Psychonomic Bulletin & Review
2010, 17 (5), 662-672
doi:10.3758/PBR.17.5.662

Evaluating the random representation assumption of lexical semantics in cognitive models

Brendan T. Johns and Michael N. Jones
Indiana University, Bloomington, Indiana

A common assumption implicit in cognitive models is that lexical semantics can be approximated by using randomly generated representations to stand in for word meaning. However, the use of random representations contains the hidden assumption that semantic similarity is symmetrically distributed across randomly selected words or between instances within a semantic category. We evaluated this assumption by computing similarity distributions for randomly selected words from a number of well-known semantic measures and comparing them with the distributions from random representations commonly used in cognitive models. The similarity distributions from all semantic measures were positively skewed compared with the symmetric normal distributions assumed by random representations. We discuss potential consequences that this false assumption may have for conclusions drawn from process models that use random representations.

A model of a cognitive phenomenon typically requires accounts of representation, of process, and of how the two interact (Estes, 1975). These two aspects of a model are interdependent, with the process requiring a representation on which to operate and the representation requiring a process to simulate behavior. For example, Rumelhart and McClelland (1982) created a model in which 16-feature vectors were used to represent capital letters in which each feature was the presence or absence of a line at a particular orientation, and they evaluated the process of interactive activation. The use of this representation was justified by research on how the visual system responds to primitive features. Similarly, some models use realistic representations of faces (e.g., Dailey & Cottrell, 1999) and orthographic and phonological characteristics of words (e.g., Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001) or digits (e.g., Hinton, 2007).

If insufficient research exists to point to the correct representation, a common practice in cognitive modeling is to use randomly generated representations to stand in for psychological structure. This practice makes it unlikely that the representation is biased toward supporting the process mechanism, and the model can be refined later when further research reveals the correct representation. An example is Hintzman's (1986) use of random representations to simulate schema abstraction using Posner and Keele's (1968) stimuli: Stimuli were random dot patterns, and exemplars of the same category were random perturbations of a prototype pattern. Hintzman (1986) was able to create equivalent structure in his simulation by generating prototypes as random vectors and generating exemplars within a category as distortions of a prototype.

However, caution is needed if random representations are used. The performance of cognitive models is dependent largely on valid representational assumptions. For example, Daugherty and Seidenberg's (1992) connectionist model was only able to correctly simulate past tense verb processing when the model was trained on representations that contained the correct distributional structure. Alternatively, it is also possible for a process model to give a good account of the human behavior *only* when random representations are used, but not when the correct representational structure is encoded. Cree, McRae, and McNorgan (1999) argued that using a plausible representational structure, rather than random representations, constrains the modeling exercise by reducing degrees of freedom.

Random representations are commonly used in models of episodic memory. In global matching models of recognition memory (e.g., Hintzman, 1988; Murdock, 1982; Shiffrin & Steyvers, 1997), decisions are made by assessing the similarity of the probe word to the (usually noisy) study items with particular processing mechanisms. The use of random representations in these models produces a hidden assumption that the distribution of similarity across randomly selected words is symmetric and is approximately Gaussian. The distributional assumption comes from the design of a typical memory experiment in which random words are used. Because words are selected randomly, they are assumed to have only random similarity on dimensions extraneous to the experimental manipulation (e.g., orthography, phonology, and semantics); however, this assumption is unlikely to be valid. Hence, it is common to explicitly control confounding factors, such as frequency. In this examination, we focus on lexical

B. T. Johns, johns4@indiana.edu

semantics—a factor often ignored, because it is difficult to quantify and control. In assuming that two randomly selected words have only a random expected semantic similarity, random representations seem appropriate.

However, the use of random representations in many cognitive models assumes that semantic similarity is symmetrically distributed across all sampled words. We demonstrate that this is unlikely to be the case with actual words used in experiments and that this may produce serious consequences for conclusions drawn from process models that use random representations. Multiple techniques have recently emerged that allow convergent tests of the distributional characteristics of semantic similarity to test the assumptions made by models that use random representations.

ANALYSIS

Comparison distributions are needed for evaluating the assumption of random similarity. Our analysis uses three types of semantic similarity measures to create distributions: measures computed from free-association data, a hand-coded lexical database (WordNet), and corpus-based co-occurrence models. These semantic similarity measures are compared against common methods for constructing random representations.

Semantic Measures

Word association space (WAS). Steyvers, Shiffrin, and Nelson (2004) applied singular value decomposition (SVD) to the free-association data from Nelson, McEvoy, and Schreiber's (1998) norms. Steyvers et al. demonstrated that semantic similarity measures from the resulting reduced vectors offer a good predictor of similarity effects in recognition and recall.

WordNet similarity. WordNet (Miller, 1990) is a hand-coded lexical network in which nodes contain one or more synonymous words and are linked together via lexical relationships (e.g., hypernymy and holonymy). A variety of methods have been proposed to compute similarity from the network—we use the Jiang–Conrath distance measure (JCN) here, because it maps best onto human similarity ratings (Maki, McKinley, & Thompson, 2004). JCN is a distance metric that basically counts the number of nodes and edges between two concepts in the network.

Latent semantic analysis (LSA). This method (and those that follow) differs from the WAS of Steyvers et al. (2004), in that it does not use data from human experiments to infer a semantic representation but, instead, uses statistical regularities computed from a large text corpus. In LSA (Landauer & Dumais, 1997), a word $\times$ document matrix is decomposed using SVD so that each word is represented by a vector containing the ~300 dimensions with the largest eigenvalues. LSA representations have seen success in accounting for a variety of semantic behaviors.

BEAGLE. In Jones and Mewhort's (2007) BEAGLE model, a representation of a word is built through experience with a text corpus. Words are initially represented by random Gaussian vectors, and a word's semantic representation is created by summing and convolving words that occur in sentences with a target word. The use of convolution allows order information to be included (the sentential position of the word relative to other words), as well as the co-occurrence information in LSA. BEAGLE has successfully accounted for a variety of semantic behaviors, including semantic typicality, categorization, and sentence completion (Jones & Mewhort, 2007), as well as a range of semantic priming data (Jones, Kintsch, & Mewhort, 2006).

The COALS model. COALS (Rohde, Gonnerman, & Plaut, 2005) uses a word $\times$ word co-occurrence matrix, but uses correlations instead of raw frequency. This matrix is subsequently reduced in dimensionality with SVD similarly to LSA. Across a wide variety of tasks, Rohde et al. demonstrated that COALS outperforms LSA, WordNet, and thesaurus-based distance measures.

Pointwise mutual information (PMI). PMI (Recchia & Jones, 2009) is a pure co-occurrence metric, computed as the probability of observing two words together divided by the probability of observing each independently. Recchia and Jones computed PMI values over a large corpus of Wikipedia articles (~400,000 articles) and found that PMI produced a significantly better fit to human rating data than did LSA or other similarity metrics.

Random Representations

To compare the distributions created by the semantic measures, we explored five common types of random vectors that have been used to represent semantics in influential cognitive models.

Gaussian vectors. A word's representation is created by randomly sampling vector elements from a Gaussian distribution, $N(\mu, \sigma)$. This type of representation has been used in a variety of models of recognition (e.g., Murdock, 1982). In the following analysis, vectors are created as in Murdock, with a vector dimensionality of 250 ($\mu = 0$ and $\sigma = \sqrt{1/250}$).

Gamma vectors. A word vector is created by sampling integers from a gamma distribution. This type of representation is used in the retrieving effectively from memory (REM) model of recognition memory (Shiffrin & Steyvers, 1997) and in related models, such as the REM lexical decision model (REM–LD; Wagenmakers et al., 2004), that are designed to explain lexical decision. We constructed these vectors as is specified in Shiffrin and Steyvers, with a length of 20 and with $g = 0.45$.

MINERVA vectors. In the influential MINERVA 2 model of memory (Hintzman, 1986, 1988), vector elements are selected randomly from the set $\{-1, 0, 1\}$. A value of 1 represents a positive link between the word and that feature, -1 represents an inhibitory link, and 0 is defined as either irrelevant or unknown for that particular word and feature. Vectors were constructed with a length of 20, as in Hintzman (1988).

Sparse binary vectors. In this type of representation, the majority of entries are zeros, with some entries set to one at random locations. For instance, in Plaut (1995), elements in a word's semantic representation had a 10%

probability of being nonzero. Sparse binary vectors have been used to model lexical priming (Plaut, 1995) and recognition memory (Dennis & Humphreys, 2001; Norman & O'Reilly, 2003), among other domains. Similar to Plaut's simulations, we generated vectors with a length of 100, with each element having a 10% probability of being nonzero. In addition, we tested binomial distributions (sparseness of 50%) to examine the effect of sparseness on the similarity distributions.

Dichotomous vectors. A common representation used in connectionist modeling is a random vector composed equally of 1 or −1 elements. These are similar to MINERVA vectors, but without zero-valued elements. Dichotomous vectors have been used in a variety of models, such as connectionist models of semantic priming (e.g., Masson, 1995). Vectors with a length of 100 were used in the following simulations.

METHOD

To calculate similarity distributions using the semantic measures, we selected 1,000 words from the Toronto word pool (Friendly, Franklin, Hoffman, & Rubin, 1982) and computed the similarity between each pair of words. Next, to examine the similarity distribution under each representation, we randomly sampled 50,000 of these semantic comparison values. For WAS, LSA, BEAGLE, and COALS representations, the similarity metric used was a vector cosine. To compute PMI values, we used a tool made available by Recchia and Jones (2009) and computed PMI values for the word pairs across 400,000 documents from Wikipedia.

For the randomly generated representations, we created a distribution of 100,000 similarity comparisons for each representation type. The distribution was constructed by randomly generating two vectors from the given representation type and computing the similarity between them. Similarity was measured with a vector cosine for all random representations.

To evaluate distribution shape, we used two methods of assessing normality: skewness and normal quantile–quantile (Q–Q) plots. *Skewness* is the third moment about the mean, and it signals asymmetry in a distribution. Q–Q plots are used to assess the standardized difference between an observed distribution and a theoretical (in this case, Gaussian) distribution.

RESULTS

The skewness values for the similarity distributions of both the semantic spaces and random representations are plotted in Figure 1. As the figure shows, all the semantic spaces create positively skewed similarity distributions. That is, there tends to be a greater number of low-similarity scores and a small number of high-similarity scores in a given distribution of randomly selected words. Co-occurrence models (LSA, BEAGLE, and COALS) produced the lowest skew (from 1.06 for BEAGLE to 2.01 for COALS). The PMI distribution produced the largest skew, likely due to the fact that the metric is a pure co-occurrence count with a broader possible range. In the middle was the JCN measure, with a skewness of 2.61, and the WAS of Steyvers et al. (2004), with a skewness of 8.04.

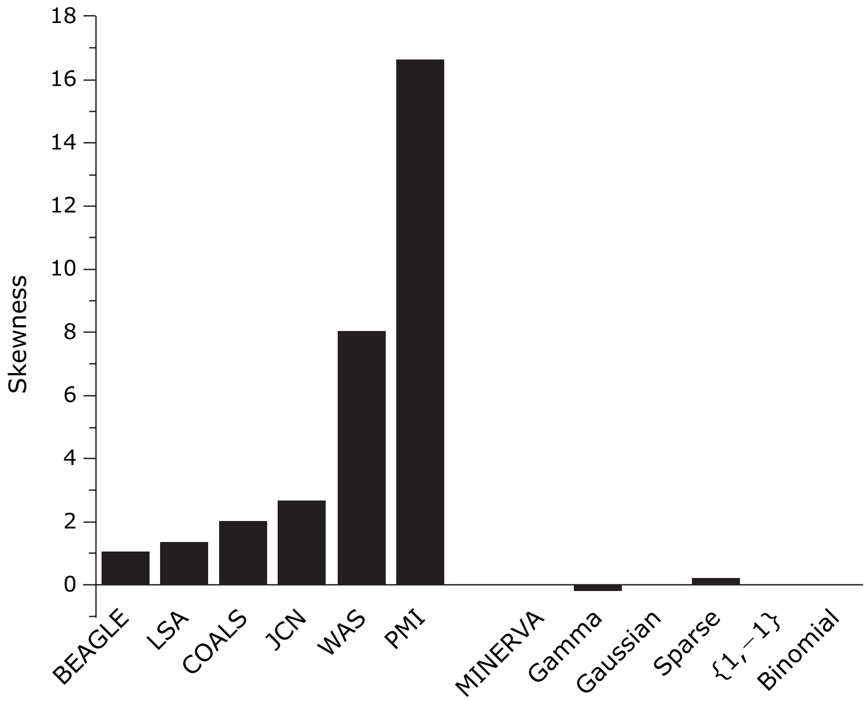

Figure 1. Levels of skewness for the semantic metric distributions and the random representation distributions. BEAGLE, bound encoding of the aggregate language environment; LSA, latent semantic analysis; COALS, correlated occurrence analogue to lexical semantic; JCN, Jiang–Conrath distance measure; WAS, word association space; PMI, pointwise mutual information.

In contrast, all of the random representations produced skewness values of essentially zero (this is expected by their construction). The only distribution that was mildly positively skewed was the sparse binomial distribution, with a skewness of 0.21, whereas the Gamma distribution was mildly negatively skewed, with a value of -0.17. This simple analysis demonstrates that the similarity distributions created by semantic metrics and randomly generated representations are considerably different. Two randomly selected words are likely to be less similar (relative to the other values in the distribution) for semantic models than for random representations.

The Q–Q plots are displayed in Figure 2 for the semantic space distributions and in Figure 3 for the distributions computed from the random representations. Again, the semantic distributions display positive skew, with all of the models having fewer than expected high-similarity values. In four of the distributions (WAS, JCN, LSA, and BEAGLE), there also tend to be fewer than expected low-similarity values. Again, the random representation distributions produce very different results—there is little deviation from normality.

As a final method of visualizing the data, we plot the probability of sampling a similarity value on an equated scale for the different semantic and random representation distributions. To make the different distributions easily comparable, we normalized each distribution with a linear transformation within the range of 0–1. The individual probability distributions for both the semantic space and the random representation distributions are displayed in the top panel of Figure 4, whereas the average distributions for each type (with a polynomial curve fit) are displayed in the bottom panel. This figure depicts what was implied by the two previous analyses: Randomly generated representations tend to be approximately normally distributed, whereas the semantic space models are positively skewed.

DEMONSTRATIONS

To demonstrate how the false assumptions made by random representations might affect a process mechanism, we conducted two brief simulations using recognition memory data—the first using classic signal detection and the second using a particular process model to fit data from a false-memory experiment. Both simulations clearly demonstrate that the use of random representations can allow a process model an unnecessary degree of freedom for fitting human data. However, when the representation is fixed to reflect plausible semantic similarity structure, a process model can have considerably more difficulty reproducing the correct human behavior. This pattern indicates that a process mechanism is incorrect if it relies on an incorrect structural representation to simulate the human data.

Demonstration 1: Signal Detection in Recognition Memory

In this demonstration, we explore signal detection as a recognition process by simulating the task of encoding a list and later making *old/new* judgments for probe words. The simulation varies the type of representation used for words—random versus semantic. Recognition is then simulated by using signal detection as a process model and by fitting an optimum criterion to the human data. The simulation demonstrates the effect that skewed similarity distributions have on a signal-detection decision mechanism.

To fairly compare the random and semantic similarity distributions, we equated them to a normalized scale. Similarity distributions from each of the semantic metric and random representations were normalized within a range of 0 and .5 and were centered on a mean of .25: This procedure allows us to evaluate the shape of the distribution, while centering the distributions on the same mean and within the same range.

Evidence distributions for new and old items were simulated for lists of 20 words. The evidence for a probe was the similarity of the probe to the 20 items on the list. For "new" probes, this evidence was simply the mean of 20 randomly sampled similarity values (because new probes are randomly similar to the contents of memory). For "old" probes, this evidence was the average of the similarity of the item to itself and to the other items on the list (simulated as the mean of 19 randomly sampled similarities and the value of 1, representing the similarity of the word to itself). This similarity sampling process is essentially equivalent to how Murdock (1982) simulated recognition using composite vectors following his theory of distributed associative memory (TODAM). The sampling process was repeated 50,000 times for each similarity distribution.

To compare the resulting evidence values, we calculated the discriminability (measured with d') for each simulation (d' is a measure of how distinct studied items are from nonstudied items). Figure 5 displays the d' values for the different similarity distributions compared with the d' from a simple recognition experiment that used a list length of 20 (Dennis, Lee, & Kinnel, 2008). As the figure shows, all of the semantic distributions have higher d' than do the random distributions. In addition, the d' values for the random representations are much closer to the behavioral data from Dennis et al. The difference in magnitude demonstrated for d' values for semantic and random similarity was statistically reliable [$t(11) = 4.75$, $p < .001$]. To evaluate the effect of skew in the similarity distributions on the resulting d' values, we computed the partial correlation between d' and skewness (controlling for kurtosis and variance) for the distributions, which yielded robust results ($r = .913$, $p < .001$).

The skewness of the similarity distribution has a large effect on the calculation of evidence distributions, because the probability of sampling lower similarity values is much greater than it is in a symmetric distribution. Hence, with "true" semantic representations, an old item tends to be more distinct from other random items on the list, producing a greater difference between old and new evidence distributions. This demonstration is certainly not

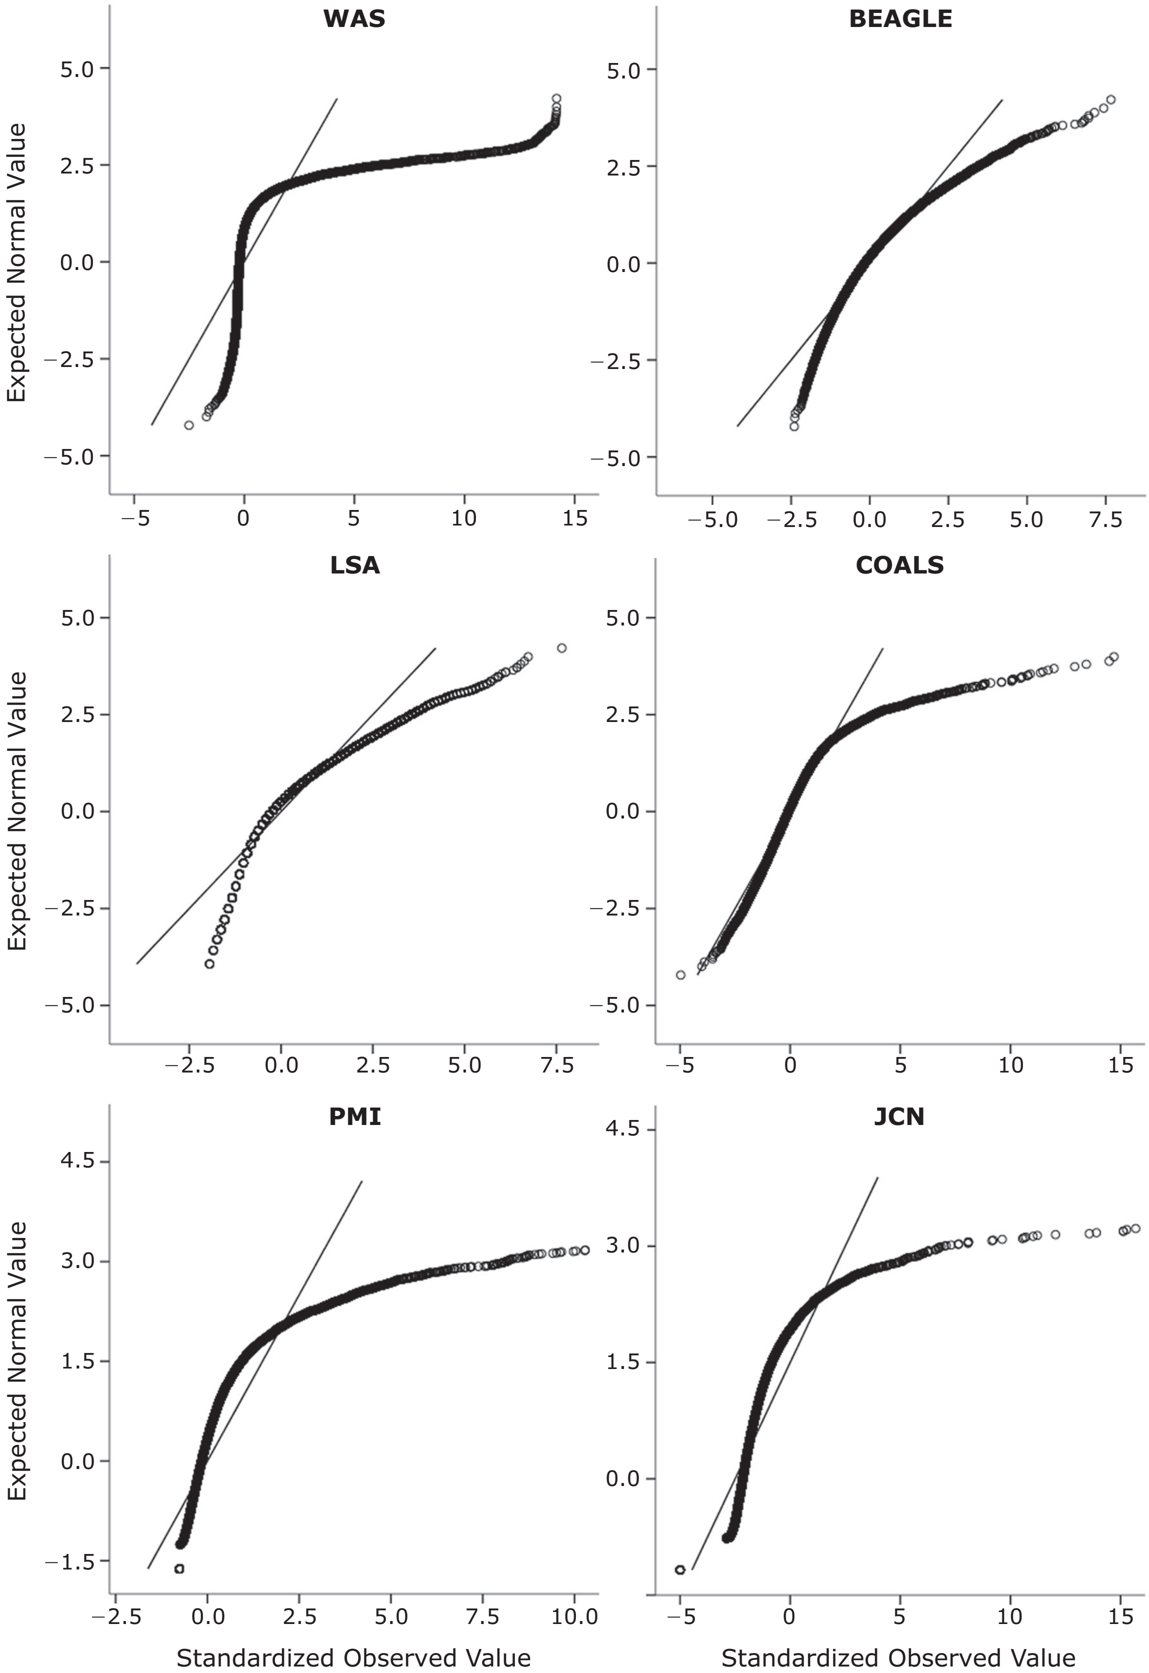

Figure 2. Normal Q–Q plots for the semantic metrics.

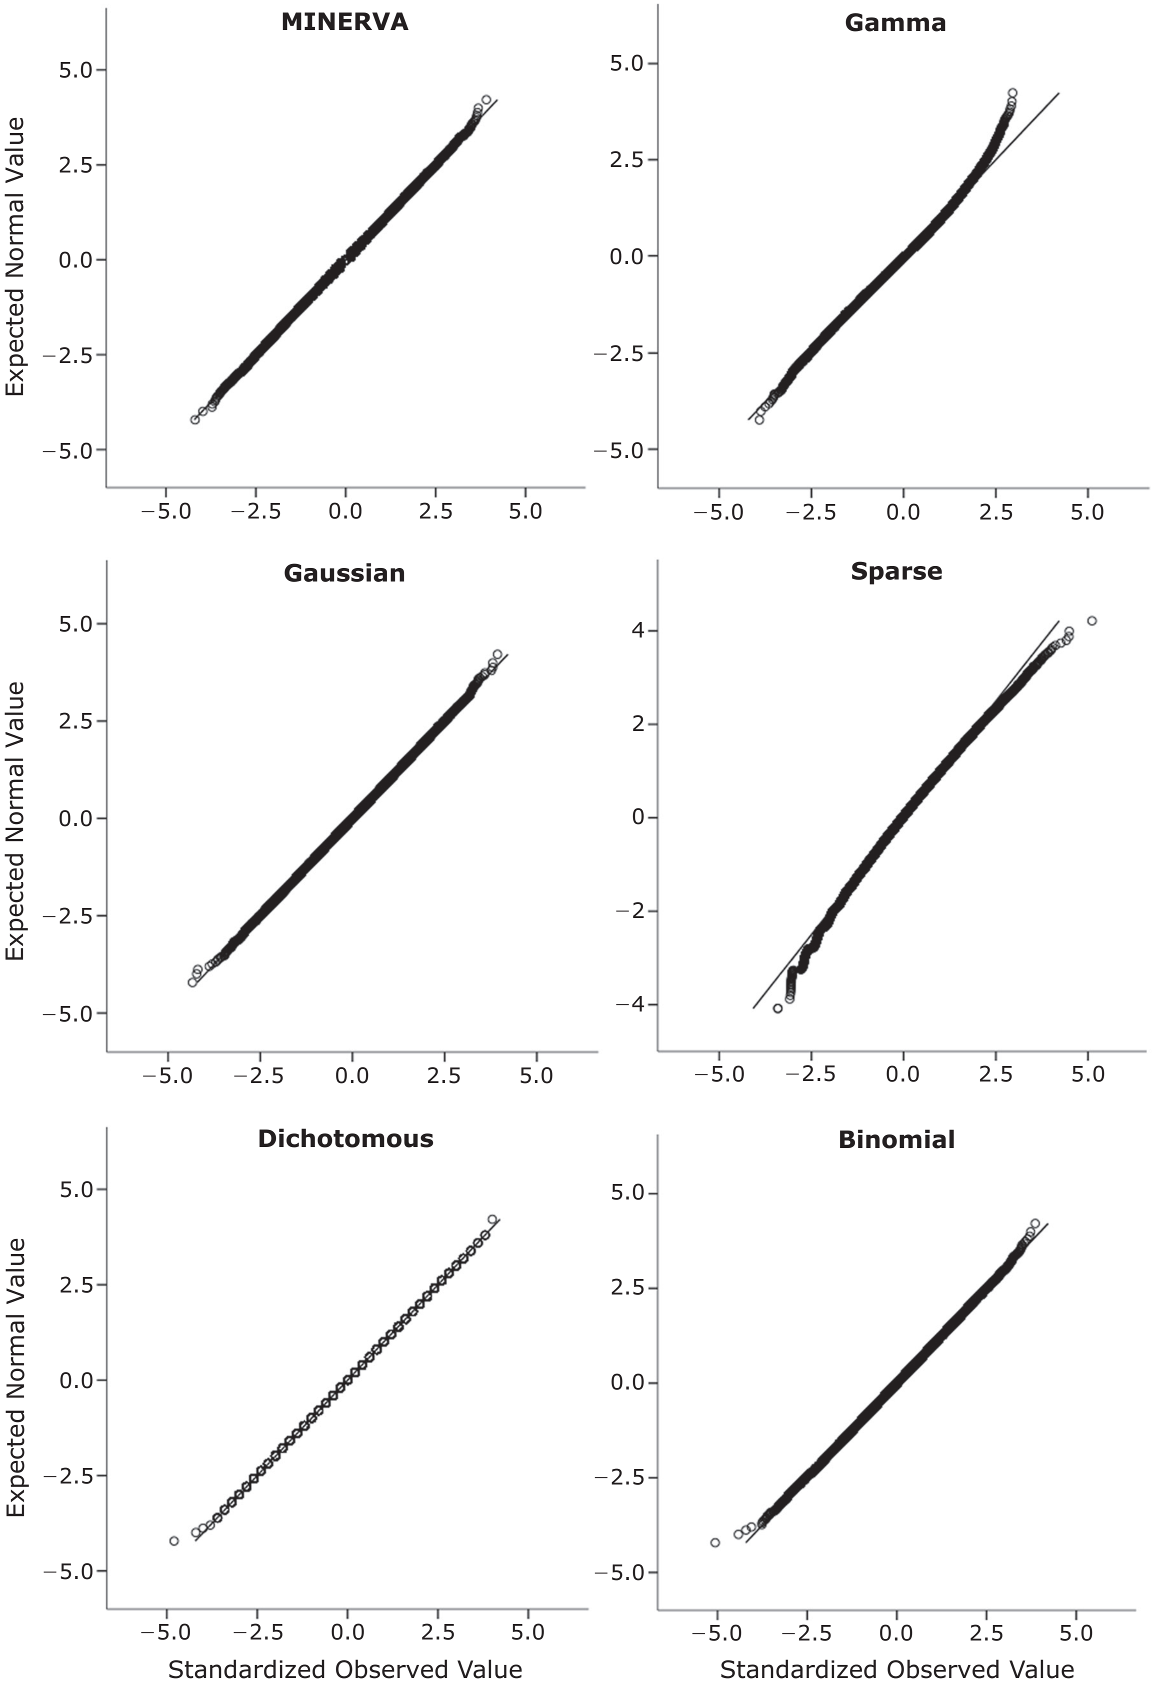

Figure 3. Normal Q–Q plots for the random representations.

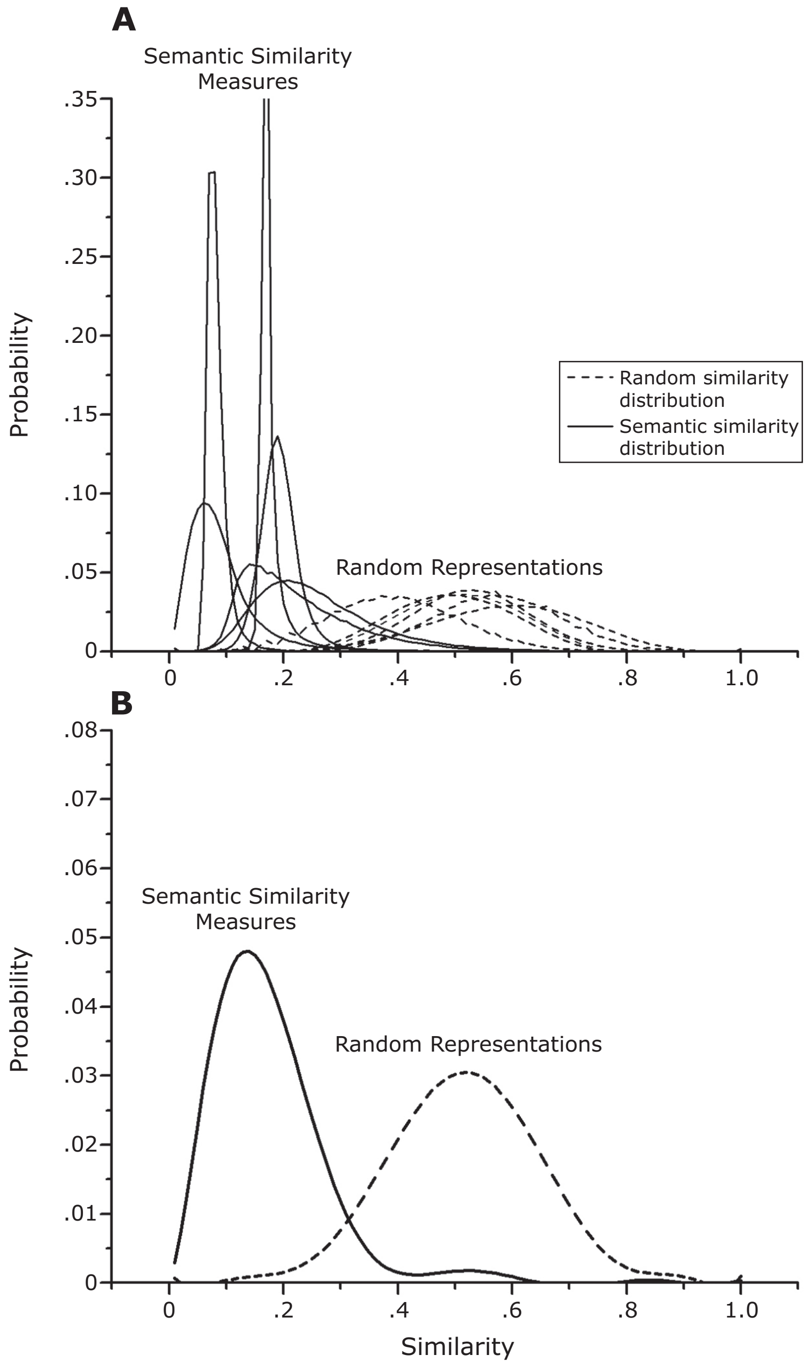

Figure 4. Normalized probability distributions for the individual distributions (A) and the averaged distributions (B).

meant as a refutation of signal detection theory, but instead demonstrates that using realistic representations of semantics imposes significant constraint on a processing model's ability to simulate data.

Demonstration 2: MINERVA and False Memory

The MINERVA 2 model of Hintzman (1986) has been used to successfully account for a variety of categorical false-recognition effects (Arndt & Hirshman, 1998). Here, we simulate associative false recognition with the model, using both random and structured representations of semantics. Robinson and Roediger (1997) found that, as the number of studied items that are related to a critical lure is increased, so is the probability of falsely recognizing that critical lure. The purpose of this demonstration is to compare the ease with which a simple process model like MINERVA is able to model this effect when it is using random representations versus when it is using representa-

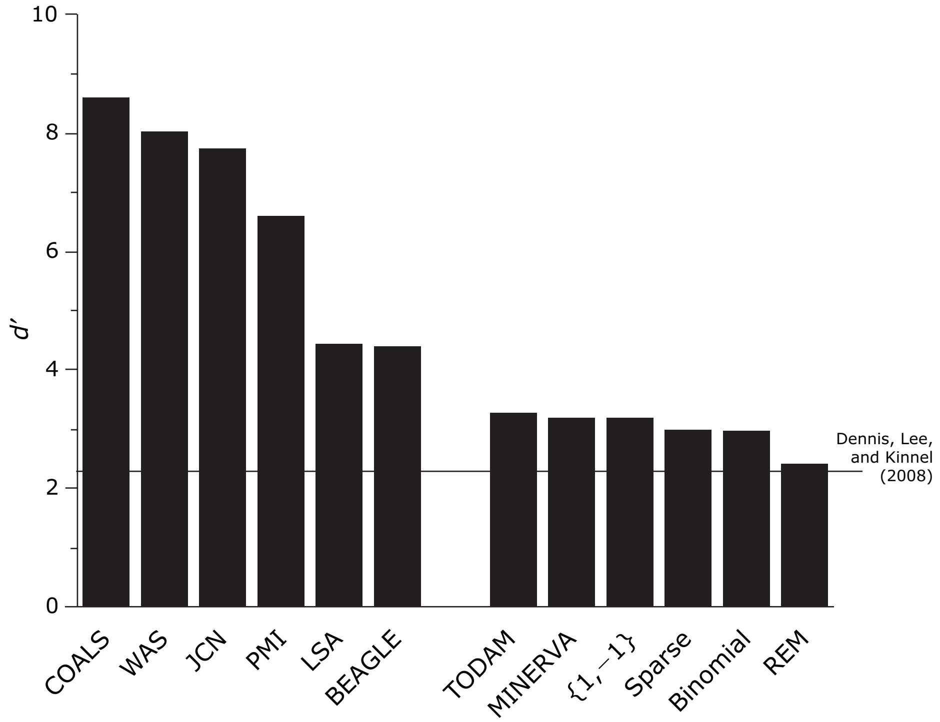

Figure 5. Levels of discriminability (d') for signal detection theory simulations for each representation type, along with behavioral data from Dennis et al. (2008).

tions that contain knowledge about the similarity structure of the actual words.

To construct MINERVA vectors that contain plausible semantic structure, we transformed the WAS representations from Steyvers et al. (2004). Typical applications of MINERVA use ternary vectors with a fairly low dimensionality. Hence, WAS vectors were collapsed from 400 to 20 dimensions by summing every 20 successive elements in a WAS vector into a single element in the reduced vector. This reduced vector was then transformed into a ternary vector, with values of the set $\{-1, 0, 1\}$; the magnitude of the summed WAS values was recoded so that the highest third were assigned $+1$ (representing a high weighting on that feature), the middle third were assigned 0, and the lowest third were assigned -1. To ensure that the MINERVA-transformed vectors still reflected the semantic structure in the original WAS vectors, we computed the word × word cosines between vectors in both representations and correlated the two matrices: The original vectors and their ternary transformed versions were highly correlated ($r = .67$, $p < .001$), indicating that the transformed vectors contain an arrangement of elements that reflects the semantic structure in the original WAS vectors. Using the false-recognition lists from Stadler, Roediger, and McDermott (1999) and Gallo and Roediger (2002), there was a high average similarity of the critical word's representation to the representations of the list items across the 52 word lists ($r = .35$, $p < .001$).

Random representations for critical words and their corresponding lists were created as in Arndt and Hirshman (1998), by using prototype and exemplar vectors. A prototype vector (representing the critical word) is first generated by randomly sampling elements from the set $\{1, 0, -1\}$ with equal probability. Each item in the word list is then created by randomly perturbing elements in the prototype vector. This process requires a distortion parameter, which determines the probability of switching elements from the prototype vector when creating a list item vector. The distortion parameter determines how similar the list items are to the critical word. The important point is that both the semantic and random representations contain the same elements (same quantities of -1s, of 0s, and of 1s). The difference is that the elements are arranged independently for the random representations, whereas they are arranged to respect the interword similarity structure from WAS in the semantic version.

For MINERVA with a semantic representation, the results of Robinson and Roediger (1997) were modeled by randomly selecting three word lists and by adding 3, 6, or 9 items from one of the lists into a study list. Because the word lists in Robinson and Roediger were longer (they also used 12 and 15 associates), 27 randomly selected words from the Toronto word pool were added into the study list. To simulate this with MINERVA using random representation, we created 3, 6, or 9 exemplars for three random prototypes and added them to the study list. We also added 27 random vectors into the study list to make the two simulations equivalent. Decisions are based on activation levels of a probe to the studied items (echo intensity; Hintzman, 1986), calculated by summing the similarity across all items in the study list.[1]

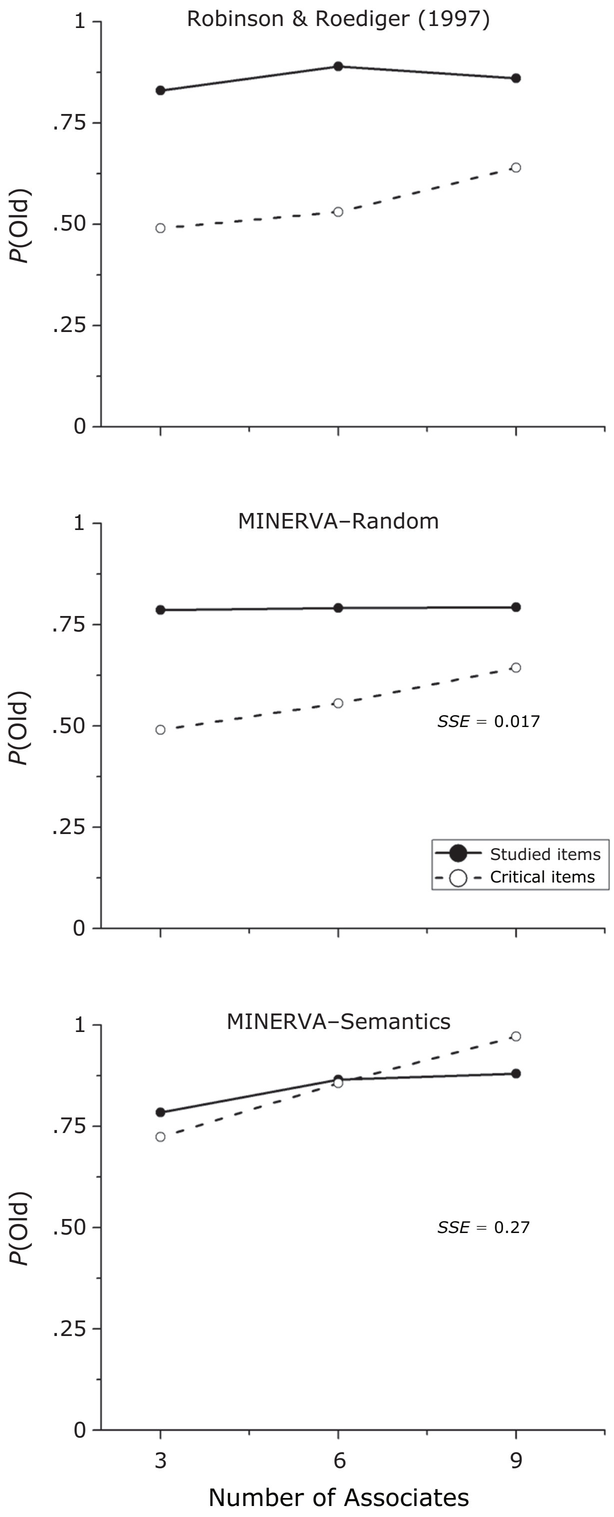

Figure 6. Model fits using MINERVA to simulate the false-recognition results of Robinson and Roediger (1997) with random versus realistic representations of lexical semantics.

For MINERVA with semantic representations, there are two free parameters: a criterion for making an *old/new* decision on the basis of activation levels and a forgetting parameter that determines the probability of a nonzero element switching to zero during study. The simulation with random representations includes an additional distortion parameter (described above) to create the semantic structure. These parameters were fit to the data from Robinson and Roediger (1997) using a Nelder–Mead simplex algorithm. The results of the simulation are displayed in Figure 6: The MINERVA model that uses random representations was able to reproduce the overall trend in the data. However, this was not the case with the MINERVA model that used semantic representations—this model tended to falsely recognize critical items over studied items, which is not the case with the human data. The random representation version of the model produced an excellent account of the data ($R^2 = .98$, $p < .001$). However, the version based on the true semantic similarity of the words used fit no better than chance ($R^2 = .05$, $p = .45$).

This simulation provides a simple demonstration of how a process model that has false representation assumptions may be incorrectly accepted as a plausible model. The only difference between the two models is in their representation structure—the process is identical. Whereas the semantic version contains the "true" semantic structure for the exact words used in the experiment, the random version uses the distortion parameter to create the semantic structure that is most likely if this process account is correct. It is exclusively the incorrect inferred semantic structure that allows the process account to fit these data. If the correct representational structure were used, the process account would be rejected.

DISCUSSION

The use of randomly generated representations contains the assumption that semantic similarity is normally distributed over randomly selected pairs of words.[2] This assumption was shown to be false across many different semantic metrics that have demonstrated success at accounting for human data. In experiments using words, two randomly selected words are likely to be relatively less similar (compared with the distribution of all possible pairs) than is implied by using random representations. Because similarity plays a central role in the processing mechanisms used by many cognitive models, the use of random representations may have consequences for conclusions drawn from simulations using these models. As McClelland (2009) has noted, "simplification is essential, but it comes at a cost, and real understanding depends in part on understanding the effects of simplification" (p. 18). The use of random representations in the development of cognitive models has been a necessary simplification for our understanding of cognitive processes. In making this simplification, researchers have made use of representations whose assumptions may not be entirely accurate, but, through the use of this simplification, modelers have made fundamental discoveries about how memory processes work. It has been only within the last decade that researchers have had access to realistic representations of lexical semantics. The task for the future is to integrate semantic representations with a cognitive process model for a fuller understanding of how they work together to produce observable behavior.

Recent models have begun to explore this type of integration. For example, Monaco, Abbott, and Kahana (2007) have created a neural network model of the mirror effect of frequency, using lexical semantic representations taken from the WAS of Steyvers et al. (2004). Kimball, Smith, and Kahana (2007) have developed a model of false recall that uses the association strengths taken from WAS. Ideally, future models will combine a learning process that builds a representation through exposure to environmental information that can then feed into a processing mechanism. For example, Johns and Jones (2009) have used representations built through a co-occurrence learning process to drive a processing model of both false recognition and false recall. Similarly, Howard, Shankar, and Jagadisan (in press) have recently used the temporal context model to build semantic representations from text; these representations can naturally feed back into the same mechanisms that were used to build them to simulate memory retrieval processes. These models suggest that it is no longer necessary to assume random representations for lexical semantics when modeling cognitive phenomena, but that item-specific semantic representations are now available and offer additional modeling constraints about the structure of semantic similarity on which a process mechanism must operate to produce behavior in a given task.

AUTHOR NOTE

The first author was supported by a postgraduate scholarship from NSERC. We thank Mark Steyvers, Rich Shiffrin, Doug Mewhort, and Simon Dennis for their comments on the manuscript. Address correspondence to B. T. Johns, Department of Psychological and Brain Sciences, 1101 E. 10th Street, Indiana University, Bloomington, IN 47404 (e-mail: johns4@indiana.edu).

REFERENCES

Arndt, J., & Hirshman, E. (1998). True and false recognition in MINERVA2: Explanations from a global matching perspective. *Journal of Memory & Language*, **39**, 371-391.

Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, **108**, 204-256.

Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, **23**, 371-414.

Dailey, M. N., & Cottrell, G. W. (1999). Organization of face and object recognition in modular neural network models. *Neural Networks*, **12**, 1053-1073.

Daugherty, K., & Seidenberg, M. S. (1992). Rules or connections? The past tense revisited. In *Proceedings of the 14th Annual Conference of the Cognitive Science Society* (pp. 259-264). Hillsdale, NJ: Erlbaum.

Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, **108**, 452-478.

Dennis, S., Lee, M. D., & Kinnel, A. (2008). Bayesian analysis of recognition memory: The case of the list-length effect. *Journal of Memory & Language*, **59**, 361-376.

Estes, W. K. (1975). Some targets for mathematical psychology. *Journal of Mathematical Psychology*, **12**, 263-282.

Friendly, M., Franklin, P. E., Hoffman, D., & Rubin, D. C. (1982). The Toronto Word Pool: Norms for imagery, concreteness, orthographic variables, and grammatical usage for 1,080 words. *Behavior Research Methods & Instrumentation*, **14**, 375-399.

Gallo, D. A., & Roediger, H. L., III (2002). Variability among word lists in eliciting memory illusions: Evidence for associative activation and monitoring. *Journal of Memory & Language*, **47**, 469-497.

Hinton, G. E. (2007). Learning multiple layers of representation. *Trends in Cognitive Sciences*, **11**, 428-434. doi:10.1016/j.tics.2007.09.004

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

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, **95**, 528-551.

Howard, M. W., Shankar, K. H., & Jagadisan, U. K. K. (in press). Constructing semantic representations from a gradually changing representation of temporal context. *Topics in Cognitive Science.*

Johns, B. T., & Jones, M. N. (2009). Simulating false recall as an integration of semantic search and recognition. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2511-2516). Austin, TX: Cognitive Science Society.

Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory & Language*, **55**, 534-552.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, **114**, 1-37.

Kimball, D. R., Smith, T. A., & Kahana, M. J. (2007). The fSAM model of false recall. *Psychological Review*, **114**, 954-993.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*, **104**, 211-240.

Maki, W. S., McKinley, L. N., & Thompson, A. G. (2004). Semantic distance norms computed from an electronic dictionary (WordNet). *Behavior Research Methods, Instruments, & Computers*, **36**, 421-431.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 3-23.

McClelland, J. L. (2009). The place of modeling in cognitive science. *Topics in Cognitive Science*, **1**, 11-38.

Miller, G. A. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography*, **3**, 225-312.

Monaco, J. D., Abbott, L. F., & Kahana, M. J. (2007). Lexico-semantic structure and the word-frequency effect in recognition memory. *Learning & Memory*, **14**, 204-213.

Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, **89**, 609-626.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms. Retrieved June 18, 2010, from www.usf.edu/FreeAssociation.

Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review*, **110**, 611-646. doi:10.1037/0033-295X.110.4.611

Plaut, D. C. (1995). Semantic and associative priming in a distributed attractor network. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 37-42). Hillsdale, NJ: Erlbaum.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, **77**, 353-363.

Recchia, G. L., & Jones, M. N. (2009). More data trumps smarter algorithms: Comparing pointwise mutual information to latent semantic analysis. *Behavior Research Methods*, **41**, 647-656.

Robinson, K. J., & Roediger, H. L., III (1997). Associative processes in false recall and recognition. *Psychological Science*, **8**, 231-237.

Rohde, D. L. T., Gonnerman, L., & Plaut, D. C. (2005). *An improved model of semantic similarity based on lexical co-occurrence.* Unpublished manuscript. Retrieved March 18, 2010, from www.cnbc.cmu.edu/~plaut/papers/abstracts/RohdeGonnermanPlautSUB-CogSci.COALS.html.

Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: Part 2. The context enhancement effect and some tests and extensions of the model. *Psychological Review*, **89**, 60-94.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—Retrieving effectively from memory. *Psychonomic Bulletin & Review*, **4**, 145-166.

Stadler, M. A., Roediger, H. L., III, & McDermott, K. B. (1999). Norms for word lists that create false memories. *Memory & Cognition*, **27**, 494-500.

Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications* (pp. 237-249). Washington, DC: American Psychological Association.

Wagenmakers, E.-J., Steyvers, M., Raaijmakers, J. G. W., Shiffrin, R. M., van Rijn, H., & Zeelenberg, R. (2004). A model for evidence accumulation in the lexical decision task. *Cognitive Psychology*, **48**, 332-367.

NOTES

1. Similarity was cubed, as in Hintzman's (1986) version of the model.

2. It is important to note that these results may depend on the similarity metric used. Here we use a vector cosine to compare all representations, but some models use nonlinear similarity metrics (e.g., Shiffrin & Steyvers, 1997), which may produce different behavior.

(Manuscript received November 6, 2009;
revision accepted for publication April 29, 2010.)