

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The Natural Input Memory Model

Permalink

<https://escholarship.org/uc/item/2cd6x26h>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 26(26)

ISSN

1069-7977

Authors

Lacroix, Joyca P.W.
Murre, Jaap M.J.
Postma, Eric O.
et al.

Publication Date

2004

Peer reviewed

The Natural Input Memory Model

Joyca P.W. Lacroix (j.lacroix@cs.unimaas.nl)

Department of Computer Science, IKAT, Universiteit Maastricht, St. Jacobsstraat 6, 6211 LB Maastricht, The Netherlands

Jaap M.J. Murre (jaap@murre.com)

Department of Psychology, Universiteit van Amsterdam, Roeterstraat 15, 1018 WB Amsterdam, The Netherlands

Eric O. Postma (postma@cs.unimaas.nl)

H. Jaap van den Herik (herik@cs.unimaas.nl)

Department of Computer Science, IKAT, Universiteit Maastricht, St. Jacobsstraat 6, 6211 LB Maastricht, The Netherlands

Abstract

A new recognition memory model is proposed which differs from the existing memory models in that it operates on natural input. Therefore it is called the natural input memory (NIM) model. A biologically-informed perceptual pre-processing method takes local samples from a natural image and translates these into a feature-vector representation. The feature-vector representations reside in a similarity space in which perceptual similarity corresponds to proximity. By using the similarity structure of natural input, the model by-passes assumptions about distributional statistics of real-world input. Our simulations on the list-strength effect, the list-length effect, and the false memory effect support the validity of the proposed model. In particular, we conducted a face recognition simulation with the NIM model and found that it is able to replicate well-established recognition memory effects that relate to the similarity of the input.

Memory Representation

Many computational memory models represent an item by a vector of abstract features (e.g., the SAM model, Raaijmakers & Shiffrin, 1981; the REM model, Shiffrin & Steyvers, 1997, the model of differentiation, McClelland & Chappell, 1998). The feature values are usually drawn from a mathematical distribution (e.g., a geometric distribution). Since the computational models artificially generate vector representations, they do not address the contribution of the similarity structure intrinsic to natural data. However, we believe that the similarity structure contains important information. Therefore, we propose a memory model that operates on natural data and represents the similarity structure of these data.

The similarity structure of natural data can be represented in any type of space that fulfills the compactness criterion (Arkadev & Braverman, 1966). This criterion is fulfilled when similar objects in the real world are close in their representations. Several researchers developed so called ‘similarity spaces’, in which representations of similar items are in close proximity of each other (e.g., Nosofsky, 1986; Steyvers, Shiffrin, & Nelson, in press). An analysis of human similarity judgments or of free association data often forms the basis of a similarity space. However, we propose to derive the similarity space from the natural data by employing a biologically-informed transformation.

In the next section, a new recognition memory model that operates on natural images is introduced and described. We call this model the natural input memory (NIM) model. We will conduct a face recognition simulation with the NIM model and will evaluate its ability to replicate findings from

recognition-memory studies. Finally our main conclusion will be given.

The NIM Model

The NIM model encompasses the following two stages.

1. A perceptual pre-processing stage that translates a natural image into a number of feature vectors.
2. A memory stage comprising two processes:
 - (a) a storage process that simply stores feature vectors;
 - (b) a recognition process that compares feature vectors of the image to be recognized with previously stored feature vectors.




Figure 1: The natural input memory (NIM) model.

Figure 1 presents a schematic diagram of the NIM model. The face image is an example of a natural image. The two boxes correspond to the perceptual pre-processing stage and the memory stage.

The Perceptual Pre-Processing Stage

In this section, we first provide some background on the sources of biological inspiration and on the computational considerations. Then, we discuss some relevant implementation details.

Biological Inspiration and Computational Considerations

The human visual system is our main source of biological inspiration. The eye sequentially fixates on those parts of a visual scene that are most informative for recognition (e.g., Yarbus, 1967). Early visual processing in the brain leads to the activation of millions of optic nerve cells (Palmer, 1999). The nerve-cell activations may be conceived as a high dimensional vector. The high dimensionality enables the representation of a large amount of information without suffering from interference (Rao & Ballard, 1995), but it also hampers the memory performance, as the number of examples

that is necessary for a reliable generalization performance grows exponentially with the number of dimensions. This phenomenon is known as the ‘curse of dimensionality’ (Bellman, 1961; Edelman & Intrator, 1997). In coping with the curse of dimensionality, subsequent stages in the visual system are assumed to reduce the dimensionality of the high-dimensional input (e.g., Hubel, 1988; Tenenbaum, Silva, & Langford, 2000). The assumption is supported by findings of Edelman and Intrator (1997), who showed that the human visual system is able to retrieve the intrinsic low-dimensional structure of the high-dimensional visual input.

In the NIM model, dimension reduction of high-dimensional natural input is achieved in two sequential steps: (1) a biologically-informed feature-vector extraction (Freeman & Adelson, 1991) followed by (2) a principal component analysis (Pearson, 1901). The feature-vector extraction method employed by the NIM model operates directly on a high-dimensional natural image. The image has a high dimensionality because it is treated as a vector, the elements of which are the constituent pixel values. Motivated by eye fixations in human vision, the feature-vector extraction method takes samples from randomly-selected locations along the contours in the image. To emphasize the parallel with human vision, we refer to the samples as ‘fixations’. For each fixation, the NIM model extracts features (i.e., a feature vector) from the image area centered at the fixation location. Since the feature vector contains a limited number of features, it is of a much lower dimensionality than the image. The feature-vector extraction method is based on the visual processing generally believed to occur in the visual area V1. The responses of neurons in V1 are modeled by a multi-scale wavelet decomposition (described later). Several studies showed that the biologically-informed multi-scale wavelet decomposition results in a representation space that accurately represents similarities as perceived by humans (e.g., Kalocsai, Zhao, & Biederman, 1998; Lyons & Akamatsu, 1998; Bartlett, Littleworth, Braathen, Sejnowski, & Movellan, 2003). After extraction of feature vectors, principal component analysis represents the feature vectors by their projection onto a number of orthogonal basis vectors which are ordered according to the amount of variance they explain. The dimensionality of the feature vectors is reduced by taking the projection onto the first p basis vectors. The low-dimensional feature vectors reside in a similarity space where visual similarity translates to proximity of feature vectors. Translating a two-dimensional image using a multi-scale wavelet decomposition followed by a principal component analysis, is an often applied method in the domain of visual object recognition to model the first three stages of processing of information in the human visual system (i.e., retina/LGN, V1/V2, V4/LOC; Palmeri & Gauthier, 2004). In contrast, existing memory models lack such a pre-processing method and often make simplifying assumptions about object representations.

Implementation The input image is translated into a multi-scale representation at four spatial scales. At every scale, the image is processed by four oriented filters in the orientations 0° , 45° , 90° , and 105° using the steerable-pyramid transform (Freeman & Adelson, 1991). This processing results in sixteen (four scales times four orientations) filtered images. From each of the sixteen images a 7×7 window is

selected at a fixation point and the 16×49 pixel values are placed in a vector. In addition, the pixel values of a 7×7 low-resolution subimage centered at the fixation point are appended to the vector. Fixation points are randomly drawn from the contours of the face. The result is a feature vector for each fixation. As mentioned before, a principal component analysis was used to reduce the dimensionality of the feature vectors by taking the projection onto the first p basis vectors.

The Memory Stage

The Storage Process In the NIM model, the storage process straightforwardly stores an item (i.e., a pre-processed natural image). A pre-processed natural image is represented by a number of low-dimensional feature vectors in the similarity space, each corresponding to an eye fixation. The storage strength, S , is defined as the number of feature vectors stored for an image.

The Recognition Process In the NIM model, the recognition process determines the familiarity of an image to be recognized by comparing feature vectors of the image to be recognized with previously stored feature vectors. Models with a recognition process based on comparing items to previously stored exemplars can provide an accurate quantitative account of recognition performance (Medin & Schaffer, 1978; Nosofsky, 1986; Nosofsky, Clark, & Shin, 1989). In the NIM model, the recognition process uses a nearest neighbor classifier method, which takes each feature vector of the image to be recognized and then determines the number of previously stored feature vectors, f , that fall within a hypersphere with radius r , centered around the feature vector of the image. The familiarity, F , of the image is defined as $\sum f_i / T$, with f_i the value of f for the i^{th} feature vector of the image, and T the total number of feature vectors of the image.

We expect that the similarity-space representations employed by the NIM model will deepen our understanding of human recognition memory. Moreover, they may effectively support a number of memory effects often obtained in recognition memory studies. The latter studies are described in the next section.

Human Recognition Memory Studies

Three recognition memory effects often found in recognition memory studies are: the list-strength effect, the list-length effect, and the false memory effect. In general, recognition memory studies provide subjects with a study list of items and test their recognition memory for (some of) the studied items (i.e., targets) and a number of non-studied items (i.e., lures). We will emphasize the relation between the similarity structure of the targets and the lures used in the experiments on the one hand and the memory effects on the other hand.

The List-Strength Effect

A list-strength effect is defined as: a decrease in memory performance for a given set of study list items when other items of the study list are “strengthened” (i.e., the amount of time or the number of times the items are studied is increased) (Ratcliff, Clark, & Shiffrin, 1990). While some researchers failed

to find a list-strength effect for recognition memory (e.g., Ratcliff et al., 1990), recent findings showed that a list-strength effect can be obtained when there is a high degree of similarity between targets and lures. Norman (2002) tested whether strengthening some words of the study list affected a subject's recognition performance for other (non-strengthened) studied words. In the experiments, a significant list-strength effect was obtained only when targets and lures were similar. For dissimilar targets and lures, no list-strength effect was found. Moreover, recognition scores were significantly higher for dissimilar targets and lures than for similar targets and lures.

The List-Length Effect

A list-length effect is defined as: a decrease in memory performance for the items of the study list when additional items are added to the study list (Ratcliff et al., 1990). List-length studies yielded contradictory results. While some researchers failed to find a list-length effect (e.g., Dennis & Humphreys, 2001), others did obtain it (e.g., Cary & Reder, 2003). Recent experimental results indicate that the similarity between targets and lures can affect the degree to which a list-length effect occurs (MacAndrew, Klatzky, Fiez, McClelland, & Becker, 2002). In a study examining the effect of phonological similarity on recognition memory, MacAndrew et al. (2002) tested subjects' recognition memory for letters on a study list of four or six letters. The results showed that a larger list-length effect occurred for similar targets and lures than for dissimilar targets and lures. Also, overall recognition scores were higher for dissimilar targets and lures than for similar targets and lures.

The False Memory Effect

A number of experimental studies showed a false memory effect, which holds that the recognition of a lure (i.e., a false memory or a false alarm) is more likely to happen when the lure is similar to (one of the) studied items (e.g., Postman, 1951; Dewhurst & Farrand, 2004). For instance, the results by Dewhurst and Farrand (2004) show that the number of false memories increases together with the number of targets on the study list that are similar to the lures.

In a similarity space, representations of similar targets and lures show more overlap than representations of dissimilar targets and lures. Similar targets and lures are thus more difficult to discriminate than dissimilar targets and lures. Therefore, we expect that list-strength effects and list-length effects will be more pronounced and there will be more false alarms when targets and lures are similar than when targets and lures are dissimilar.

We hypothesize that the similarity structure of the perceived targets and lures can give rise to the recognition-memory effects discussed above. To test this hypothesis, we conducted a face recognition simulation with the NIM model, which employs similarity-space representations of perceived natural images.

Simulation

In our simulation, we investigated the ability of the NIM model to produce the following effects: (1) the effect of the

similarity between targets and lures on the list-strength effect, (2) the effect of the similarity between targets and lures on the list-length effect, and (3) the false memory effect. The NIM model was repeatedly provided with a study list of face images and tested for recognition of the studied images (i.e., targets) and a number of non-studied images (i.e., lures). The images were gray-scale images of human faces taken from the FERET database (Phillips, Wechsler, Huang, & Rauss, 1998) of facial images. Male and female Caucasian faces without beards or glasses were selected. An example of such an image is shown on the left hand side of Figure 1. In this simulation, recognition memory was tested in two different conditions: (1) the dissimilar condition that employed lures dissimilar from the targets, and (2) the similar condition, that employed lures similar to one of the targets. In the NIM model, similar images are separated by a small distance in the similarity space. List-strength effects and list-length effects were assessed in both conditions and compared to determine whether the similarity between targets and lures had affected the degree to which the list effects occurred. Moreover, a comparison of the recognition results in the dissimilar condition and the similar condition revealed whether a false memory effect had occurred. Below we describe the calculation of recognition scores, the paradigms, the conditions, the procedure, and the results.

Calculation of Recognition Scores

The familiarity values, F , were used in a signal detection analysis to determine the recognition scores. The appropriate measure for the recognition score (d_a) was based on the normalized difference between the average F values of the targets ($\bar{F}(I_T)$) and the average F values of the lures ($\bar{F}(I_L)$):

$$d_a = \frac{\bar{F}(I_T) - \bar{F}(I_L)}{\sqrt{\frac{\sigma^2_{[F(I_T)]} + \sigma^2_{[F(I_L)]}}{2}}}$$

(Simpson & Fitter, 1973). Each d_a value was calculated on the basis of the familiarity values for targets ($F(I_T)$) and the familiarity values for lures ($F(I_L)$) of ten recognition tests.

Paradigms

The List-Strength Effect We used the mixed-pure paradigm first proposed by Ratcliff et al. (1990). It is used in many list-strength studies. The mixed-pure paradigm employs three types of study lists: pure weak lists (N weak images), pure strong lists (N strong images), and mixed lists ($N/2$ strong and $N/2$ weak images). A list-strength effect is said to occur (1) when the recognition score for weak images on a pure list is higher than the recognition score for weak images on a mixed list or, (2) when the recognition score for strong images on a mixed list is higher than the recognition score for strong images on a pure list. The pure/mixed ratio for weak images (i.e., the recognition score for weak images on a pure list divided by the recognition score for weak images on a mixed list) thus is an indication for the degree to which a list-strength effect occurs for weak images. Likewise, the mixed/pure ratio for strong images is an indication for the degree to which a list-strength effect occurs for strong images.

The List-Length Effect A list-length effect is said to occur when the recognition score for images on a shorter list is higher than the recognition score for images on a longer list. To assess the occurrence of a list-length effect we compared recognition scores for images on study lists of different lengths.

The False Memory Effect A higher false alarm rate (together with no difference in the hit rate) for the similar condition than for the dissimilar condition is said to indicate the occurrence of a false memory effect. However, using the general performance measure d_a (as described in the previous subsection) to determine recognition scores, the NIM model produces no false memories (and thus no false memory effect), simply because no recognition decisions are made. Most computational memory models, however, make recognition decisions based on the comparison of an obtained familiarity value to a given criterion (e.g., Bussey, 2001). When the familiarity value exceeds the criterion, the item is recognized, if not, the item is not recognized. To assess the false memory effect, we also applied a decision criterion to the familiarity values, F , obtained for the dissimilar condition and for the similar condition. As a criterion we used: $C = S \times (0.02 + N/500)$, with S the storage strength of the images, and N the number of images on the study list.

Conditions

We distinguished two conditions: the dissimilar and the similar condition. For the dissimilar condition, recognition performance for targets versus dissimilar lures was tested. Targets and lures were selected from a subset of dissimilar images. The images in the subset of dissimilar images were selected in such a way that the clusters of their feature vectors in the similarity space showed relatively little overlap. Hence, dissimilarity for a subset of images, D , is defined as: $\sum f_{B,A_i}/T_A \leq d_1, \forall A, B \in D$, with f_{B,A_i} the number of feature vectors of image B that fall within a hypersphere with radius r centered around the i^{th} feature vector of image A , T_A the total number of feature vectors of image A , and d_1 a dissimilarity constant. For the similar condition, recognition performance for targets versus similar lures was tested. Pairs of similar targets and lures were selected in such a way that the clusters of their feature-vector representations in the similarity space showed relatively much overlap. Hence, similarity for two images, A and B , is defined as: $\sum f_{B,A_i}/T_A \geq d_2$, with f_{B,A_i} the number of feature vectors of image B that fall within a hypersphere with radius r centered around the i^{th} feature vector of image A , T_A the total number of feature vectors of image A , and d_2 a similarity constant, with $d_2 > d_1$.

Procedure

We provided the NIM model with (1) pure weak study lists, (2) pure strong study lists, and (3) mixed study lists of lengths $N = 4, 8$, and 12 , in both the dissimilar and the similar condition. Weak images were stored with storage strength $S = 5$ (i.e., five feature vectors were stored, corresponding to five fixations) and strong images were stored with storage strength $S = 10$. For each feature vector, the first $p = 50$ dimensions were stored. After the last image of a study list had been presented to the model, the N images of the study list (i.e., targets) along with N new images (i.e., lures) were presented for




Figure 2: (a) Norman's (2002) results, and (b) Recognition scores (d_a) for weak images on pure lists (black bars) and on mixed lists (white bars) of length $N = 12$ for the dissimilar condition and for the similar condition.

recognition. Lures in the dissimilar condition were selected with dissimilarity constant $d_1 = 0.26$ and lures in the similar condition were selected with similarity constant $d_2 = 0.8$. Recognition tests and the selection of targets and lures were performed using the radius parameter $r = 5.0$.

Results

Table 1 presents the results for the dissimilar and similar conditions, respectively. The rows show the recognition results for lists of lengths $N = 4, 8$, and 12 . The columns labeled w show the recognition scores for the weak images and the columns labeled s show the recognition scores for the strong images. The columns labeled pw/mw show the pure/mixed ratio for weak images and the columns labeled ms/ps show the mixed/pure ratio for strong images (both of which are indications of the degree to which a list-strength effect occurred). Figure 2(a) presents a bar graph of the results re-

Table 1: The average recognition scores produced by the NIM model for the dissimilar and the similar condition.

| Dissimilar condition | | | | | | |
|----------------------|------------|------|-------------|------|---------|---------|
| N | Pure lists | | Mixed Lists | | ratios | |
| | w | s | w | s | pw/mw | ms/ps |
| 4 | 1.81 | 2.39 | 1.78 | 2.41 | 1.01 | 1.01 |
| 8 | 1.65 | 2.18 | 1.54 | 2.28 | 1.07 | 1.05 |
| 12 | 1.59 | 2.11 | 1.48 | 2.15 | 1.07 | 1.02 |
| Similar condition | | | | | | |
| N | Pure lists | | Mixed Lists | | ratios | |
| | w | s | w | s | pw/mw | ms/ps |
| 4 | 1.38 | 1.83 | 1.14 | 1.97 | 1.21 | 1.08 |
| 8 | 1.12 | 1.52 | 0.87 | 1.78 | 1.29 | 1.17 |
| 12 | 0.98 | 1.35 | 0.73 | 1.61 | 1.36 | 1.20 |

ported by Norman (2002) (described previously). Figure 2(b) presents a bar graph of the recognition scores produced by the NIM model in conditions analogous to the conditions in Norman's experiment. Since results were similar for lists of different lengths N , only the results for lists of length $N = 12$ are shown. A comparison of the graphs in Figures 2(a) and 2(b) reveals a close correspondence between Norman's results and the results produced by the NIM model.

Similarity and the List-Strength Effect List-strength effects for the dissimilar condition were significantly smaller than list-strength effects for the similar condition as indicated by the higher pw/mw and ms/mw values for the similar condition compared to those for the dissimilar condition. This was supported in a two-way analysis of variance (ANOVA) by the interaction between list type (pure or mixed) and condition. Calculated F values for lists of lengths $N = 4, 8$, and 12 ranged from $F(1, 159) = 4.97$ to $F(1, 159) = 9.62$, $p < 0.05$ for weak images and $F(1, 159) = 4.52$ to $F(1, 159) = 12.02$, $p < 0.05$, for strong images.

Similarity and the List-Length Effect The list-length effects for the dissimilar condition were significantly smaller than those for the similar condition. This was indicated in a two-way ANOVA by the interaction between list length and condition for pure weak lists, $F(2, 239) = 4.61$, $p < 0.05$, and for pure strong lists, $F(2, 239) = 3.68$, $p < 0.05$.

The False Memory Effect Table 2 presents the hit rates and false alarm rates for pure strong lists of lengths $N = 4, 8$, and 12 for both the dissimilar and the similar condition. Since the results were similar for pure weak lists and pure strong lists, we only present the results for pure strong lists. The

Table 2: The average hit rates and false alarm rates produced by the NIM model for the dissimilar and the similar condition.

| N | Dissimilar condition | | Similar condition | |
|-----|----------------------|----------|-------------------|----------|
| | Hit rate | F/A rate | Hit rate | F/A rate |
| 4 | 0.84 | 0.01 | 0.86 | 0.14 |
| 8 | 0.76 | 0.02 | 0.78 | 0.17 |
| 12 | 0.69 | 0.02 | 0.70 | 0.15 |

results show that a false memory occurred: false alarm rates were higher for lists in the similar condition than for lists in the dissimilar condition (while hit rates were not significantly different). In an ANOVA, calculated F values for the false alarm rates ranged from 163.38 to 384.74, $p < 0.05$, while F values for the hit rates ranged from 2.08 to 2.24, $p > 0.05$.

Discussion

Based on recent experimental findings (Norman, 2002), we assumed that the degree to which a list-strength effect and a list-length effect occur varies with the degree of similarity between targets and lures. The NIM model produces this effect, as well as a false memory effect. Below we discuss the single-process NIM model in relation to other memory models and the ability of the NIM model to simulate mirror effects.

Comparison to Other Models

The NIM model differs from existing memory models in that it operates on natural input and employs a single process for recognition.

A Perceptual Process Operating on Natural Input The NIM model encompasses a transformation that yields the similarity structure of natural images. So far, existing memory models have been tested with artificial data (e.g., the REM

model, Shiffrin & Steyvers, 1997), with predefined similarity spaces (e.g., the SimSample model, Busey, 2001), or with synthesized natural images (Kahana & Sekuler, 2002). The predictions these models make for recognition memory performance can be similar to the predictions the NIM model makes, provided that a representation space is employed that accurately reflects the similarity structure of the input. However, these models fall short in constructing a representation in an *a priori* manner. In contrast, the NIM model remedies this. Therefore, we expect that the NIM model provides us with a useful tool for making predictions about the effects of varying similarity of natural input on memory.

Single versus dual-process models Several memory models assume two processes for recognition to explain recognition results. These dual-processing models assume that recognition involves (1) a familiarity process, i.e., a context insensitive automatic process, and (2) a recollection process, i.e., a context sensitive strategic process (see Yonelinas, 2002, for a review of dual-processing models). Norman (2002) explains his experimental findings on the similarity effect by a dual-processing approach. He argues that the degree to which a list-strength effect occurs depends on the extent to which recollection drives recognition. While there might be good biological evidence that more than one process is involved in recognition, our results show that a single straightforward process for recognition suffices to produce Norman's and other findings on recognition memory.

Mirror Effects

In addition to the list-strength effect and the list-length effect, memory models are often tested for two related effects consistently obtained in experimental studies: the strength-mirror effect and the length-mirror effect (e.g., Murnane & Shiffrin, 1991). Simulation results, reported elsewhere (Lacroix, Murre, Postma, & Herik, submitted), showed that the NIM model exhibits these effects.

Conclusion

We have seen that the NIM model is able to build a similarity space from perceived natural data. Moreover, it successfully replicated recognition findings on list-strength effects, list-length effects, false memory effects, and mirror effects. Though it is at present not clear to what extent these results emerge from the use of natural images, it does increase the validity of the model by by-passing assumptions about distributional statistics of real-world perceptual features. Future studies aim at extending the NIM model to simulate a wider variety of findings on recognition memory.

Acknowledgments

The research project is supported in the framework of the NWO Cognition Program with financial aid from the Netherlands Organization for Scientific Research (NWO). It is part of the larger project: 'Events in memory and environment: modeling and experimentation in humans and robots' (project number: 051.02.2002).

References

Arkadev, A. G., & Braverman, E. M. (1966). *Computers and pattern recognition*. Washington, DC: Thompson.

- Bartlett, M. S., Littleworth, G., Braathen, B., Sejnowski, T. J., & Movellan, J. R. (2003). A prototype for automatic recognition of spontaneous facial actions. In S. Becker & K. Obermayer (Eds.), *Advances in neural information processing systems* (Vol. 15). Cambridge, MA: The MIT Press.
- Bellman, R. (1961). *Adaptive control processes: a guided tour*. Princeton, NJ: Princeton University Press.
- Busey, T. (2001). Formal models of familiarity and memorability in face recognition. In M. Wenger & J. Townsend (Eds.), *Computational, geometric, and process perspectives on facial cognition: Contexts and challenges*. Hillsdale, NJ: Erlbaum Associates.
- Cary, M., & Reder, L. M. (2003). A dual-process account of the list-length and strength-based mirror effects in recognition. *Journal of Memory and Language*, 49, 231-248.
- Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, 108, 452-478.
- Dewhurst, S. A., & Farrand, P. (2004). Investigating the phenomenological characteristics of false recognition for categorised words. *European Journal of Cognitive Psychology*, 16, 403-416.
- Edelman, S., & Intrator, N. (1997). Learning as extraction of low-dimensional representations. In R. Goldstone, D. Medin, & P. Schyns (Eds.), *Mechanisms of perceptual learning* (Vol. 36, pp. 353-380). San Diego, CA: Academic press.
- Freeman, W. T., & Adelson, E. H. (1991). The design and use of steerable filters. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13, 891-906.
- Hubel, D. H. (1988). *Eye, brain, and vision*. New York, NY: WH Freeman.
- Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: A noisy exemplar approach. *Vision Research*, 42, 2177-2192.
- Kalocsai, P., Zhao, W., & Biederman, I. (1998). Face similarity space as perceived by humans and artificial systems. In *Proceedings, third international conference on automatic face and gesture recognition* (pp. 177-180). Nara Japan.
- Lacroix, J. P. W., Murre, J. M. J., Postma, E. O., & Herik, H. J. van den. (submitted). Modeling recognition memory using the similarity structure of natural input. *Psychological Review*.
- Lyons, M., & Akamatsu, S. (1998). Coding facial expressions with gabor wavelets. In *Proceedings, third international conference on automatic face and gesture recognition* (pp. 200-205). Nara Japan.
- MacAndrew, D. K., Klatzky, R. L., Fiez, J. A., McClelland, J. L., & Becker, J. T. (2002). The phonological-similarity effect differentiates between two working memories. *Psychological Science*, 13, 465-468.
- McClelland, J. L., & Chappell, M. (1998). Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. *Psychological Review*, 105, 724-760.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
- Murnane, K., & Shiffrin, R. M. (1991). Interference and the representation of events in memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 17, 855-874.
- Norman, K. A. (2002). Differential effects of list strength on recollection and familiarity. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 28, 1083-1094.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- Nosofsky, R. M., Clark, S. E., & Shin, H. J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 282-304.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: The MIT Press.
- Palmeri, T. J., & Gauthier, I. (2004). Visual object understanding. *Nature Reviews Neuroscience*, 5, 291-303.
- Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science*, 2, 559-572.
- Phillips, P. J., Wechsler, H., Huang, J., & Rauss, P. (1998). The FERET database and evaluation procedure for face recognition algorithms. *Image and Vision Computing Journal*, 16, 295-306.
- Postman, L. (1951). The generalization gradient in recognition memory. *Journal of Experimental Psychology*, 42, 231-235.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93-134.
- Rao, R. P. N., & Ballard, D. H. (1995). An active vision architecture based on iconic representations. *Artificial Intelligence*, 461-505.
- Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). The list-strength effect: I. data and discussion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16, 163-178.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: Rem: Retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145-166.
- Simpson, A. J., & Fitter, M. J. (1973). What is the best index of detectability? *Psychological Bulletin*, 80, 481-488.
- Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (in press). Word association spaces for predicting semantic similarity effects in episodic memory. In A. Healy (Ed.), *Cognitive psychology and its applications: Festschrift in honor of lyle bourne, walter kintsch, and thomas landauer*. Washington, DC: American Psychological Association.
- Tenenbaum, J. B., Silva, V. de, & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290, 2319-2323.
- Yarbus, A. L. (1967). *Eye movements and vision*. New York: Plenum Press.
- Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, 46, 441-517.