

Research Article

Longitudinal Analysis of Early Semantic Networks

Preferential Attachment or Preferential Acquisition?

Thomas T. Hills,¹ Mounir Maouene,² Josita Maouene,³ Adam Sheya,³ and Linda Smith³

¹Cognitive and Decision Sciences, University of Basel; ²UFR: Artificial Intelligence and Bioinformatics, ENSAT, Tangier, Morocco; and ³Department of Psychological and Brain Sciences, Indiana University

ABSTRACT—Analyses of adult semantic networks suggest a learning mechanism involving preferential attachment: A word is more likely to enter the lexicon the more connected the known words to which it is related. We introduce and test two alternative growth principles: preferential acquisition—words enter the lexicon not because they are related to well-connected words, but because they connect well to other words in the learning environment—and the lure of the associates—new words are favored in proportion to their connections with known words. We tested these alternative principles using longitudinal analyses of developing networks of 130 nouns children learn prior to the age of 30 months. We tested both networks with links between words represented by features and networks with links represented by associations. The feature networks did not predict age of acquisition using any growth model. The associative networks grew by preferential acquisition, with the best model incorporating word frequency, number of phonological neighbors, and connectedness of the new word to words in the learning environment, as operationalized by connectedness to words typically acquired by the age of 30 months.

Early noun learning is slow at first and becomes fast (Bloom, 2000; Dale & Fenson, 1996), a fact that suggests that already-learned nouns may help in learning new nouns. Analyses of adult semantic networks support such a more-gets-more growth pattern, revealing structure consistent with growth via preferential attachment. According to the principle of preferential attachment, more highly connected words at Time 1 are more

likely to receive new links at Time 2. This article provides evidence that the developmental growth of early noun networks does *not* accord with the principle of preferential attachment, but rather accords with a principle of preferential acquisition. Preferential acquisition emphasizes not the semantic connectedness of the nouns already known—as in the case of preferential attachment—but the semantic connectedness of the nouns in the learning environment.

With the advance of graph theory and its application to cognitive representations (see, e.g., Dorogovtsev & Mendes, 2001; Steyvers & Tenenbaum, 2005; Vitevitch, 2008), there has been increasing study of large-scale semantic networks. In these networks, words (e.g., nouns) are nodes that are connected to each other via links that indicate semantic relatedness. Researchers have used a variety of indices of semantic relatedness, including the associates generated by adults in free-association studies, dictionary definitions, and thesaurus entries. These large-scale semantic networks display a particular pattern of connectivity suggestive of the growth processes through which they have been formed. In particular, they show a power-law structure in the distribution of links across nodes, such that most nodes in the network are of low degree (i.e., have few links to other nodes), but a few are hublike, with high degree (i.e., have many links to other nodes).

The power-law distribution of links over nodes has been shown to emerge when networks grow following the principle of preferential attachment (Barabasi & Alberts, 1999; Steyvers & Tenenbaum, 2005). This kind of growth has been observed, for example, in the growth of the Internet and the development of protein-interaction networks. In both cases, it has been hypothesized that network growth involves duplication of existing nodes, followed by rearrangement of some subset of the existing links (Kumar et al., 2000; Wagner, 2001). This duplication-and-divergence growth pattern is known to generate scale-free (or power-law) degree distributions (e.g., Pastor-Satorras, Smith, &

Address correspondence to Thomas T. Hills, Cognitive and Decision Sciences, University of Basel, Switzerland, e-mail: thomas.hills@unibas.ch.

Sole, 2003), and is also the form of growth that Steyvers and Tenenbaum (2005) postulated for semantic networks.

If the psychological processes that underlie early noun learning adhere to the principle of preferential attachment, then nouns that are learned earliest in these networks should show proportionally more connections during later stages of network development. Steyvers and Tenenbaum (2005) tested this hypothesis using several measures of age of acquisition and found a consistent pattern of higher semantic connectivity in an associative network (measured by degree) for early nouns than for later nouns. This is a potentially profound result because it suggests how growing expertise in one domain (a high-degree node) may select the entry of new information—one way in which early word learning could guide later word learning.

There is, however, an alternative account of the relation between age of acquisition and degree of connectivity in the adult network. Early-learned words may tend to be highly connected because words with greater connectivity in the learning environment are more noticeable and readily learned; that is, children may learn first the most well-connected words in the speech stream to which they are exposed. This is possible because the adult semantic network is both a product of learning and the input (the material to be learned) for the next generation of learners. We call this hypothesis *preferential acquisition*. A third growth process that may also explain the observed results is what we call *the lure of the associates*: Unknown words may be highlighted by known words to which they are related and learned in proportion to those relations.

We tested these three hypotheses—preferential attachment, preferential acquisition, and the lure of the associates—by examining the longitudinal development of networks of nouns that are normatively in the productive vocabularies of children at monthly intervals from 16 to 30 months. If these networks grow by preferential attachment, then one should be able to predict the nouns learned at time $t + 1$ from those known at time t . In particular, nouns learned at time $t + 1$ should be those that connect to the higher-degree known nouns. In contrast, if these networks grow by preferential acquisition, then the nouns learned at any developmental period should be the highest-degree nouns left to learn in the learning environment, regardless of how they connect to already-known nouns. Finally, if these networks grow via the lure of the associates, then the newly learned nouns should be those related to the largest number of nouns in the existing network, regardless of how the known words connect to each other.

THE NETWORKS

We built networks from 130 nouns selected from the MacArthur-Bates Communicative Developmental Inventory (MCDI; Dale & Fenson, 1996), Toddler version. This inventory provides month-by-month norms indicating the proportion of children—ages 16

to 30 months—who have each noun in their productive vocabulary. The 130 words we investigated were those words from the MCDI that also had feature norms available (see the next section). These nouns included names for animals (33 nouns), foods (17 nouns), small household objects (25), vehicles (12), outdoor things (5), clothing (15), furniture and rooms (14), toys (7), and places to go (2).

Creating semantic networks from these nouns required an index of semantic relatedness that reflected the semantic relations likely to be relevant to young children. We used two such indices, creating two series of networks. Each approach was based on dependent measures of semantic relatedness taken from adults, and each has its potential strengths and weaknesses.

Features

Many theories of category organization are based on shared features, such as having eyes and being furry (see, e.g., Rogers & McClelland, 2005). Moreover, developmental studies indicate that children are sensitive to and know about the features that are characteristic of common categories (e.g., Keil & Batterman, 1984; Sheya & Smith, 2006; Younger, 1990). To construct developmental feature networks, we used the feature norms reported by McRae, Cree, Seidenberg, and McNorgan (2005), who collected feature norms for 541 nouns (including the 130 early-learned nouns in our networks) from a total of 725 adult participants. The generated features included perceptual features (e.g., has wheels), functional features (e.g., can be eaten), taxonomic features (e.g., is an animal), and encyclopedic features (e.g., was invented by the Wright brothers). We used only perceptual and functional features, excluding encyclopedic and taxonomic features as unlikely to be available to children's direct experience.

Associates

Word associations are also a plausible index of semantic relatedness because they have proven robust in predicting adults' semantic judgments (e.g., Nelson, Zhang, & McKinney, 2001). Associates are collected by providing subjects with a word (the cue) and asking them to provide the first word that comes to mind (the target). This establishes a cue-target relationship (e.g., when provided with the cue word *cat*, many subjects provide the target word *dog*). In summarizing a large body of work, Deese (1965) concluded that word associations reflect the contiguity, semantic, and frequency properties of words in the language. Co-occurrence, in particular, seems to be a primary factor (Lund, Burgess, & Audet, 1996; Spence & Owens, 1990), with words that appear together in language more frequently also having a higher likelihood of appearing in associative pairs. We used the adult-generated University of South Florida Free Association Norms (Nelson, McEvoy, & Schreiber, 1999), the same norms used by Steyvers and Tenenbaum (2005). They consist of

approximately 5,000 cue words (including the 130 early-learned nouns in our networks) and their related targets.

The Networks

In our networks, nodes represented nouns, and links between nodes represented relationships between nouns. For each monthly interval from 16 to 30 months of age, we included only nouns that were in the productive vocabulary of more than 50% of the children in the MCDI at that month. This generated a developmentally ordered set of 15 feature networks and 15 associative networks. Links between nodes were added as follows. For the feature networks, a link existed between two nouns if they shared one or more features.¹ For example, *ball* and *apple* had a link because they share the feature “is round.” For the associative networks, a noun pair was connected by a directed link from the cue word to the target word if that cue-target relationship was reported in the association norms. In the feature networks, the total degree of a word was calculated by adding the number of links to that word. For the directed-associative networks, indegree was calculated by adding the number of links for which the word was the target of a cue-target pair.² To avoid redundancy when discussing both kinds of networks, we use “degree,” noting that in the case of the associative networks, we mean “indegree.” Figure 1 presents the feature and associative networks at 30 months. (For purposes of visual comparison, Fig. 2 presents the feature and associative networks at 20 months.) The feature network at 30 months contained 130 nodes with 2,433 undirected links (mean degree = 37.4, *SD* = 15.0). The associative network at 30 months contained 130 nodes with 235 directed links (mean degree = 1.8, *SD* = 2.1).

RESULTS

Age of Acquisition, Power Laws, and Static Networks

We first wanted to establish that the structure of the static (i.e., not growing) 30-month network was statistically comparable to the adult networks examined by Steyvers and Tenenbaum (2005), and thus could have grown by a process of preferential attachment. This would provide grounds for the comparison of alternative growth models, and would also allow other researchers to generalize the conceptual utility of our growth

models to other networks that may—because of the statistical properties of their structure—appear to have grown via preferential attachment.

A signature of network growth via preferential attachment is the power-law distribution of degree in the network (i.e., most nodes are connected to only a few other nodes, but a few nodes are hublike and have many links to other nodes). More precisely, the probability that a node in the network connects with k other nodes is roughly proportional to $k^{-\gamma}$: $P(k) \sim k^{-\gamma}$, where γ is a constant called the scaling parameter. This relation is revealed by a linear pattern when the degree distribution of the nodes is plotted on a log-log plot; γ is equal to the slope of the line. To avoid the potential bias associated with how data are binned in log-log plots (Clauset, Shalizi, & Newman, 2007), we plotted degree relative to its cumulative distribution—showing the probability that a randomly chosen node was of degree equal to or higher than k . Figure 3 shows the log-log plots for the 30-month feature and associative networks. The feature-based network did not provide strong evidence of a power-law structure, but the associations-based network ($\gamma = 1.97$) was consistent with the power-law distribution found by Steyvers and Tenenbaum (2005; $\gamma = 1.79$).

If a network grows by preferential attachment or by preferential acquisition, then the words learned during the earliest stages of development should have a higher degree than those learned during later stages. Figure 4 shows the degree for each word as a function of the age of acquisition for that word, separately for the 30-month feature and associative networks. A linear regression with degree at 30 months as the dependent variable and age of acquisition as the independent variable found no significant relationship between the two variables in the feature network ($b = 0.03$, $R^2 = .001$, $p = .28$). A similar regression using the associative network did find a significant relationship between degree at 30 months and age of acquisition ($b = -0.77$, $R^2 = .18$, $p < .001$). Words acquired earlier had a higher degree at 30 months, on average, than words acquired later. This finding is consistent with the pattern reported by Steyvers and Tenenbaum (2005), but our analysis shows that this pattern is also found over a much shorter time scale of only 14 months. Because age of acquisition and number of free associates are correlated with word length (Hutchinson, Balota, Cortese, & Watson, 2007; Storkel, 2004), we conducted a second analysis with word length as a covariate. Age of acquisition still predicted degree at 30 months, but word length was a significant, though more modest, predictor of age of acquisition ($b = 1.74$, $R^2 = .05$, $p < .01$).

This evidence—based on the structure of the static network (i.e., a single point in time)—is consistent with a model of preferential attachment. However, it is also consistent with alternative growth processes, such as preferential acquisition or the lure of the associates. We formalized the models based on these growth processes and then—using longitudinal data—implemented two tests for distinguishing among them.

¹We report analyses for feature networks that used a threshold of one or more shared features. We also analyzed fully weighted feature networks (in which the weight of each link was defined as the number of shared features), networks using different threshold values (more than zero or more than two shared features), networks based on the most frequently produced features for each word, and networks that used only perceptual or functional features (see McRae et al., 2005). In all cases, the pattern of results was the same as reported here.

²In the case of the associative networks, we present results based on indegree calculations only. Results for the outdegree calculation (outdegree is the sum of the links where the word is the cue in a cue-target pair) and total degree (outdegree plus indegree) were in all cases weaker than those calculated for indegree.




Fig. 1. Networks of the nouns used in this study: (a) feature-based network and (b) associations-based network for nouns normatively acquired by 30 months of age. For visual clarity, only links representing two or more shared features are shown in (a). The arrows indicating links in (b) point from the stimulus (cue) to the response (target). See the text for details on network construction.

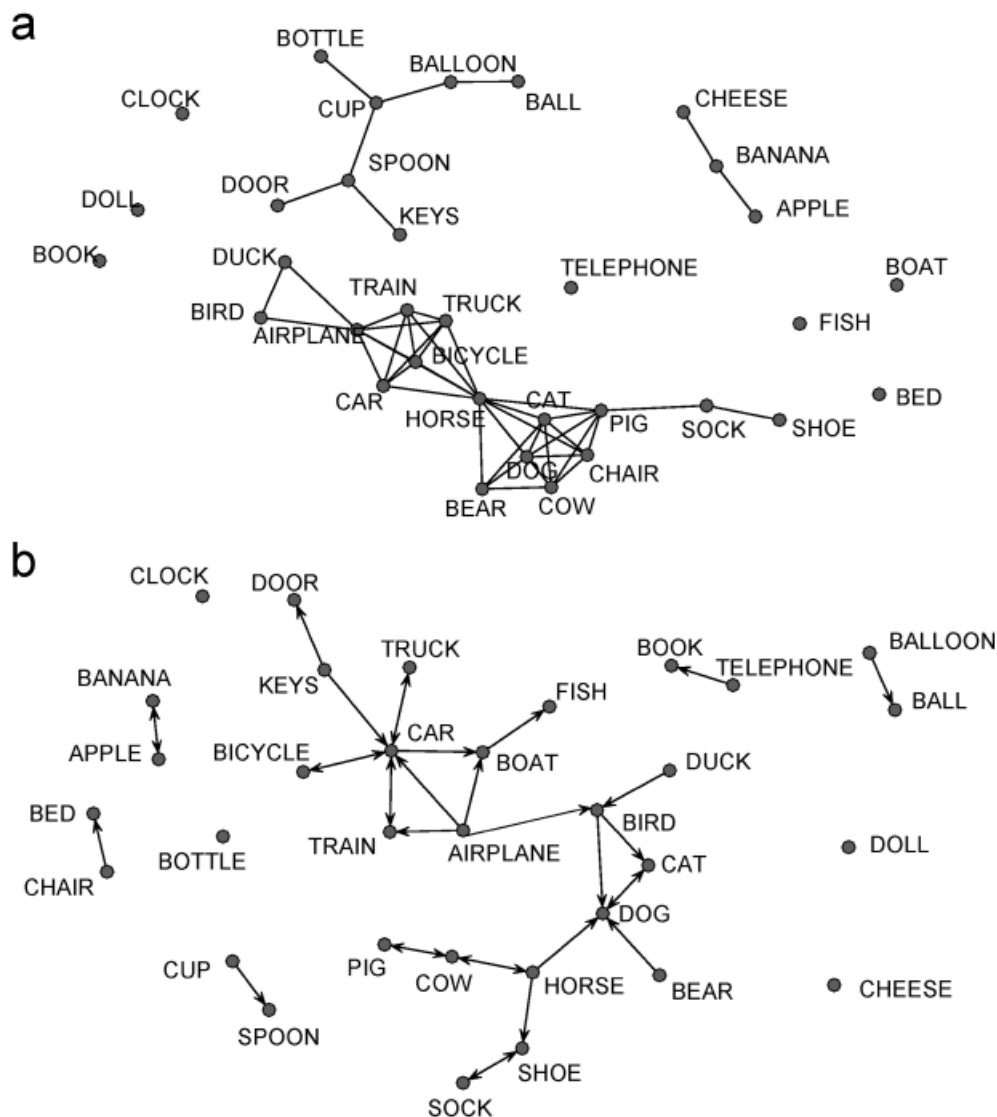


Fig. 2. Networks of the nouns used in this study: (a) feature-based network and (b) associations-based network for nouns normatively acquired by 20 months of age. For visual clarity, only links representing two or more shared features are shown in (a). The arrows indicating links in (b) point from the stimulus (cue) to the response (target). See the text for details on network construction.

Comparing Growth Hypotheses Using Longitudinal Data

We developed three formal models of growth based on our alternative hypotheses and tested them against one another in two separate ways. First, for each model, at each consecutive month of the networks' growth, we calculated a *z* score for each learned word, comparing the model-appropriate growth value of the learned word with the distribution of growth values for all words that could have been learned.³ The growth value of each word was indexed by the following: the mean degree of the known words that the new word attached to (for preferential attachment); the degree of the word, counting only links from known

words at the time of acquisition (for the lure of the associates); or the degree of the word in the presumed learning environment (for preferential acquisition).⁴ Figure 5 uses a simplified network to illustrate the differences among these models. To test the preferential-acquisition hypothesis in the case of the associative networks, we used two different networks to represent the learning environment—one based on child language (the 30-month network) and one based on adult language (the adult network). Thus, for each learned word, we found the mean and standard deviation of growth values for all possible words that could have been learned at the same time, and calculated the learned word's *z* value with respect to the resulting distribution.

³To avoid biasing the results with an overly limited learning set for the later developmental periods, we analyzed growth only over the first 10 months (16 to 26 months), while including all words acquired by the age of 30 months.

⁴Using the maximum degree (instead of the mean) for the preferential-attachment model did not alter the results.




Fig. 3. Log-log plot of the cumulative degree distribution for the 30-month feature network (left) and associative network (right). The graphs show the probability that a randomly chosen node is of degree equal to or higher than k .

Then, the z scores for each model were subjected to a one-sample t test to determine if the observed distribution of z scores was different from zero. The results for these analyses are presented in Table 1.

The main results were as follows: First, none of the growth models fit the growth of the feature networks. Second, the preferential-attachment model did not fit the developmental growth of the associative networks. Third, the preferential-acquisition and the lure-of-the-associates models, when applied to the associative networks, did describe the observed growth pattern of children's early noun vocabularies. Nouns with higher indegree at the time of acquisition or in the learning environment—either in the network of nouns commonly known by 30-month-olds or in the adult network—were more likely to be learned at any given time during the growth of the networks than were words with lower indegree. Only the lure-of-the-associates model considers the state of the child's current knowledge, and this model fit less well than the preferential-acquisition model, whether the learning environment in the latter model was represented by the 30-month network or by the adult network. The failure of the preferential-attachment model to fit these data (z scores and t values were close to zero) is remarkable. Known words of higher indegree are not more likely to receive links from new words entering the network than are known words of lower indegree. The conclusion is clear: The associative networks do not grow because known hub nouns attract new nouns. Rather, they grow because unknown words are themselves hubs that have many associative relationships with words in the learning environment.

As a second test, we asked which model for growth best fit the associative-network data and whether including additional parameters in the preferential-acquisition model would improve its predictive power. Our test determined, for each growth model, how strongly the growing network weighed the growth value of new words at each sequential month. It did this using a

parameter, β , which we fitted to an exponential ratio-of-strengths rule. We calculated the probability that a word, w_i , would be added to the network at a given month on the basis of its growth value, d_i , using the following equation:

$$P(w_i) = \frac{e^{\beta d_i}}{\sum_j e^{\beta d_j}}, \quad (1)$$

where β represents the sensitivity of the acquisition process to d_i . Positive values of β mean that words with higher values of d_i were more likely to be acquired early, whereas negative values of β mean that words with lower values of d_i were more likely to be acquired early.

We calculated d_i for each word with respect to each model, at the month of acquisition (e.g., for the lure-of-the-associates model with the associative network, d_i was equivalent to the indegree of word i at the time of acquisition). The denominator in Equation 1 was calculated using all words that were not yet learned at the start of the month in which the word in the numerator was acquired. Thus, for each model, we calculated a probability of acquisition for each acquired word at the month it was acquired. The logs of the $P(w_i)$ values for all acquired words were then summed to produce the log likelihood for that model:

$$-\log(L(\beta)) = -\sum_i \log(P(w_i))$$

We then found the β^* that minimized the log-likelihood function for each model using a standard optimization procedure.

This framework allowed us to compare models. For example, by replacing βd_i in Equation 1 with $\beta_{\text{LOA}} d_{i,\text{LOA}} + \beta_{\text{30mts}} d_{i,\text{30mts}}$, we could compute β s for a model that combined the lure-of-the-associates (LOA) model with the preferential-acquisition model that represented the learning environment as the network at 30 months (30mts). To determine when models provided additional explanatory power, we compared nested models using the likelihood-ratio test.




Fig. 4. Degree for each word in the 30-month network as a function of the age of acquisition. Results are shown separately for the feature network (top panel) and associative network (bottom panel). Best-fitting regression lines are shown.

Figure 6 shows the log-likelihood results for each model that performed better than a random model.⁵ The results indicated that the best model was the preferential-acquisition model with a learning environment represented by child vocabulary at 30

⁵Models using preferential attachment and word length are not shown in Figure 6, as neither was significantly different from the random model, with β set to 0.

months. This model was not significantly improved by including the β for either the preferential-acquisition model in which adult associations were the measure of the learning environment or the lure-of-the-associates model. These results show that what matters is not connectivity of words in the child's head, but rather connectivity in the learning environment—and, in particular, connectivity in a reduced learning environment most consistent with what children produce at 30 months.

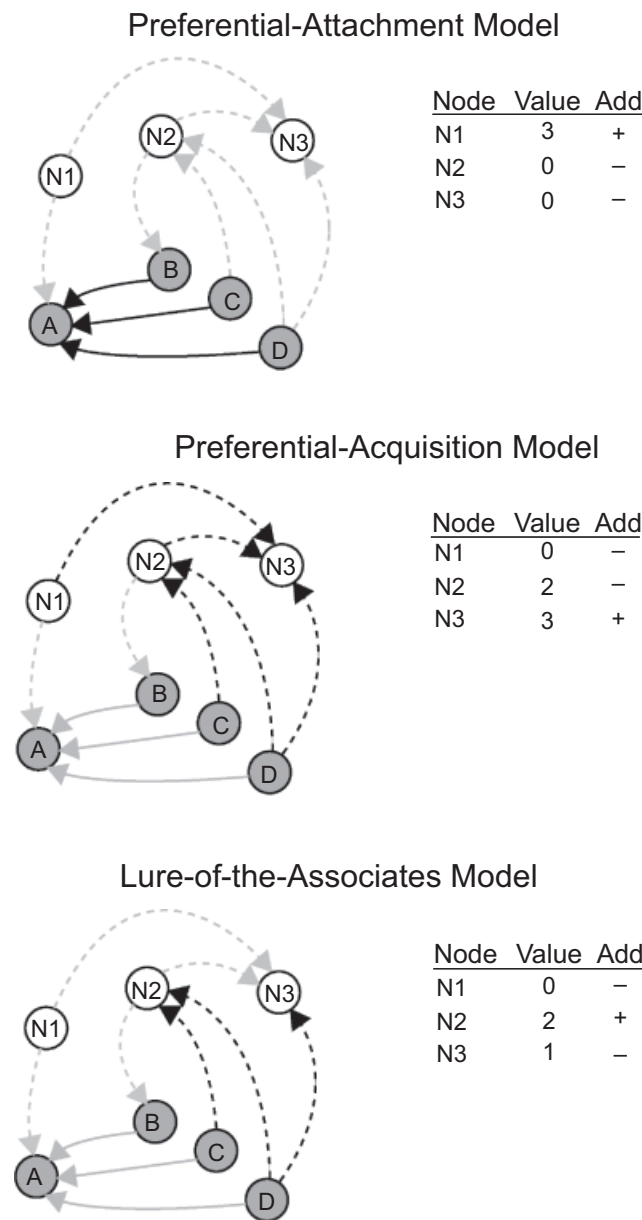


Fig. 5. The three growth models depicted in a simplified network. The network is the same in all three models, but the models assign different values to the unknown words. Gray shading and solid lines indicate nodes and links in the existing network (known words: A–D); no shading and dotted lines indicate nodes and links not yet incorporated into the known network (possible new nodes: N1, N2, and N3). Black lines indicate links relevant to the growth models, and gray lines indicate unimportant links. The “Add” column in each illustration indicates which node is favored for learning by the growth model in question; this is determined by the relative growth values of the possible new nodes. The growth values computed in this example are based on indegree for a directed network; arrow direction is important. For undirected networks, such as a feature network, arrow direction is not relevant. In the preferential-attachment model, the value of a new node is the average degree of the known nodes it would attach to (e.g., N1 is attached to A, which has an indegree of 3). In the preferential-acquisition model, the value of a new node is its degree in the learning environment—that is, the full network (e.g., N3 has an indegree of 3, which includes one link from a known node and two links from unknown nodes). In the lure-of-the-associates model, the value of a node is its degree with respect to known words (e.g., N3 has a value of 1, based on its one connection from a known node).

TABLE 1

Tests of the Distribution of Learned Words for Each Growth Model

Model	Associative networks			Feature networks		
	Mean z	$t(86)$	p	Mean z	$t(86)$	p
Preferential attachment	0.09	0.88	.38	-0.05	-0.57	.56
Preferential acquisition (30 months)	0.47	3.40	< .01	-0.002	-0.03	.98
Preferential acquisition (adult)	0.38	3.22	< .01	—	—	—
Lure of the associates	0.28	2.08	< .05	-0.07	-0.74	.46

Note. For the preferential-acquisition models, the designation in parentheses indicates the network used to represent the environment.

We conducted this study to examine how the structure of information influences early noun learning, and the results strongly suggest that the semantic connectedness of nouns in the learning environment is a significant contributor. This does not mean that other factors do not matter. Because frequency, phonotactic probability, and number of phonological neighbors (e.g., Charles-Luce & Luce, 1995; Landauer & Streeter, 1973; Pisoni, Nusbaum, Luce, & Slowiczek, 1985; Storkel, 2001) could be related to both semantic connectedness and age of acquisition, we tested these additional factors, both independently and in combination with the best model obtained using Equation 1. Phonotactic probability, number of phonological neighbors, and word frequency were from the Child Language Data Exchange System corpus (CHILDES) and Kučera and Francis (1967; values were taken from the following sources: Balota et al., 2007; MacWhinney, 2000; Vitevitch & Luce, 2004). The best statistical model for which the additional parameters were justified (using the likelihood-ratio test) included preferential acquisition (learning environment based on the 30-month vocabulary), word frequency in CHILDES, and the number of phonological neighbors (learning environment based on 30-month vocabulary: $\beta = 0.13$; CHILDES frequency: $\beta = 0.0005$; number of phonological neighbors: $\beta = 0.024$; $-\log(L) = 419.67$). Thus, the order in which nouns are learned early in word acquisition appears to be a combined consequence of semantic, frequency, and phonological factors.

DISCUSSION

These findings provide a first look at network growth during development. As in all studies of semantic networks, inferences to psychological processes are limited by the appropriateness of the indices of semantic relatedness—in the present case, the appropriateness of adult free associations and feature generation. Ultimately, the best evidence for the psychological relevance of any such index is the successful prediction of behavioral data in psychological experiments. With these limitations in mind, we conclude that the present study makes three




Fig. 6. Negative log likelihoods of the lure-of-the-associates (LOA) and preferential-acquisition (PA) models and their combinations. Two preferential-acquisition models were tested: one in which the learning environment was represented by the child's network at 30 months ("30 mts") and one in which the learning environment was represented by the adult network ("Adult"). Brackets and asterisks indicate nested models that were significantly different from each other in a likelihood-ratio test, $p < .05$.

contributions to an understanding of semantic networks and their implications for models of language acquisition.

First, the results challenge the principle of preferential attachment as a description of developmental processes. Although the links in adult semantic networks may be characterized by distributions that approximate a power law, and although such structures *may* be generated by a preferential-attachment growth rule, this does not mean that the psychological processes that give rise to adult semantic networks implement a preferential-attachment rule, or that the psychological processes that give rise to adult semantic networks are the same processes involved in child language acquisition. For example, the more-gets-more organization of links in adult semantic networks may reflect processes of retrieval, accessibility, and use, as well as the evolution of language (e.g., Ferrer i Cancho & Solé, 2003; Griffiths, Steyvers, & Firl, 2007; Steyvers & Tenenbaum, 2005; Zipf, 1949). Thus, the structure of the speech information presented to children may arise via one set of psychological processes, and be learned by children using another set of psychological processes.

Second, the results suggest that early in vocabulary growth, the order of acquisition of nouns is driven by the connectivity (both semantic and phonological) of words in the learning environment to each other, not by the connectivity of the words actually known by the child. That order of word acquisition reflects the statistical structure of the learning environment is consistent with many aspects of language learning (Hart &

Risley, 1995; Huttenlocher, Haight, Bryk, Seltzer, & Lyons, 1991). The implication of this finding is that children most readily learn those words that are central figures in the semantic and phonological landscape (i.e., words that share both semantic and phonological similarity with a variety of other words). This may be because such words are more noticeable than others; they are the figures in a many-word ground. Specifically, with respect to semantics, the relevant meaning or reference of such words may be more discernible because of the many shared contexts with other related words. Alternatively, child-directed speech may emphasize words of high degree, as it is known that emphasizing words in speech influences early word learning (Johnson & Jusczyk, 2001). One limitation of our approach is that the networks reflect the order of development in the normative child; the next step is to longitudinally examine developing networks in individual children.

Third, although developmentally interpretable patterns of growth characterized our associative networks, they did not characterize our feature networks. Perhaps the adult-generated features we used are not representative of the psychological features that characterize early categories. The features were generated by adults using their own—not a child's—perspective, and the features generated in feature-generation studies (e.g., McRae et al., 2005) have known limitations, including the exclusion of many properties that would seem deeply important to differentiate certain kinds (e.g., linguistically hard-to-describe properties such as the shape of a cow and essential but

ubiquitous properties such as “breathes”). Although we do not dismiss the possibility that our feature networks did not capture feature information important for word learning, another interpretation may be more consistent with the research on features, taken as a whole. First, the features from McRae et al. (2005) have predicted adult performance in a variety of studies (e.g., McNorgan, Kotack, Meehan, & McRae, 2007). Second, these features are consistent with the development of early superordinate categories (Hills, Maouene, Sheya, Maouene, & Smith, 2008). Perhaps feature relations organize knowledge (e.g., reasoning about instance-category relations), but do not order acquisition. In conclusion, the structure of known semantic knowledge and the acquisition of new knowledge may be deeply related but not identical.

Acknowledgments—This work was supported by the National Institutes of Health, T32 Grant HD 07475, and by National Institute of Mental Health Grant R01MH60200 to Linda Smith. We thank Brian Riordan, Peter M. Todd, and two anonymous reviewers for a number of challenging questions.

REFERENCES

- Balota, D.A., Yap, M.J., Cortese, M.J., Hutchinson, K.A., Kessler, B., Loftis, B., et al. (2007). The English Lexicon Project. *Behavior Research Methods*, *39*, 445–459.
- Barabasi, A.L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509–512.
- Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.
- Charles-Luce, J., & Luce, P.A. (1995). An examination of similarity neighborhoods in young children's lexicons. *Journal of Child Language*, *17*, 205–215.
- Clauset, A., Shalizi, C.R., & Newman, M.E.J. (2007). Power-law distributions in empirical data. Retrieved January 1, 2008, from <http://arxiv.org/abs/0706.1062>
- Dale, P.S., & Fenson, L. (1996). Lexical development norms for young children. *Behavior Research Methods, Instruments, & Computers*, *28*, 125–127.
- Deese, J. (1965). *The structure of associations in language and thought*. Baltimore: Johns Hopkins University Press.
- Dorogovtsev, S.N., & Mendes, J.F.F. (2001). Scaling properties of scale-free evolving networks: Continuous approach [Electronic version]. *Physical Review E*, *63*, Article 056125.
- Ferrer i Cancho, R., & Sole, R.V. (2003). Least effort and the origins of scaling in human language. *Proceedings of the National Academy of Sciences, USA*, *100*, 788–791.
- Griffiths, T.L., Steyvers, M., & Firl, A. (2007). Google and the mind: Predicting fluency with PageRank. *Psychological Science*, *18*, 1069–1076.
- Hart, B., & Risley, T.R. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore: Brookes.
- Hills, T., Maouene, J., Sheya, A., Maouene, M., & Smith, L. (2008). Categorical structure in early semantic networks of nouns. In V. Sloutsky, B. Love, & K. McRae (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 161–166). Mahwah, NJ: Erlbaum.
- Hutchinson, K.A., Balota, D.A., Cortese, M.J., & Watson, J.M. (2007). Predicting semantic priming at the item level. *The Quarterly Journal of Experimental Psychology*, *61*, 1036–1066.
- Huttenlocher, J., Haight, W., Bryk, A., Seltzer, M., & Lyons, T. (1991). Early vocabulary growth: Relation to language input and gender. *Developmental Psychology*, *27*, 236–248.
- Johnson, E., & Jusczyk, P.W. (2001). Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language*, *44*, 548–567.
- Keil, F.C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, *23*, 221–236.
- Kučera, H., & Francis, W. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.
- Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., & Upfal, E. (2000). Random graph models for the web graph. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)* (pp. 57–65). Redondo Beach, CA: IEEE CS Press.
- Landauer, T.K., & Streeter, L.A. (1973). Structural differences between common and rare words: Failure of equivalence assumptions for theories of word recognition. *Language and Cognitive Processes*, *16*, 565–581.
- Lund, K., Burgess, C., & Audet, C. (1996). Dissociating semantic and associative word relationships using high-dimensional semantic space. In G.W. Cottrell (Ed.), *Proceedings of the 18th Annual Conference of the Cognitive Science Society* (pp. 603–608). Mahwah, NJ: Erlbaum.
- MacWhinney, B. (2000). *The CHILDES project: Tools for analyzing talk* (3rd ed.). Mahwah, NJ: Erlbaum.
- McNorgan, C., Kotack, R.A., Meehan, D.C., & McRae, K. (2007). Feature-feature causal relations and statistical co-occurrences in object concepts. *Memory & Cognition*, *35*, 418–431.
- McRae, K., Cree, G.S., Seidenberg, M.S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods, Instruments, & Computers*, *37*, 547–559.
- Nelson, D.L., McEvoy, C.L., & Schreiber, T.A. (1999). *The University of South Florida word association, rhyme, and word fragment norms*. Retrieved July 24, 2007, from <http://www.usf.edu/freeassociation/>
- Nelson, D.L., Zhang, N., & McKinney, V.M. (2001). The ties that bind what is known to the recognition of what is new. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1147–1159.
- Pastor-Satorras, R., Smith, E., & Sole, R. (2003). Evolving protein interaction networks through gene duplication. *Journal of Theoretical Biology*, *222*, 199–210.
- Pisoni, D.B., Nusbaum, H.C., Luce, P.A., & Slowiaczek, L.M. (1985). Speech perception, word recognition and the structure of the lexicon. *Speech Communication*, *4*, 75–95.
- Rogers, T.T., & McClelland, J.L. (2005). *A parallel distributed processing approach to semantic cognition: Applications to conceptual development*. Mahwah, NJ: Erlbaum.
- Sheya, A., & Smith, L.B. (2006). Perceptual features and the development of conceptual knowledge. *Journal of Cognition and Development*, *7*, 455–476.
- Spence, D.P., & Owens, K.C. (1990). Lexical co-occurrence and association strength. *Journal of Psycholinguistic Research*, *19*, 317–330.

- Steyvers, M., & Tenenbaum, J.B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science*, *29*, 41–78.
- Storkel, H.L. (2001). Learning new words: Phonotactic probability in language development. *Journal of Speech, Language, and Hearing Research*, *44*, 1321–1337.
- Storkel, H.L. (2004). Do children acquire dense neighborhoods? An investigation of similarity neighborhoods in lexical acquisition. *Applied Psycholinguistics*, *25*, 201–221.
- Vitevitch, M.S. (2008). What can graph theory tell us about word learning and lexical retrieval? *Journal of Speech, Language, and Hearing Research*, *51*, 408–422.
- Vitevitch, M.S., & Luce, P.A. (2004). A web-based interface to calculate phonotactic probability for words and nonwords in English. *Behavior Research Methods, Instruments, & Computers*, *36*, 481–487.
- Wagner, A. (2001). The yeast protein interaction network evolves rapidly and contains few redundant duplicate genes. *Molecular Biology and Evolution*, *18*, 1283–1292.
- Younger, B. (1990). Infants' detection of correlations among feature categories. *Child Development*, *61*, 614–620.
- Zipf, G.K. (1949). *Human behavior and the principle of least effort: An introduction to human ecology*. Cambridge, MA: Addison-Wesley.

(RECEIVED 5/20/08; REVISION ACCEPTED 11/7/08)

This document is a scanned copy of a printed document. No warranty is given about the accuracy of the copy. Users should refer to the original published version of the material.