for new items, even though the items are old. Failure of discrimination for excluded old items is the matter of interest. The level of “yes” responses to these items has been taken as a measure of an automatic or unconscious component of their memory strength (cf. Jacoby, 1991).

In all of these discriminations, whether the measure is percentage correct or a signal detection measure like d' , the judgment has to be unidimensional, even though the memory representation may be multidimensional. The proposal offered here is that each of these unidimensional judgments is made by projection of the multidimensional configuration onto an appropriate unidimensional axis that can serve as the basis of the yes/no decision.

Consider a hypothetical set of data in which recognition memory (old/new discrimination) for seen items has a d' of 1.0, recognition memory for heard items has a d' of 0.75, and source discrimination between heard and seen items has a d' of 0.60. These interpoint distances do not allow the three points to lie on the same line, and thus make it impossible to place the three distributions on a common decision axis.

The solution is to construct a two-dimensional representation, as is shown in Figure 1, where the three distances form a triangle in the plane (see Tanner, 1956, for one representation of this sort, and Thomas, 1985, for another; see also Ashby & Townsend, 1986). Once the space is constructed, modeling is performed by projecting from the space onto unidimensional decision axes, and applying SDT analyses

Address correspondence to William P. Banks, Psychology Department, Pomona College, Claremont, CA 91711; e-mail: wbanks@pomona.edu.

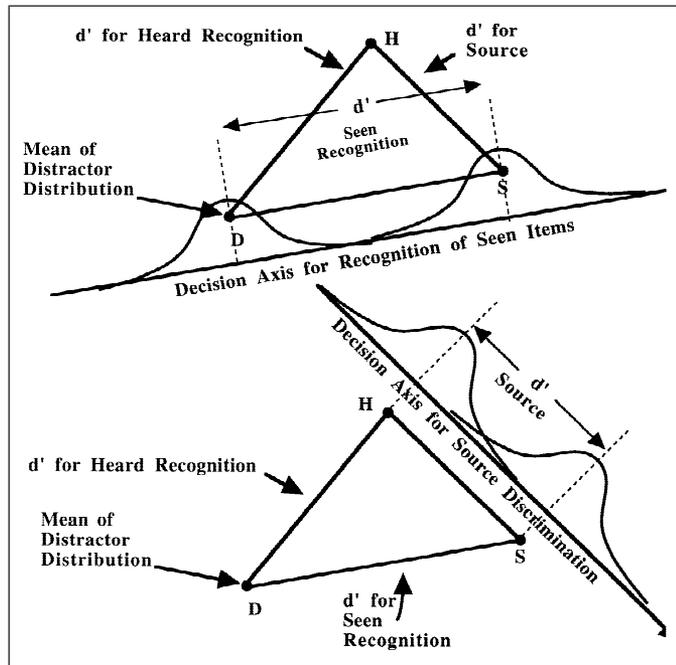


Fig. 1. Illustration of a two-dimensional signal detection model necessitated by the triangle inequality. The means of the distributions are noted as points S (seen), H (heard), and D (distractor). Because the d' for discriminating D and S is less than the sum of d' s for discriminating D and H and H and S, the mean of the H distribution cannot lie on the axis from D to S, and a second dimension is needed. Here, the three distances are arranged in a bivariate normal space. Projections are shown for constructing a distance axis for seen recognition and for source discrimination.

to the resulting distributions. The angle taken by the decision axis through the space determines the pattern of projections of the distributions and consequently which memory task is predicted. Figure 1 illustrates projections of the distributions for recognition of seen items and for source discrimination between seen and heard items. Likelihood axes are drawn between the means, and the bivariate distributions are projected onto these axes, which function as the underlying representations for the decisions.

These plots assume orthogonal axes and equal variance along both axes. They also assume zero correlation between the distributions and orthogonal projection of distributions. These assumptions amount to severe restrictions on freedom of modeling when compared with the full range of assumptions considered within GRT. As will be seen, however, the fits with these assumptions are adequate, and the fitting process is greatly simplified under these constraints.

In Figure 1, the space is simply a by-product of the triangle inequality. In fitting real data, there may be several ways to create the space. The method used for the present data is multidimensional scaling (MDS). MDS provides an objective and repeatable method for generating a space, but the test should be how well it predicts the different tasks rather than the method by which it was created. Another approach, as in most of the chapters in Ashby (1992), is to plot the results of one task against the results of the other on orthogonal axes. Whatever approach is taken, it is up to the ingenuity of the investigator to devise the multidimensional representation, which is a theoretical construct intended to predict the results.

METHOD

The strategy of this experiment was to derive a two-dimensional representation from source and old/new memory performance by a first group of participants, and then to use this space to predict performance on other tasks by two new groups. One new group was tested only on seen- and heard-exclusion tasks. The other new group was tested on source and item memory for restricted target sets. The restricted tasks were (a) source memory judgments for only foils and heard items (no seen items) and (b) item memory tests with only foils and seen items (no heard items). The restricted sets provide an additional test of the model that examines its ability to handle criterial changes.

Participants

There were three groups, comprising 24, 12, and 12 participants, respectively. All participants were male or female students at the Claremont Colleges. They ranged in age from 18 to 25 and were paid \$5 each for their participation.

Materials

Two types of item and two sources were used. The two item types were words and men's first names. Two sources were created by presenting some items visually (seen items), on a computer screen, and other items auditorily (heard items), through the computer loudspeaker.

Word lists for targets and foils were balanced for frequency (mean of 57.1 occurrences per million; Francis & Kuçera, 1982) and counterbalanced across foil and target positions, as well as across tasks. Names were chosen for recognizability as men's common first names. All presentations were controlled by SuperLab. Seen items were presented in uppercase on a 14-in. color monitor, at a distance of 50 to 60 cm, with a font approximately 1 cm high. Heard items were read by a male and digitized so that they could be presented through the loudspeaker of a Macintosh PPC computer. After receiving instructions and practice items, participants were presented with 20 items (10 words and 10 names randomly intermixed) on the screen and then 20 items (10 words and 10 names, also randomized) auditorily.

After study, participants were tested with response sheets, each with 30 items (15 names and 15 words), and they were asked to respond to each item with a number that reflected their confidence that it was old or new (item memory) or seen or heard (source memory, seen exclusion, heard exclusion). Each sheet of 30 items listed 5 seen names, 5 heard names, 5 name foils, 5 heard words, 5 seen words, and 5 foil words. Targets and foils were distributed between sheets in such a way that no item was repeated across sheets for a given participant. The order of test sheets was counterbalanced over participants. Candidate words or names, drawn from target and foil sets, were listed in random order along the left margin, with a space by each item for a response on a scale from 1 to 8.

The 24 participants in the initial group were tested only for source and item memory. On one sheet, they gave a higher number to an item to indicate their confidence that it was old; on the other, they gave higher ratings to seen items and lower ratings to heard items. Exclusion (Jacoby, 1991) for the same items was tested with a second group of 12 participants, who were given the following instructions: "Please

choose only the items in this list you heard [saw] in this experiment. Give each item a rating of 8 if and only if you are sure it was heard [seen]; reject it if it was either seen [heard] or not presented at all by giving it a rating of 1." A final group of 12 participants was given the same source and item tests as the first group of 24, except that for each test foils and items came from only one target set. Thus, recognition testing included only seen and new items, and source memory testing included only heard and new items.

Procedure

Participants were told that they were to be tested on memory for words and male names, some seen on a computer screen and some heard from the computer's loudspeaker. They were instructed to read each visually presented word aloud and press the space bar on the keyboard after they had finished. The next item was presented when the space bar was pressed, or after 2 s, whichever came first. While each auditory word was being played, SuperLab presented a blank screen for 2 s. No auditory word required more than 2 s, and the screen stayed in view until the 2 s had elapsed.

RESULTS AND DISCUSSION

Analysis

The first step in applying the model was to generate a memory space from source and item memory judgments using the data of the first group of participants. This space is an MDS representation of the interitem distances as measured by *da*. Table 1 presents the interitem distances used for the MDS analysis. The number in any given cell is the distance between the corresponding row and column categories of items in *da* units. For each cell, a separate receiver operating characteristic (ROC) was constructed, with cumulative proportions for one set of items plotted against the other. The area under the ROC was computed with ROCFIT (Metz, 1991). The area score was then transformed to *da* to give a detection metric (cf. Simpson & Fitter, 1973).

Some cells contain measures that are not normally reported. One of these is the comparison between word foils and name foils, with a *da* of 0.283. This *da* is derived from an ROC based on the judgments of likelihood that the two sorts of foil are targets, and it measures the difference between these two categories of nontargets in their likeli-

hood of seeming like a target. Measurements like this could be considered scale values of discrimination as much as a detection measure. They measure the distance between items on the memory attribute under test.

The interitem distances of Table 1 were submitted to the ALSCAL routine of the SPSS package (Mac version 6.1.1) for MDS. This routine arranges the items in a spatial layout so that the distances among items in the space correspond as closely as possible to the distances in the matrix. The fitting algorithm operates to reduce stress, which is a measure that compares interitem distances in the space with interitem distances in the original matrix (cf. Kruskal & Wish, 1978; Young, 1970). A perfect fit would have a stress of 0.0. In this case, a two-dimensional fit was good, with stress of .064, and an r^2 for the relation between original and derived distances of .973. A spatial portrayal of the MDS space is seen in Figure 2.

Several obvious features of the data are captured in this plot. The distance between seen and heard names is less than the distance between seen and heard words, reflecting the greater source confusion for names than words. Name foils are closer than word foils to the target means, as is to be expected if there are higher false alarms for name foils. More regularities are seen when projections onto decision axes are computed.

Fitting the results of the different memory tasks was accomplished by projecting the points onto axes drawn through the space, with a different angle for the axis for each task. The graphic representation of this process is shown in Figures 3, 4, and 5. Table 2 describes how fitting was done mathematically, using as an example the seen-exclusion results. In this table, the first two columns are the values of the points on the two dimensions found by ALSCAL for the data set. The projection is the dot product ($x \cdot \sin\theta + y \cdot \cos\theta$) of the values on the two dimensions, which is a unidimensional number. The test of the projection is the measure of how well it compares with the data from the task being predicted, using in this case r as the measure of fit. The value of θ was varied until the best fit was found. The best-fitting value of θ is shown in the table.

Predictions

Table 3 shows the mean probability of "yes" as a response to the instructions for each of the six tasks. For exclusion, a positive response means that the item was judged to have been heard or seen,

Table 1. Distance matrix among items for the initial group of subjects

	Heard words	Seen words	Word foils	Heard names	Seen names	Name foils
Heard words	0					
Seen words	1.592	0				
Word foils	1.633	1.553	0			
Heard names	0.035	1.454	1.527	0		
Seen names	1.294	0.127	1.896	1.226	0	
Name foils	1.459	1.372	0.283	1.332	1.697	0

Note. Cell entries are the measure *da* based on the area under receiver operating characteristics (ROCs). The values of *da* comparing items with foils were derived from the ROCs for recognition memory; all other values for *da* were derived from source memory.

Decisions in Memory Space

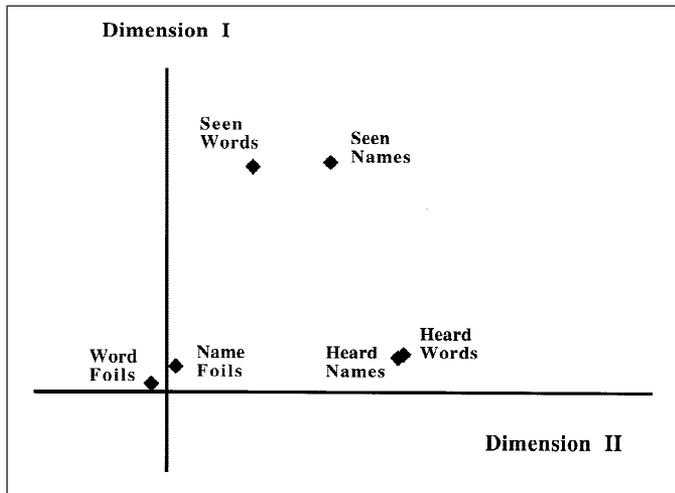


Fig. 2. Bivariate space representing the data of the initial group of subjects. This is the output of multidimensional scaling of the discriminabilities in Table 1, rotated 45° clockwise from the ALSICAL fit to this orientation for convenience.

depending on the instructions. For source memory, the higher ratings are for seen items, so the “yes” proportion reflects the proportion of “seen” responses an item received.

The raw recognition data show that the heard words had a slight advantage over the heard names. This advantage is larger in terms of the detection measure *da*, which was 1.633 for heard words and 1.332 for heard names (see Table 1 for these *das*). Seen names yielded a higher hit rate than seen words, but the higher false alarm rate for name than word foils reduces the slight overall advantage for names, to *das* of 1.697 for seen names and 1.553 for seen words.

The reason for having names as well as words in both sources was to create one category of item for which source was less discriminable than the other. As expected, source was less discriminable for names

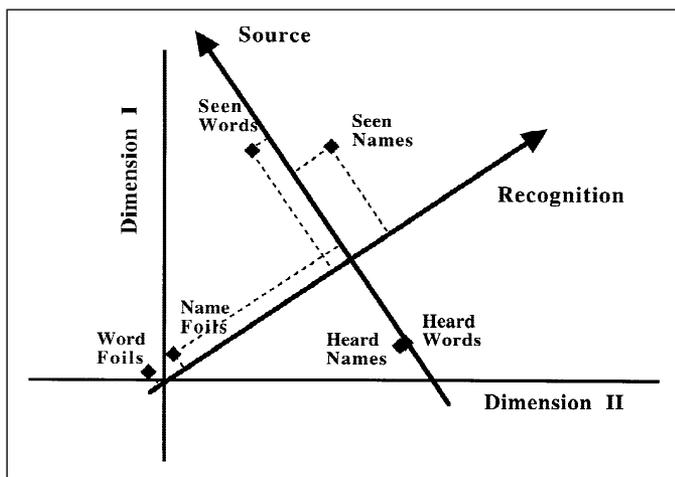


Fig. 3. Decision axes for recognition (old/new discrimination) and source memory (seen vs. heard discrimination) in the initial group of subjects, drawn in the space shown in Figure 2. Dotted lines illustrate the projections of the means of the distributions on the axes. The predictions derived from these projections are shown in Table 3.

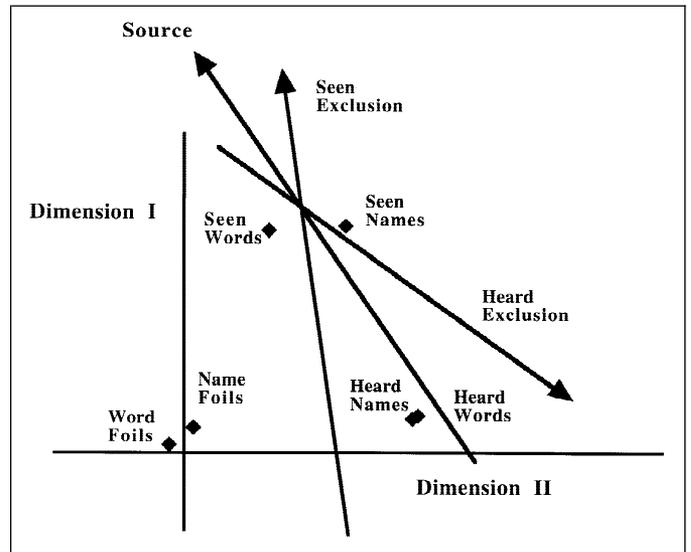


Fig. 4. Decision axes for seen exclusion and heard exclusion, drawn in the bivariate space for the initial group of subjects. The source memory axis is shown for comparison. Projections are not illustrated, but the reader can drop perpendiculars to the two exclusion axes to see intuitively the pattern of results predicted. The point predictions are in Table 3.

than words. The *da* was 1.592 for discriminating the source for words and 1.226 for discriminating the source for names. This pattern is also seen in the probabilities.

Heard exclusion (heard items are positive), the third task listed in Table 3, showed an interaction: Heard words were accepted with a probability of .617 and heard names with a probability of .567, but excluded seen items had the opposite pattern, with an acceptance rate of .350 for seen words but .400 for seen names. The seen-word acceptance rate was lower than the false alarm rate for words (.450), leading to a negative *da*, -0.26. Seen names, in contrast, had a higher acceptance rate than the name foils (.300), and their *da* was a positive 0.27.

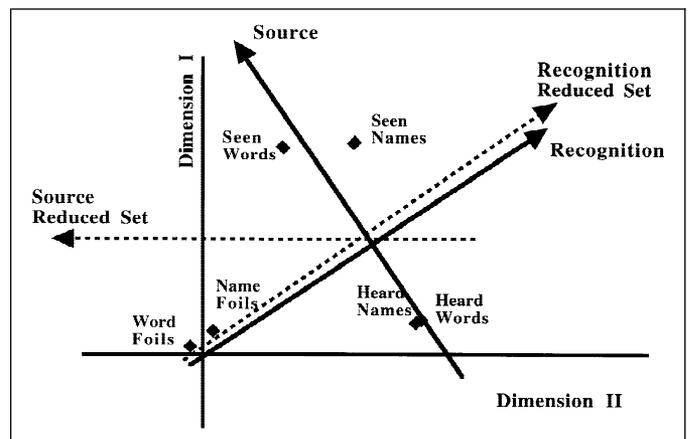


Fig. 5. Decision axes for the reduced-set tests (dotted lines) and the original source and recognition axes (solid lines). Point predictions are in Table 3.

Table 2. Example showing how predictions from the multidimensional scaling representation were made for seen exclusion

Item type	x value	y value	$x \cdot \sin\theta + y \cdot \cos\theta$	Observed seen exclusion	z score for observed seen exclusion	Predicted seen exclusion
Heard word	0.448	-1.31	-0.776	.433	-0.169	.371
Seen word	0.742	1.2	1.40	.750	0.674	.683
Word foil	-1.61	0.314	-0.716	.420	-0.202	.379
Heard name	0.382	-1.28	-0.772	.333	-0.432	.369
Seen name	1.31	0.66	1.33	.600	0.253	.673
Name foil	-1.31	0.264	-0.576	.333	-0.432	.399

Note. The x and y values were taken from the multidimensional scaling solution of the da distances in Table 1. The observed probability for seen exclusion is the probability of acceptance (confidence ratings of 5 or greater) in the seen-exclusion condition, in which seen items were positive and heard words and foils were to be rejected. The dot product ($x \cdot \sin\theta + y \cdot \cos\theta$) of the values of the two dimensions for each point is a scalar quantity that represents the projection of the point onto a vector passing through the space at angle θ to the x -axis. The value of θ was determined by regressing the dot products against the z score for the observed seen exclusion. The θ that maximized r was found with the use of Solver in an Excel spreadsheet. For this case, the optimal θ was 53° . The r relating predicted to obtained probabilities (without the z transformation) was .92.

The fourth task in Table 3 is seen exclusion (seen items are positive). The positive seen-word category achieved the same level, .750, attained in recognition, but performance was worse because the foils were accepted at a much higher rate, with a false alarm rate of .417 (resulting in a da of 0.884 in exclusion), compared with .183 in recognition (and a da of 1.58). The hit rate for seen names was .600, lower than the hit rate of .825 in recognition. The excluded category, heard items, had acceptance rates of .433 for words and .333 for names.

Several aspects of the pattern of results in exclusion are difficult to explain in terms of unidimensional notions of memory strength. The drop in performance for the target items from recognition to exclusion (also noted by Dodson & Johnson, 1996) is one of these aspects. Why would the same positive items have lower rates of acceptance and smaller das in one task than the other? If it were simply a matter of

criterion setting, measures like da or d' would not be reduced, but they are in fact reduced.

Furthermore, negative values of da or d' in exclusion for items that had large and positive discriminability measures in inclusion are not easy to explain with unidimensional strength theories. If some strategy is invoked to explain this (e.g., that the rejected target items were recognized as such and consequently given negative ratings rather than positive ones), the negative values of da or d' should be close to the positive values in the simple recognition task. However, the absolute value is much smaller for the negative da than the positive one, and this seemingly reasonable strategy cannot explain the results.

To make matters worse for a strategic explanation, the categories of items that are positive in recognition testing, such as the names and words here, often have different relative discriminabilities, whether positive or negative, in exclusion. Heard words and heard names had

Table 3. Probability of acceptance of items (confidence ratings of 5, 6, 7, or 8) and predictions of these probabilities derived from the bivariate representation

Item type	Recognition memory (initial group)		Source memory (initial group)		Heard exclusion		Seen exclusion		Recognition memory: Limited set		Source memory: Limited set	
	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted
Heard words	.800	.760	.242	.251	.617	.591	.433	.371	—	—	.283	.305
Seen words	.750	.720	.804	.795	.350	.338	.750	.683	.600	.601	—	—
Word foils	.183	.180	.513	.516	.450	.375	.417	.379	.150	.145	.500	.503
Heard names	.725	.745	.267	.254	.567	.586	.333	.369	—	—	.342	.319
Seen names	.825	.856	.708	.718	.400	.404	.600	.673	.750	.749	—	—
Name foils	.225	.240	.517	.518	.300	.386	.333	.399	.183	.190	.400	.397
θ^a	34°		124°		-33°		98°		35°		-91°	
r	.995		.999		.90		.920		.999		.980	

^aValues of θ are relative to the transformed coordinates (Figs. 2–5; rotated 45° clockwise), with positive values indicating a counterclockwise rotation from $\theta = 0.0$, the horizontal axis.

Decisions in Memory Space

das of 1.633 and 1.332 in recognition, and both dropped essentially to zero, 0.03 and 0.0, respectively, in exclusion with seen as the positive category. Likewise, the *da* of 1.697 for seen names dropped to 0.27 in exclusion with heard the positive category, and the *da* for seen words dropped from 1.58 to a negative *da*, -0.26 . It is impossible to account for these differences by a simple linear shift of the noise distributions, or by a strategic explanation that is not blatantly ad hoc. As will be seen, these results are nicely predicted by the projections on the decision axis.

Fitting was accomplished by plotting r as a function of θ , as θ was varied through 360° , and then searching for the best-fitting θ at the peak of the function. The Excel routine Solver was used to confirm the best-fitting θ . The value of θ that gave the best fit was used to compute the predicted "yes" proportions in Table 3.

Figure 3 shows the decision axes that give the best fit for recognition and source discrimination. The best-fitting axes for recognition and source memory are exactly orthogonal to each other, at 34° for recognition and 124° for source discrimination. (This was also the case for a similar data set in Banks, Chen, & Prull, 1999.) The excellent fit, $r = .995$ for recognition and $.999$ for source (see Table 3), might not be a surprise, because source and recognition *das* went into the matrix on which the scaling was based. However, there was no guarantee that these *das* would combine to constrain the space in this way, or that a single criterion applied to all six distributions would predict the "yes" probabilities so well.

The best-fitting vectors for both heard and seen exclusion are plotted in Figure 4. These exclusion fits (see Table 3) are not quite as good as the source and recognition fits, but they are based on a different, and smaller, sample. Nevertheless, the point predictions of exclusion performance seen in Table 3 capture the patterns of performance that were not compatible with unidimensional strength.

The exclusion decision axes are almost exactly symmetrical about the source memory axis, also shown in the figure. The seen-exclusion axis is 26° clockwise of source memory, and the heard-exclusion vector is 23° counterclockwise. Although many researchers have observed that exclusion seems to be a form of source memory, the pattern of results is not the same in exclusion as in source memory testing, and it has been difficult to relate the two. Here we have a precise relationship between source memory and exclusion, and one that could not be obtained from a unidimensional strength theory.

Figure 5 shows the decision axes for the limited-set tasks, with the full-set versions shown for comparison. The elimination of heard items from the set presented for recognition testing moved the recognition decision axis only 1° . In contrast, elimination of seen items from source testing created a major shift in the decision axis that seems adaptive to the task. This shift came about because the task required that seen items get a high rating and heard items a low one, with new items in between. When there were no seen items in the judgment set, this orientation of the decision vector gave a better discrimination of new and heard items.

CONCLUSIONS REGARDING SOURCE, ITEM, AND EXCLUSION TESTS

These findings show that it is possible to give an account of item recognition, source memory, and exclusion with a single spatial representation of the memory information. The different tasks are modeled as different decision axes, or vectors, onto which the distributions

are projected for decision in the various tasks. These projections are able to model relations between the tasks that are very difficult for unidimensional strength theories to account for. The directions of the vectors in the memory space also reveal some interesting relations among the tasks. Source memory and item recognition memory turn out to be orthogonal tasks, a result supporting proposals that they are independent (Johnson, Kounios, & Nolde, 1996). However, at the same time, they use the same memory database, which suggests that in a memory task one uses the same information in different ways for source and item memory decisions.

The hypothesis of separate familiarity and recollection components may be a result of the presupposition that memory strength is unidimensional. Rather than adding new memory components to save unidimensional strength, with debatable rules of combination, implausible threshold assumptions, and a certain ad hoc flavor (do these two components operate in source memory, e.g., or are they special to recognition and exclusion?), it seems better to postulate a multidimensional model of memory that can handle a number of paradigms with a change in a single parameter. When such a model was applied here, the seemingly complex relations among the results of the tasks turned out to have a simple underlying structure, intuitively represented spatially.

Spatial representations have proved useful in many areas of psychology, including perception, sensation, categorization, semantic memory, and spatial vision (see Ashby, 1992; Luce, D'Zmura, Hoffman, Iverson, & Romney, 1995; Nosofsky, 1992; Thomas, 1985; Thomas & Olzak, 1992). Spatial representations and decision models have rarely been used in the field of memory (but see Nosofsky, 1991). This is surprising because some of the central models of memory assume multidimensional components (Bower, 1996; Gillund & Shiffrin, 1984; Hintzman, 1986; Murdock, 1982), and most theorists have tested their models with signal detection measures. The present approach could easily be a natural extension of some of these models.

Acknowledgments—This research was supported by a Sontag grant and a grant from the Irvine Foundation. I thank Adolfo Rumbos and William Prinzmetal for their invaluable suggestions during the development of the model reported here and Matthew W. Prull for his careful reading of the manuscript and astute advice on the research. Edwin Wong and Andrea Chung deserve recognition for their good work in the data collection and analysis.

REFERENCES

- Ashby, F.G. (1988). Estimating the parameters of multidimensional signal detection theory from simultaneous ratings on separate stimulus dimensions. *Perception & Psychophysics*, *44*, 195–204.
- Ashby, F.G. (Ed.). (1992). *Multidimensional models of perception and cognition*. Hillsdale, NJ: Erlbaum.
- Ashby, F.G., & Townsend, J.T. (1986). Varieties of perceptual independence. *Psychological Review*, *93*, 154–179.
- Banks, W.P. (1970). Signal detection theory and human memory. *Psychological Bulletin*, *74*, 81–99.
- Banks, W.P., Chen, Y.-P., & Prull, M.W. (1999). Memory and awareness: Is memory information stratified into conscious and unconscious components? In B.H. Challis & B.M. Velichkovsky (Eds.), *Stratification in cognition and consciousness* (pp. 129–172). Amsterdam: Benjamins Publishing.
- Batchelder, W.H., & Riefer, D.M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, *97*, 548–564.
- Bower, G. (1996). Reactivating an activation theory of implicit memory. *Consciousness and Cognition*, *5*, 27–72.

- Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward unbiased measurement of conscious and unconscious memory processes within the process dissociation framework. *Journal of Experimental Psychology: General*, *124*, 137–160.
- Dodson, C.S., & Johnson, M.K. (1996). Some problems with the process-dissociation approach to memory. *Journal of Experimental Psychology: General*, *125*, 181–194.
- Francis, W.N., & Kuçera, H. (1982). *Frequency analysis of English usage: Lexicon and grammar*. Boston: Houghton Mifflin.
- Gillund, G., & Shiffrin, R.M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, *91*, 1–65.
- Green, D.M., & Swets, J.A. (1974). *Signal detection theory and psychophysics*. Huntington, NY: Krieger.
- Hintzman, D. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.
- Jacoby, L.L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Johnson, M.K., Hashtroudi, S., & Lindsay, D.S. (1993). Source monitoring. *Psychological Bulletin*, *114*, 3–28.
- Johnson, M.K., Kounios, J., & Nold, S.F. (1996). Electrophysiological brain activity and memory source monitoring. *NeuroReport*, *7*, 2929–2932.
- Kinchla, R.A. (1994). Comments on Batchelder and Riefer's multinomial model for source monitoring. *Psychological Review*, *101*, 166–171.
- Kruskal, J.B., & Wish, M. (1978). *Multidimensional scaling*. Beverly Hills, CA: Sage.
- Luce, R.D., D'Zmura, M., Hoffman, D., Iverson, G.J., & Romney, A.K. (1995). *Geometric representations of perceptual phenomena: Papers in honor of Tarow Indow*. Hillsdale, NJ: Erlbaum.
- Macmillan, N.A., & Creelman, C.D. (1991). *Detection theory: A user's guide*. New York: Cambridge University Press.
- Metz, C.E. (1991). ROCFIT (Apple Macintosh Version) [Computer software]. Chicago: University of Chicago, Department of Radiology.
- Murdock, B.B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, *89*, 609–626.
- Nosofsky, R.M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 3–27.
- Nosofsky, R.M. (1992). Similarity scaling and cognitive process models. *Annual Review of Psychology*, *42*, 25–53.
- Simpson, A.J., & Fitter, M.J. (1973). What is the best index of detectability? *Psychological Bulletin*, *80*, 481–488.
- Swets, J.A., Tanner, W.P., & Birdsall, T.G. (1961). Decision processes in perception. *Psychological Review*, *68*, 301–340.
- Tanner, W.P. (1956). Theory of recognition. *Journal of the Acoustical Society of America*, *28*, 882–888.
- Thomas, J.P. (1985). Detection and identification: How are they related? *Journal of the Optical Society of America A*, *2*, 1457–1467.
- Thomas, J.P., & Olzak, L.A. (1992). Simultaneous detection and identification. In F.G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 253–277). Hillsdale, NJ: Erlbaum.
- Wainwright, M.J., & Reingold, E.M. (1996). Response bias correction in the process dissociation procedure: Approaches, assumptions, and evaluations. *Consciousness and Cognition*, *5*, 232–254.
- Young, F.W. (1970). Nonmetric multidimensional scaling: Recovery of metric information. *Psychometrika*, *35*, 455–473.

(RECEIVED 9/14/98; ACCEPTED 1/12/99)