

Word associations: Network and semantic properties

SIMON DE DEYNE AND GERT STORMS
University of Leuven, Leuven, Belgium

A number of properties of word associations, generated in a continuous task, were investigated. First, we investigated the correspondence of word class in association cues and responses. Nouns were the modal word class response, regardless of the word class of the cue, indicating a dominant paradigmatic response style. Next, the word association data were used to build an associative network to investigate the centrality of nodes. The study of node centrality showed that central nodes in the network tended to be highly frequent and acquired early. Small-world properties of the association network were investigated and compared with a large English association network (Steyvers & Tenenbaum, 2005). Networks based on a multiple association procedure showed small-world properties despite being denser than networks based on a discrete task. Finally, a semantic taxonomy was used to investigate the composition of semantic types in association responses. The majority of responses were thematically related situation responses and entity responses referring to parts, shape, or color. Since the association task required multiple responses per cue, the interaction between generation position and semantic role could be investigated and discussed in the framework of recent theories of natural concept representations (Barsalou, Santos, Simmons, & Wilson, in press).

Recently, the study of the representation of concepts and word associations has regained interest. Three different lines of research characterize studies of this issue. First, a number of studies have explored the distributional and structural properties of word association networks. It has been shown that the structure of connections between associations adheres to special topological laws, commonly found in many natural networks (Steyvers & Tenenbaum, 2005), such as small-world properties. Examples of these natural networks include citation networks (e.g., Redner, 1998), the World Wide Web (e.g., Albert, Jeong, & Barabási, 1999), metabolic networks (e.g., Jeong, Tombor, Albert, Oltvai, & Barabási, 2000), and the network of human sexual contacts (e.g., Liljeros, Edling, Amaral, Stanley, & Åberg, 2001). The network topology of word associates has also been used to explain semantic neighborhood effects on the recognition of words (Locker, Simpson, & Yates, 2003) and recognition success in list-learning experiments (Nelson, Zhang, & McKinney, 2001).

In a second line of research, spatial models based on association data have been used to explain the function of semantic memory in tasks such as semantic similarity ratings. With this approach, free association norms are used to derive proximity measures in a high-dimensional representation. These measures have been successful in predicting semantic similarity ratings of words and in free-recall and cued-recall tasks (Steyvers, Shiffrin, & Nelson, 2004).

Third, besides the application of word associations in spatial and network models, word associations play a central role in theories of language and concept processing, such as the dual coding theory (Paivio, 1986). This line of

research has focused on further adapting and improving theoretical accounts of the function of word associations in a larger system for representing meaning. A more recently advanced account of the dual coding theory is the *language and situated simulation of conceptual processing* (LASS—Barsalou, Santos, Simmons, & Wilson, in press) theory, which will be considered in further detail in a later section of this article. According to LASS, multiple systems are used for representations of concepts. One of these systems comprises linguistically grounded information similar to representations proposed by Burgess and Lund (1997) and by Landauer and Dumais (1997). Word associations are assumed to capture most of the representations in this language system. When a word is read, associations become automatically activated. The information activated by the word associations spreads to a second system, where conceptual information is stored. This system consists of situated simulations where knowledge is grounded in a modality-specific manner (Allport, 1985; Barsalou, 1999; Glenberg, 1997). According to this account, word associations act as pointers or heuristics that aid in the retrieval of related conceptual knowledge.

In our companion article (De Deyne & Storms, 2008), a multiple response procedure was used in which each participant responded with three associations for the cues that were presented to them. The procedure differed from those used in other recent large association studies, which employed a discrete response procedure in which only one response was collected per cue. The use of a multiple-production procedure allows weak associations to be represented in the norms and thus results in a denser representa-

S. De Deyne, simon.dedeyne@psy.kuleuven.be

tion of word associations. Apart from the representational properties that capture association responses, an additional hypothesis can be derived that focuses on the generation process, since the present task used a multiple response procedure. Since, unlike in most studies, participants were asked to generate three associates instead of one, it is possible that second and third association responses reflect different information. The reasoning behind this hypothesis is that the second and third association responses are presumably less available and require different processing than first associates. For instance, it is possible that certain association responses are generated late (i.e., second or third response) because of their semantic role. To investigate this, the semantic role of the first, second, and third association responses from the multiple response procedure will be systematically compared in this article. Similarly, the tendency to give responses that belong to the same word class as the cue (i.e., paradigmatic responses) can be attenuated when responses given early are compared with those given late.

The goal of the present study is to apply the different research lines described above to this large set of word associations (De Deyne & Storms, 2008) in order to provide more insight into the nature of word associations in the following ways. First, the current Dutch norms allow researchers to generalize their findings to languages other than English. For instance, Steyvers and Tenenbaum (2005) noted that it would be desirable to verify whether the large-scale statistical patterns found in associative networks correspond in different languages. Second, the methodology used in the present study allows a different approach to address questions about the origin and status of the associations themselves and their functions in semantic representations in general. For instance, the LASS theory has concrete predictions on how the time course of the generation of continuous associations interacts with the semantic content of these associations.

Association responses are characterized by a number of factors that all capture different but related aspects of word representations. In this article, three word representation aspects are considered. The first aspect concerns the influence of the word class of the cue on the association response. Previous studies of word associations have found that cues generally evoke paradigmatic responses (i.e., responses from the same word class) instead of responses that are syntagmatically related. Most of these studies were based on discrete tasks, which could give rise to a different pattern of results than continuous tasks would.

Second, we present a macroscopic description of the general network topology of word associations. More specifically, we investigated whether the availability properties of words render a specific topology called *small-world networks* (Steyvers & Tenenbaum, 2005). An important property of small-world networks is that only a small minority of the words are connected with a significant portion of the other words in the network. The connectivity and centrality of certain nodes that result from this topology can also explain processing advantages of words in terms of well-known availability measures such as word frequency. For instance, it is possible that words that occur frequently in language also have many word associates. This hypothesis

refers to the connectivity characteristics of the nodes in relation to the availability of words in tasks requiring word processing. Of course, word frequency is not the only measure that determines the centrality of words. We also compared the effect of word frequency with two other availability measures: age of acquisition (AoA) and word imageability.

Third, the associations are described in terms of their semantic properties. The number of studies that have investigated the semantic composition of word associates is quite small, whereas there is an abundance of semantic tasks that use word associates in one form or another. For instance, in the literature on semantic priming, a long-lasting debate has focused on the distinction between semantic priming and associative priming (e.g., Livesay & Burgess, 1999; Lupker, 1984; Shelton & Martin, 1992). One of the problems is that semantic relations and the interpretation of associations is often underspecified.

Although it has previously been suggested that word associations comprise semantic responses, further distinctions between these semantic responses along the lines proposed by Barsalou and colleagues (see, e.g., Barsalou et al., in press) could have extensive consequences for our understanding of semantic or conceptual processing. Before turning to the semantic analysis of word associations, the word class properties of these associations will be investigated. The distinction along word class also forms the basis of some of the analyses in the second section of this article, which deals with the analysis of network connectivity.

PARADIGMATIC AND SYNTAGMATIC ASSOCIATIONS

In early research concerning the nature of word associations, the distinction between paradigmatic and syntagmatic associations was the subject of extensive study (Deese, 1962; Glanzer, 1962). Paradigmatic associations are responses that belong to the same word class as the cue (e.g., noun–noun, verb–verb). Syntagmatic responses, in contrast, do not preserve the stimulus's word class and include, for instance, cue–response pairs with noun–verb or adjective–verb form class relations. In a review of the literature, Cramer (1968) came to the conclusion that in most association tasks, paradigmatic responses are the modal response. In terms of frequency of paradigmatic responses, the following order was observed: nouns > pronouns > adjectives > adverbs > verbs. Of course, to some extent syntagmatic responses also occur. The syntagmatic responses that did occur were found most frequently in the following order: adverbs > adjectives > verbs > nouns (Cramer, 1968). This means that when the cue is an adverb, there is a tendency to respond with other parts of speech (POSSs), and this tendency is stronger for this word class than for other word classes, such as adjectives, verbs, or nouns. Most of the research on paradigmatic and syntagmatic associations has focused on a discrete association task (Cramer, 1968). Since more than one association response is given in our data, it is possible that a shift from paradigmatic to syntagmatic associations occurred for the second and third responses. In the present analysis, two hypotheses are of importance. First, we investigate whether the dominant re-

sponse type is paradigmatic or syntagmatic. Second, we investigate whether this pattern holds in comparisons of first, second, and third word association responses.

Method

Materials and Procedure. A detailed description of the data and the procedure is presented in De Deyne and Storms (2008). For ease of understanding, a short summary of the procedure is represented here. The word associations were collected for a total of 1,424 words in two phases. During the first phase, word associations were collected for cues belonging to numerous semantic categories. In the second phase, the most generated association responses were presented as cues to a new group of subjects. These cues thus covered a wide variety of words. For each cue, the participants had to give three different word associations resulting in minimum 246 (3×82) and maximum 591 (3×197) association responses. The resulting data set of 381,909 responses is the starting point for analyzing the paradigmatic and syntagmatic properties of the association responses. To identify the POS of the associations in the De Deyne and Storms (2008) data set, we used the CELEX (Baayen, Piepenbrock, & van Rijn, 1993) POS tags. In cases in which more than one POS tag was available, the most frequent tag was used. Although this method is not perfect, it is sufficient for the present purposes.

Results and Discussion

For 327,833 responses (i.e., 86% of the association responses), the word class could be determined. The large majority of response tokens were nouns (72%) and adjectives (18%). Less frequently generated were verbs (9%), adverbs (0.4%), prepositions (0.34%), pronouns (0.13%), numerals (including quantifiers such as *more* and words referring to numeric values) (0.10%), interjections (e.g., *congratulations*) (0.08%), and expressions (0.08%). To investigate paradigmatic and syntagmatic responses, the contingencies between different POSs are investigated. We will limit this analysis to nouns, adjectives, and verbs. Since the cue set includes nouns, verbs, and adjectives, the 3×3 contingency table represents paradigmatic associations in the diagonal entries and syntagmatic associations in the off-diagonal entries. These contingencies are shown in Figure 1 for first (R1), second (R2), and third (R3) responses.

From Figure 1, it can be seen that nouns are given often as an initial response irrespective of whether the cue is a verb or an adjective. Additional evidence for syntagmatic associations is found for nouns when first, second, and third responses are compared. These responses show a de-

crease of noun responses in favor of adjectives and verbs. This finding provides evidence that a paradigmatic to syntagmatic shift occurs among the total numbers of first, second, and third responses. In order to investigate these contingencies, a 3 (cue word class) $\times 3$ (response word class) $\times 3$ (response position) log-linear model was formulated. The only fitting three-way log-linear model was a saturated model ($G^2 = 0$, $df = 0$). A more parsimonious model did not fit the data well, which is not surprising given the large number of observations in this data set. In other words, all possible main effects as well as all possible interaction effects were significant. Therefore, these effects will not be discussed in more detail.

In conclusion, nouns constitute the most frequent word class for the responses even if the cues are adjectives or verbs. The results for adjectives and verbs thus support the syntagmatic hypothesis and are in line with the results reported by Deese (1962), who also found that adjective and verb cues produce a majority of noun responses. Since nouns elicit mostly noun responses, this indicates that for this word class, paradigmatic responses are the modal response. The resulting pattern is more subtle when later responses (second and third), which are less paradigmatic and shift toward syntagmatic responses, are taken into account. In summary, the present findings contradict the previously reported results (Cramer, 1968) that paradigmatic responses dominate association responses for verbs and adjectives. On the other hand, the pattern observed for syntagmatic responses for verbs and adjectives was similar to the patterns of previous findings (Cramer, 1968).

The following section deals with different properties of centrality and connectivity in a word association network and goes into further detail on how the word class of word associations and cues interacts with the connectivity in a word association network.

SMALL-NETWORK ANALYSIS

Steyvers and Tenenbaum (2005) have examined the large-scale structure of semantic networks using data from word associations, WordNet, and Roget's *Thesaurus*. One of the major findings was the small-world network structure, characterized by sparse connectivity, short average

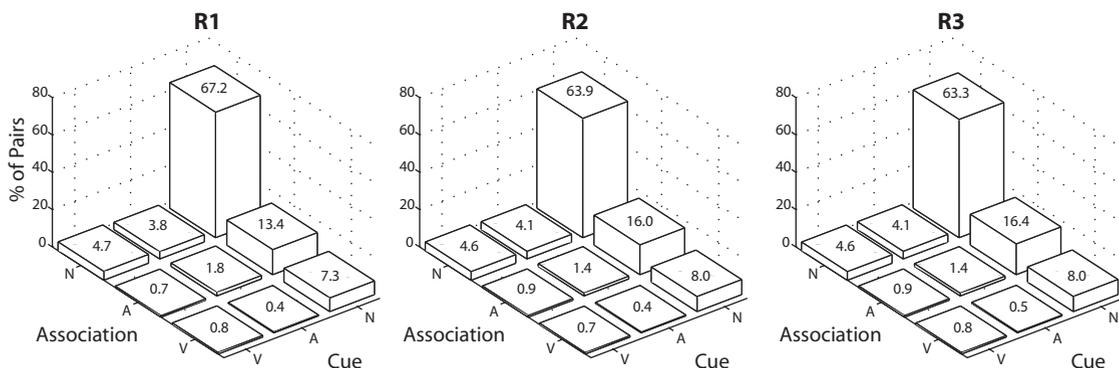


Figure 1. Contingency relationships between nouns (N), verbs (V), and adjectives (A) as cues and association responses, separately for first (R1), second (R2), and third (R3) responses.

path lengths between words, and strong local clustering in networks built upon these data. These network properties provided fundamental insights into the structure and processes in semantic memory. For example, Steyvers and Tenenbaum used the small-world network properties of associations to propose a model of semantic growth. On the basis of this model, an intuitive account was given for the origin of word frequency or AoA. In a similar analysis of our association data, we investigated the generalizability of these finds to a different language (Dutch vs. English) and to data from a continuous (instead of discrete) association task.

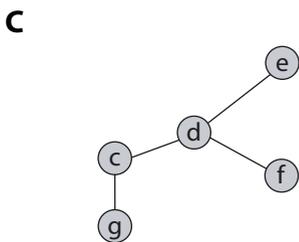
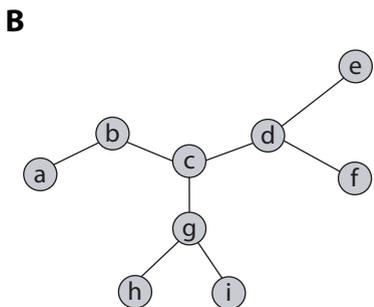
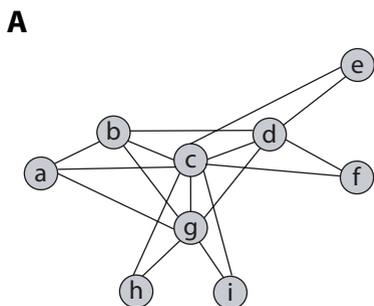
A Short Introduction to Network Analysis

First, a brief description of the relevant concepts of network analysis is given (see Newman, 2003, for a recent overview). A real example of a portion of the semantic network built from our word association data is shown in

Figure 2. This figure also illustrates some of the concepts that are discussed below.

An associative network is a graph (G) that consists of nodes or vertices (V) linked by connections or edges (E). In this case, the nodes correspond to cues connected by a weighted (directed) arc or (undirected) edge to other associated cues. A weight can be assigned to both the nodes as well as to the edges and arcs. However, for present purposes only the edges and arcs will be weighted by the frequency of association. Two nodes are neighbors if they are directly connected by an arc or edge. Nonadjacent nodes are connected by paths. These paths connect two edges or arcs with the restriction that the nodes on this path occur only once. There are many possible paths between two nodes, but the shortest path connecting two nodes is called the *geodesic path*.

From each graph, a number of statistics can be derived for the nodes and the connections. Each node has an *in-degree* k^{in} and an *out-degree* k^{out} corresponding to the



Average Path Length L for network B

$cg = 1$ $de = 1$ $eg = 3$ $fg = 3$
 $cd = 1$ $df = 1$ $ef = 2$
 $ce = 2$ $dg = 1$
 $cf = 2$

$$\frac{1+1+2+2+1+1+1+3+2+3}{10} = 1.8$$

Node degrees k for network B

$a = 1$ $d = 3$ $g = 3$
 $b = 2$ $e = 1$ $h = 1$
 $c = 3$ $f = 1$ $i = 1$

Closeness c for node c in network C

$c(c) = \frac{5-1}{6} = 0.67$

Betweenness b for node c in network C

$b(c) = \frac{\#(gd+ge+gf)}{\#(gd+ge+gf+de+df+ef)} = 0.5$

Clustering Coefficient CC1 for node c in network B

The numerator corresponds to the number of connections between c's neighbors, multiplied by 2. The denominator corresponds to the maximum number of possible connections between c's neighbors.

$CC1(c) = \frac{2(0)}{3(3-1)} = 0$

Clustering Coefficient CC1 for node a in network A

$CC1(a) = \frac{2(3)}{3(3-1)} = 1$

Figure 2. Examples of networks with strong (A) and weak (B) local clustering, and example calculations of network statistics for these networks. The statistics include node degree and average path length (B), clustering coefficients (A, B), and closeness and betweenness (C).

number of incoming and outgoing arcs in a directed network. For the corresponding undirected network, nodes have a certain degree, k , which is the number of edges of a node. The degree of the nodes is therefore also a measure of the importance and centrality of a certain node in the network. More advanced measures for the centrality of nodes in the network are the closeness and betweenness of a node or vertex. To make these measures easier to understand, Figure 2 provides an illustration and some example calculations of the measures that are discussed below.

Closeness centrality of a vertex is the number of other vertices divided by the sum of all shortest path lengths between the vertex and these other vertices. In comparison with degree centrality, closeness centrality has the advantage that it takes into account direct and indirect connections between nodes. The normalized form (Freeman, 1977) is

$$cl(u) = \frac{(n - 1)}{\sum_{v \in V} d(u, v)}, \quad (1)$$

where $d(u, v)$ is the geodesic (i.e., length of the shortest path) from u to v . The values of $cl(u)$ are bounded between 0 and 1. In a directed network, two closeness measures can be calculated: the first version is based on the incoming arcs, and the second version is based on the outgoing arcs.

Betweenness centrality measures how often a node is located on the geodesic path between other nodes in the network. If a node with a high level of betweenness were to be deleted from a network, the network would fall apart into otherwise coherent clusters. Unlike degree, which is a count, betweenness is normalized, by definition, as the proportion of all geodesics that include the vertex under study. The betweenness centrality (b) for a vertex u is defined as (adapted from Freeman, 1977):

$$b(u) = \sum_i \sum_j \frac{g_{ij}}{g_{ij}}, \text{ with } i \neq j \neq u. \quad (2)$$

In this formula, g_{ij} is defined as the number of geodesic paths between i and j , and g_{ij} is the number of paths between i and j that pass through u . The measure $b(u)$ has values between 0 and 1.

A final measure of centrality is the *clustering coefficient*, which is an important way to identify small-world properties of a network. For each node, a clustering coefficient (CC) can be calculated. For a node u , this is the proportion of edges with neighboring nodes divided by the maximal number of edges within this neighborhood. Intuitively, the clustering coefficient measures the likelihood that two nodes that have a mutual neighbor are themselves neighbors. The neighborhood is defined for direct adjacent neighbors (CC1) and neighbors that are related through an intermediate node (CC2). For a given undirected graph $G = (V, E)$, CC1 is derived as follows:

$$CC1_{(u)} = \frac{2\{E[G^1(u)]\}}{\deg(u) \cdot [\deg(u) - 1]}. \quad (3)$$

The Neighborhood 2 clustering coefficient CC2 follows from the previous equation:

$$CC2(u) = \frac{E[G^1(u)]}{E[G^2(u)]}. \quad (4)$$

The numerator corresponds to the number of lines among vertices in Neighborhood 1 of u , and the denominator corresponds to the number of lines among vertices in Neighborhoods 1 and 2 of u . Normalized coefficients CC1' and CC2' are obtained by multiplication with the degree of u divided by the maximal possible degree of u . Figure 3 illustrates direct neighbors as well as nodes in Neighborhoods 2 and 3 for the example of the word *strawberry* based on actual data.

The remaining statistics pertain to the distribution of edges or arcs in the network. Networks are described in terms of their *density*, or the proportion between the number of edges between the nodes and all possible edges if all these nodes were connected to each other. A second measure is the *average distance L* between two nodes. This is the average of the shortest path lengths between all node pairs in a network. A calculation example is shown in Figure 2. The measure is related to the *diameter D* of the network, which gives the maximum distance between any two nodes in the network. Finally, each network has a certain distribution $P(k)$, which is the probability that a random node will have degree k .

With the eight statistics introduced above, properties (e.g., small-world networks) that are typical for certain types of networks can be identified. Recent studies have shown that associative networks contain clusters of nodes with tightly coupled neighborhoods. These clusters or hubs of nodes differ from other nodes in that they maintain very short distances among nodes across the entire network. This pattern of connections is said to give rise to a small-world architecture. The degree k to which nodes are connected forms a distribution that decays as a power law, producing a scale-free architecture characterized by the existence of highly connected nodes or hubs (Barabási & Albert, 1999).

In summary, semantic networks have small-world properties if they exhibit short average path lengths and high clustering. A specific form of these small-world networks is scale free if the distribution of the degrees follows a power-law distribution. Using the previously introduced terminology for graphs and the above statistics, the small-world properties of our association network can be investigated and compared with the Steyvers and Tenenbaum (2005) analysis of the English association data collected by Nelson, McEvoy, and Schreiber (2004).

Small-World Properties

The first aim of the following analysis was to replicate the small-world network properties found by Steyvers and Tenenbaum (2005). Since the cues in our study had a fixed presentation rate (minimum 80 participants for each cue), cues and associations should not be presented simultaneously in one network. These cues and associations form a two-mode network and will have different properties. If no distinction between cues and associations is made, the cue representations will be biased due to their fixed presentation frequency and will make the comparison between nodes for cues and nodes for associations impossible. For example,

Table 1
Network Statistics for the Undirected and Directed Associative and Comparable Random Networks for the Leuven (De Deyne & Storms, 2008) and Florida (Nelson et al., 2004) Data Sets

Data Set	Density	k	k^{out}	L	D	CC1	CC2	cl	cl^{out}	b
Leuven ($N = 1,424$)										
Original										
Undirected	0.020	29	–	2.46	4	0.226	0.001	0.409	–	0.001
Directed	0.012	17	17	3.43	9	0.147	0.002	0.267	0.257	0.001
Random										
Undirected	0.020	29	–	2.52	4	0.020	0.001	0.397	–	0.001
Directed	0.012	17	17	2.86	4	0.012	0.001	0.350	0.012	0.001
Florida ($N = 5,018$)										
Original										
Undirected	0.004	22	–	3.04	5	0.186	0.001	0.330	–	0.000
Directed	0.003	13	13	4.26	10	0.121	0.000	0.236	0.228	0.001
Random										
Undirected	0.004	22	–	3.03	4	0.004	0.001	0.330	–	0.000
Directed	0.003	13	13	3.63	6	0.003	0.001	0.276	0.276	0.001

Note—See the text for an explanation of abbreviations.

both networks less than 2% of all possible associations between words are instantiated. This shows the extreme sparseness of these networks. The Florida network has an even sparser density (0.4% for the undirected version, 0.3% for the directed version), which can be explained by the larger variety in the cue set and the use of a discrete procedure for collecting the association responses. The average degree or number of different connections in the network was between 17 (directed) and 29 (undirected), and each node was reachable through on average of three connections (i.e., there was average path lengths L of 2.46 for the directed network and 3.43 for the undirected network). Analysis of the network diameter D showed that the most distant nodes were separated by four (undirected) and nine (directed) connections. The values for L and D found in the Florida network are slightly higher, which is hardly surprising due to lower density of the network.

The remaining statistics measure the average centrality of the nodes in the network, which allowed us to identify the small-world properties by comparison with a comparable random network. The neighborhood clustering coefficients CC1 show that the associates of a word were directly associated with each other as well. For the undirected network, this happened approximately 23% of the time, whereas for the directed network it occurred about 14% of the time. Again, Table 1 shows that these values were higher in the Leuven network than in the Florida network. Due to the high increase of neighboring nodes, the average clustering coefficient that includes indirect associations in Neighborhood 2 were still quite small. Both the directed and the undirected networks indeed showed a much higher clustering coefficient than did the corresponding random graph. Together with the average short path lengths L , the present network statistics corresponded closely to Steyvers and Tenenbaum's (2005) conclusion that the Florida network exhibits small-world properties.

To investigate the scale-free structure of the networks, the distribution $P(k)$ can be plotted as function of k . In the case of the undirected network, k is equal to the number of neighbors and does not make a distinction between incoming arcs (in-degree) or outgoing arcs (out-degree).

The plot of the degree distribution, together with the distribution in the Nelson data set, is shown in Figure 4.¹

The plot in Figure 4 shows the best-fitting regression line of the power-law distribution for the directed network based on the in-degree distribution.² The distributions were estimated by grouping all values of k into bins of consecutive values and computing the mean value of k for each bin. The mean value of each bin corresponds to one point, whereas the boundaries were spaced logarithmically to ensure approximately equal numbers of observations per bin.³ As can be seen in Figure 4, the value of the gamma parameter (1.86) was around the same value (1.79) that was originally reported by Steyvers and Tenenbaum (2005). However, in our own simulation of the Florida network, we obtained slightly larger gammas, which could be due to small variations in setting up the network. The plot shows a slight deviation from a power-law distribution, since the points in both networks did not follow a straight line. These small deviations might be explained by the fact that our network was denser and conceptually more coherent than the Florida network.

In the analysis of Steyvers and Tenenbaum (2005), no differentiation was made between nodes from different word classes. One possibility is that the hubs observed in these networks are primarily members of a certain word class, such as nouns or adjectives. We investigated the grammatical class of the most connected nodes in the network. The hubs were identified as the nodes that corresponded to the 10% most connected nodes based on their cumulative in-degree distribution. As can be expected, the number of these hubs was fairly low: 53 out of a total of 1,424 nodes. From these 53 nodes, 28 were adjectives, 24 were nouns, and 1 was a verb. The finding of a high proportion of adjectives is even more pronounced when one considers that the cues selected in the association task were primarily nouns, and that it has been shown that, in making their responses, people tend to generate nouns from the same word class (Cramer, 1968). One explanation for this finding can be found in the large variability in occurrences for the different word classes. For instance, the CELEX count for Dutch words contains 95,657 nouns, 13,912 adjectives, and 11,837 verbs. When a response is

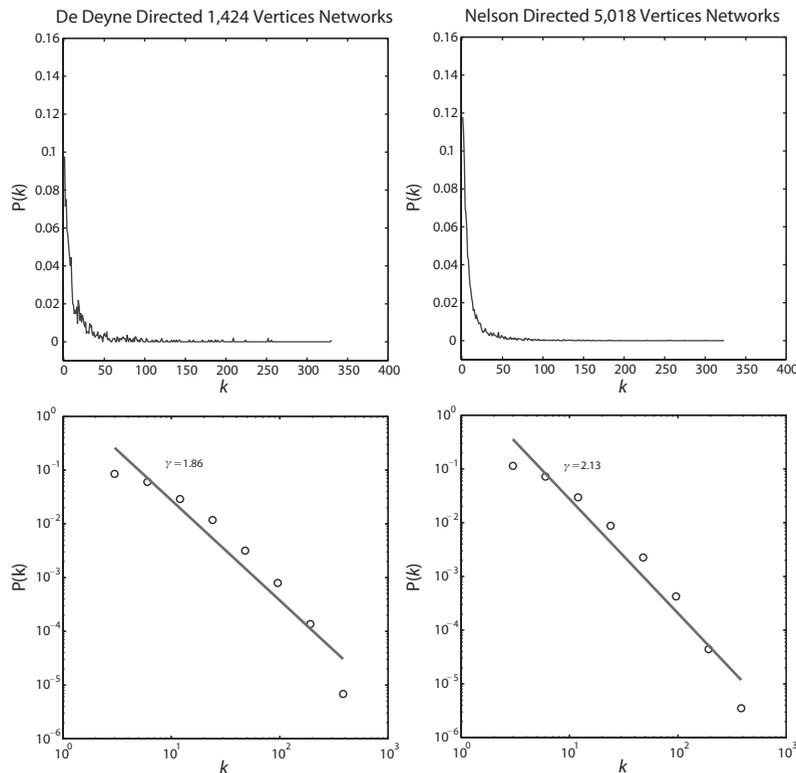


Figure 4. The degree distribution for the directed network in log–log coordinates and best-fitting power-law distribution regression line for the present (Leuven) and Nelson et al. (2004) association networks.

made, the number of choices of adjectives and verbs is limited in comparison with that of the nouns. This can explain the hub qualities of the adjectives. On the other hand, a different explanation might be needed to interpret the findings for the verbs, since members of this word class did not occur as hubs, although the size of the word class is also considerably smaller than that of nouns.

In the following section, the cue centrality statistics of the network are compared with measures of word availability based on word frequency, AoA, and imageability. A different but complementary analysis of the relationship between these word availability measures and the frequency of association responses can be found in the companion article (De Deyne & Storms, 2008).

CUE CENTRALITY

In this section, we investigate the relationship between centrality and three language utility properties of the nodes: word frequency, AoA, and imageability.

Theories of word representation hold contrasting views on which words are central in networks. A first account of the organization of semantic networks is based upon the abstract-versus-concrete distinction for words. For instance, Lambert (1955) found, for both French and English, that more responses to concrete nouns were available than responses to abstract ones. However, these results are not in line with the predictions from network models of memory

(Schwanenflugel & Shoben, 1983; Wattenmaker & Shoben, 1987). In these memory models, concrete or high-imageability words and abstract or low-imageability words are differentially represented in such a way that abstract concepts contain more information than concrete concepts. In other words, nodes representing abstract concepts have a higher number of connections and are thus more central. This view is supported by the observation that abstract concepts occur in a greater variety of contexts than do concrete concepts (Galbraith & Underwood, 1973). Further empirical support comes from de Groot (1989), who asked participants to generate word associations for high- and low-frequency words and for high- and low-imageability words. It was found that word imageability exerted a strong influence on word association, whereas the effect of word frequency was negligibly small. Associations to concrete words were made more quickly than those to abstract words. Furthermore, the association frequency of the responses to concrete words was larger, whereas the response heterogeneity was smaller. In other words, concrete words had fewer connections. In summary, different theories have given opposing views on the processing advantage of concrete over abstract words. However, regardless of which theory is correct, word concreteness or imageability might not be the only determinant of the organization of semantic knowledge.

A different account is given by van Loon-Vervoorn (1989), who suggested that the order of acquisition is the most important organizational principle in the semantic

system. Meanings of words acquired later are built upon the meanings of words acquired earlier. This conclusion was based on a timed association task in which van Loon-Vervoorn found that words acquired early had faster (RT = 1,440 msec) associations than words acquired later (RT = 1,681 msec). Although a similar effect was found for highly imageable words (RT = 1,445 msec) in comparison with abstract words (RT = 1,677 msec), this effect was not as strong. In addition, possible effects of word frequency were investigated but found not to be significant. Further support for the idea that the AoA of concepts reflects the centrality of words acquired early has been found in numerous studies (Brybaert, Van Wijnendaele, & De Deyne, 2000; De Deyne & Storms, 2007; Steyvers & Tenenbaum, 2005). Many of these accounts have been based on relatively straightforward measures for cue centrality, such as the number of different associations or the strength of the primary response. In the following section, cue centrality measures are used to test contrasting views on word availability in directed and undirected networks. First, we investigate the extent to which these language utility properties correspond to connectivity properties in the association network, using a number of different centrality measures. Next, a qualitative comparison is made for the different centrality measures for the highest connected nodes or hubs, followed by a brief comparison with hubs from a similar network based on English word associations.

Method

Materials and Procedure. The networks used in this investigation were identical to the directed network and the undirected network introduced in the previous section, each consisting of the 1,424 nodes. Calculated for each vertex were degree, clustering coefficients, closeness, and betweenness. Word frequency measures were selected from the Dutch CELEX (Baayen et al., 1993). AoA ratings were compiled from Ghyselinck, De Moor, and Brybaert (2000), Ghyselinck, Custers, and Brybaert (2003), Ruts et al. (2004), and De Deyne and Storms (2007) and measured the subjective rating (in years) of the age at which a word is learned. Imageability ratings were taken from Ghyselinck et al. (2000), van Loon-Vervoorn (1985), and newly collected ratings. The newly collected imageability judgments were ratings on a 7-point scale and are described in the companion article (De Deyne & Storms, 2008).

Results and Discussion

The resulting set, in which values for word frequency, AoA, and imageability overlapped, consisted of 1,117 vertices. The results are shown in Table 2. Note that the degree

for the undirected network is the sum of the in-degree and the out-degree. The closeness centrality measure makes the same distinction, thus providing additional columns in the upper panel of Table 2. The most important differences between the directed network and the undirected network are the different results for the in-degree and out-degree in the directed network, and the degree of the undirected network.

There was a strong correlation between the number of incoming arcs and the word availability measures of word frequency, AoA, and imageability. However, the negative correlation for the out-degree showed that a large number of outgoing arcs was weakly associated with lower values of word availability. Further elaborating on these findings, it becomes clear that cues with many different associates—that is, cues with a high out-degree—tended to have negative correlations, because there was a low agreement in responses. Summing the in-degree and out-degree dampens the correlations for the undirected network. A similar effect could be found for the closeness centrality. These findings correspond with earlier reports that found either no relationship or an inverse relationship between response heterogeneity and word frequency (Cramer, 1965; Postman, 1964). The findings also correspond to Nelson and McEvoy (2000), who found that the connections that common words have from other words tend to be more numerous than the connections they make to other words. The clustering coefficients⁴ were only calculated for the simple case of undirected networks. The clustering based on direct neighbors in CC1' shows that most dense clustering nodes were concrete and early acquired, although the magnitude of this correlation was weak. The size of the neighborhood increased greatly when neighbors of neighbors were considered (CC2'). This centrality statistic closely resembled the results of the betweenness statistic.

The lemma and word-frequency correlations were similar in magnitude to each other, as well as to the AoA correlations to most of the centrality measures. However, the low correlations for imageability did not support the hypothesis that word concreteness determines the network's structure (de Groot, 1989).

When looking at the relation between word measures and the different centrality measures in the association network, a clear division between word frequency measures and AoA on the one hand and word imageability on the other is observable. Moderate correlations with the

Table 2
Correlation Between Network Centrality Measures for the Directed and Undirected Networks— k , k^{in} , k^{out} , Clustering Coefficients in Neighborhood 1 (CC1) and Neighborhood 2 (CC2), Betweenness (b), Closeness for Directed Networks (c^{in} , c^{out})—and Word Accessibility Measures

Measure	Directed Network					Undirected Network				
	k^{in}	k^{out}	c^{in}	c^{out}	b	k	c	b	CC1	CC2
WF	.70**	-.14**	.61**	-.21**	.61**	.48**	.16**	.50**	.02	.46**
LF	.70**	-.11**	.61**	-.18**	.61**	.49**	.19**	.50**	.03	.46**
IMA	.30**	-.01	.35**	-.05**	.21**	.13**	.08**	.10**	.08**	.13**
AoA	-.64**	-.02	-.61**	.03	-.55**	-.47**	-.30**	-.43**	-.12**	-.43**

Note— $N = 1,130$. k , k^{in} , k^{out} , and b were log-transformed to reduce the effect of skew in the data. WF, CELEX log-transformed word frequency; LF, CELEX log-transformed lemma frequency; IMA, imageability; AoA, log-transformed age of acquisition. **Significant at the .01 level (two-tailed).

centrality measures were found for word frequency and AoA, but not for imageability. Although the correlations point toward a higher degree of centrality for early words, this effect also occurred for highly frequent words. In conclusion, the central network nodes tended to be early acquired and highly frequent, but were not necessarily associated with mental imagery. This supports the theoretical proposals by van Loon-Vervoorn (1989) and Brysbaert, Van Wijnendaele, and De Deyne (2000) and corroborates the findings of Steyvers and Tenenbaum (2005), but it contradicts de Groot's (1989) findings of imageability effects in the absence of word frequency effects. Notably, the strongest correlation between imageability and measures of node centrality was found for the node in-degree ($r = .30$). Nodes that frequently occurred as an association response (i.e., those with many incoming arcs) tend to be more concrete. However, the absence of a significant correlation of the out-degree also shows that word imageability does not imply that a concrete word has a larger number of different associations. Differentiating between incoming and outgoing arcs provides evidence against the claim that concrete words have more information associated with them than abstract words do (de Groot, 1989).

The global account given above does not provide any specific information about which nodes were central in the network. In addition, it did not show how measures of centrality differ. To be able to appreciate which words correspond to the central nodes, the 10 most central nodes were calculated according to previously introduced measures of centrality. Table 3 shows the results for in-degree and closeness from a directed network, since the findings reported above indicate that the interpretation based on node in-degree provided the clearest case. Since the cluster and betweenness coefficients have an easy interpretation only for undirected networks, CC1 and CC2 calculations were restricted to the undirected network.

Despite the different ways in which centrality can be formalized, there was considerable overlap concerning which nodes were most central. Many of these hubs were adjectives or corresponded to an ontological category such as *animal*. Without doubt, this topic is worthy of an entire study in itself, and could benefit from additional behavioral data.

To investigate how these results relate to the English association norms in the Florida network, the highest connected nodes were calculated in the undirected network. The 10 words with the highest degrees were *food, money, water, car, good, bad, work, school, house, and love*. It is interesting to see that some hubs correspond in English and Dutch, but it is also clear that there are important differences. Clearly, the identification of universal hubs across languages is a topic that needs further investigation.

SEMANTIC PROPERTIES

What qualifies as a semantic relationship has been the focus of much debate. Some researchers (e.g., Joordens & Becker, 1997) propose that only word pairs such as *dog–cat* are semantically related, whereas words such as *dog–house* are not. According to this view, words are semantically related if they significantly overlap in terms of their physical properties. At the other side of the spectrum, words are considered to be meaningfully related if they share similar contexts (e.g., Barsalou et al., in press; Buchanan, Westbury, & Burgess, 2001). Most likely, both types of semantic relationships are prevalent in tasks that involve the processing of words. This view closely resembles the idea behind the LASS theory (Barsalou et al., in press). In order to further study what qualifies as a semantic relationship, we investigated the semantic relationships that appear in word associations and studied their relative occurrence. Such an analysis shows the extent to which associations involve physical–conceptual semantic relation types or contextual semantic relations. For this purpose, semantic properties of the associations were coded using a taxonomy previously introduced by McRae and Cree (2002), which is an adaptation of a scheme used by Wu and Barsalou (2007). The coding scheme has a taxonomic structure that consists of four feature classes: entity features, situation features, taxonomic categories, and introspective features. Since this taxonomy is very fine-grained and hierarchical, it allows us to investigate small differences, as well as more general differences, by collapsing feature classes in one branch of the taxonomy.

Table 3
Overview of the 10 Most Central Network Nodes

Node	k^{in}	Node	c^{in}	Node	CC1'	Node	CC2'	Node	b
food	332	food	516	food	48	food	150	food	48
white	261	white	514	tasty	47	tasty	80	white	47
tasty	257	pretty	514	sweet	39	water	46	red	38
water	253	red	509	fruit	39	white	29	pretty	33
red	230	warm	504	orchestra	38	pretty	26	black	33
pretty	215	summer	495	weapon	37	red	22	water	32
black	214	water	495	cake	36	summer	21	pain	26
hot	199	black	493	fruit*	35	hot	21	hot	25
green	198	green	492	dessert	35	green	20	brown	24
summer	192	brown	486	summer	35	animal	18	green	22

Note—Centrality was measured as the network in-degree (k^{in}), closeness (c^{in} based on incoming connections), corrected clustering coefficient within Neighborhood 1 (CC1'), corrected clustering coefficient within Neighborhood 2 (CC2'), and betweenness (b). For ease of reading, c^{in} , CC1', CC2', and b are multiplied by 1,000. *This entry refers to the Dutch *vrucht*, which has a similar meaning to the Dutch *fruit* but refers to the part of a plant containing seeds.

Two questions are of particular interest in this study. First, the taxonomy should give a detailed account of the kinds and importance of semantic entities and relations given as association responses. Second, the taxonomic coding should allow us to identify qualitatively different processes underlying the generation of primary, secondary, and tertiary association responses. The general hypothesis for these process differences is that later responses are less automatic and more elaborate, indicating a contribution from the situated conceptual representations that is not available in the associative–semantic network. Under the assumption that the linguistic system produces responses faster than the simulation system, we expect to find relatively more linguistic responses for the primary associations than for the secondary and tertiary associations. According to LASS (Barsalou et al., in press), situation and entity responses should be statistically more likely to originate from describing situated simulations than from retrieving linguistic forms. The LASS theory is less clear about introspective and taxonomic features. Most often, taxonomic categories are viewed as a part of the conceptual system. However, taxonomic category labels and their subordinates are well imbedded in language as collocates (e.g., an apple is a fruit). Therefore, taxonomic properties could be part of the linguistic system or the situated system (Barsalou et al., in press). Introspective properties involve situated representations to some extent, but they differ in that they contain affective and evaluative content.

Method

Materials and Procedure. Since the entire set of association cues was too large to code exhaustively, a subset of 458 items was chosen. The items were chosen to form a representative set of concepts from a wide range of categories. The coded set consisted of two activity categories (*professions* [32] and *sports* [30]), four food categories (general *food* and *drinks* [14], *fruit* [31], and *vegetables* [32]), five animal categories (*insects* [26], *reptiles* [20], *birds* [30], *mammals* [34], and *fish* [22]), and six artifact categories (*kitchen utensils* [29], *clothing* [29], *musical instruments* [27], *vehicles* [29], *weapons* [17], and *tools* [27]). This set corresponded to the concepts in Ruts et al. (2004), with the additional categories *weapons*, *kitchen utensils*, *clothing*, *food*, and *drinks*. Besides these basic categories, 31 words corresponding to superordinate concepts, such as *fruit* or *insect*, completed our set.

The words were limited to nouns and verbs, since it is unclear how the existing taxonomies provide a semantic ontology for adjectives. The coding was completed during a first round by 2 independent coders. In a second round, diverging codes were solved through discussion. All cue–response pairs had at least one and at most three codes, since the semantic features do not always indicate exclusive categories. For example, the Dutch cue–response pair *appel* <eten> (*apple* <food/to eat>) had two codes: superordinate and situated action, corresponding to the two senses in Dutch of the associate <eten>: food and to eat. The coding definitions and examples for each coding instance can be found in the Appendix. Some small adaptations had to be made to make the Wu and Barsalou (2007) taxonomy appropriate for associations.⁵ Two particular types of responses that were obtained through word associations but that do not occur much in feature data are word completion and rhymes. To account for these phenomena and to have a taxonomy with complete coverage of the data, a fifth class of linguistic features was introduced. This class consisted of forward compound continuations (<sun> *set*), backward compound continuations (*star* <fruit>), word fragments (*swordfish* <word>), words with similar orthography (*worm* <dorm>), and mediated relations. The

mediation could be a virtue of a shared property that connects to otherwise completely different concepts (*apricot* <tiger>) or a mediating concept (*beaver* <toothpaste> mediated through <teeth>).

Results and Discussion

A total of 39,359 unique response pairs, which amounts to 30% of the total response pairs, were coded. First, the general distribution of the entity, situation, taxonomic, introspective, and lexical properties will be addressed regardless of the specific subtypes of these properties. For each of these types, the proportion of codes for the indicated cue–response pairs was calculated. Next, separate analyses were performed on the subtypes of each property type.

The distribution of the main property types is shown in Figure 5. The figure illustrates two main findings. First, it shows that the majority of codes for a total count of 128,929 association responses refer to situation properties. Second, it indicates that taxonomic responses are most often given as first responses. The conceptual content, coded in the entity and situational property types, becomes more available during second and third responses. Further evidence for the hypothesis of superficial processing comes from the lexical feature distribution, which drops off in the second and third responses. The corresponding log-linear model that corroborates these claims yields a saturated model ($G^2 = 0$, $df = 0$). To see whether the interaction between response position and feature type also occurred in the less frequent introspective and lexical features, a similar response position \times feature type log-linear model was composed for a subset consisting of introspective and lexical features only. Again, only the saturated model gave an adequate fit of the data.

For each of the basic semantic features, the subfeature distributions are shown in Figures 6, 7, 8, 9, and 10. Figure 6 shows the distribution of the entity properties. This figure shows that visual information about color, shape, and texture is very prominent regardless of whether the response was the first, second or third association.

Figure 7 shows that the most frequent taxonomic responses were superordinates and coordinates. This is not surprising given the fact that most of the concepts are basic-level concepts. For the superordinate concepts that were included in the materials, the result obviously deviated from this pattern and the most frequent taxonomic responses were subordinates. This pattern indicates that unmarked information, such as superordinates, is readily evoked when participants are asked to give associations. The number of synonyms and antonyms was a lot smaller than the numbers of super- and subordinate responses. Although the present coding protocol was not designed to code for adjectives, previous research by Deese (1965) showed that these two semantic types are the modal response for adjectives and would occur more frequently when more adjectives are represented in the data set.

Figure 8 contains the distribution of the situation properties. These properties represent the largest portion of the association responses. Responses referring to information of concrete objects especially, but also responses regarding locations embedded in a specific context, are often generated.

The distribution of the lexical codes in Figure 9 shows a small number of mediated responses. Assuming that

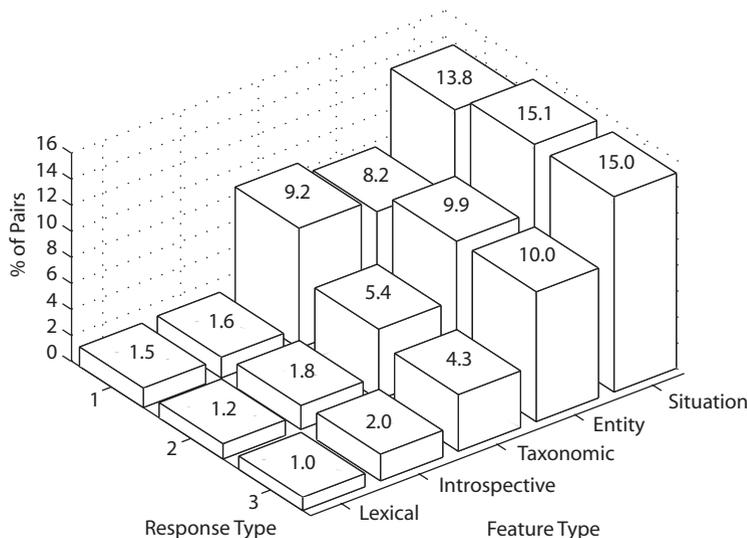


Figure 5. Three-dimensional bar chart with the percentages of each of the main features in the adapted Wu and Barsalou (2007) taxonomy.

all mediated responses indicate chaining in the association data (a very unlikely assumption, since mediated responses can occur in discrete tasks as well), the present data indicate that this is of little concern. Only 2% of the first, second, and third responses were related to the cue through a mediating concept or property. This accounts for 0.24% of the complete coded data.

The taxonomy of the introspective features in Figure 10 includes many types that code for information typically

found in word feature generation tasks (e.g., McRae & Cree, 2002), such as representational states, contingencies, cognitive operations, and negations. Due to the nature of the task, these types hardly occur at all in the association responses. The introspective responses largely consist of evaluations and emotion responses. The pattern of their distributions is on a par with the conceptual branches of the taxonomy that code for entity and situation information, and suggests the recruitment of additional sources from memory.

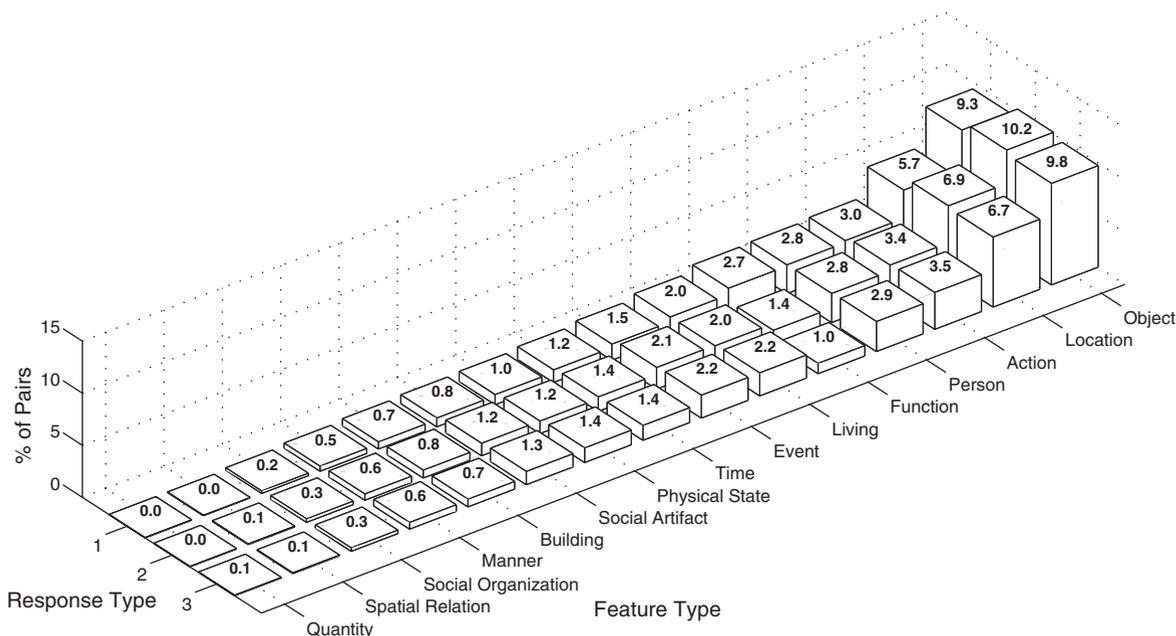


Figure 6. Bar chart with the percentages of occurrences of situation properties in the association responses.

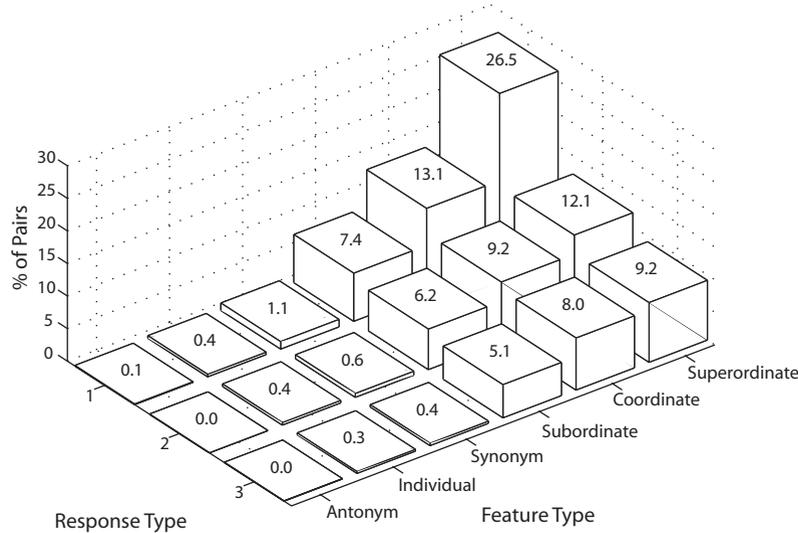


Figure 7. Bar chart with the percentages of occurrences of taxonomic properties in the association responses.

The global picture that emerges indicates that the continuous association task adds response types that are traditionally more attributed to the retrieval of information from conceptual representations. However, this evidence does not allow any dissociation between a semantic-associative system, in which words are related due to their distributional properties in language, and a conceptual system, in which word meaning and associates are mediated by retrieval of sensory representations. Nonetheless, the present data do provide strong evidence for readily available information about category membership. The data allow specific tests to investigate this kind of proposals. For instance, if language contingencies shape the semantic-associative system, the likelihood of certain cue pairs can be estimated and compared with the likelihood

of these pairs in the association data. The largest dissociations between the “language model” and the association model must be due to some conceptual involvement. This account remains speculative for now, but the tools and data to investigate its merits are accessible.

GENERAL DISCUSSION

In this article, word associations were studied on a number of different levels. What differentiates this research from previous studies that used large word association data sets is the specific procedure for generating responses (De Deyne & Storms, 2008). In contrast to the association data set by Nelson et al. (2004), our associations were gathered in a continuous task. The nature of

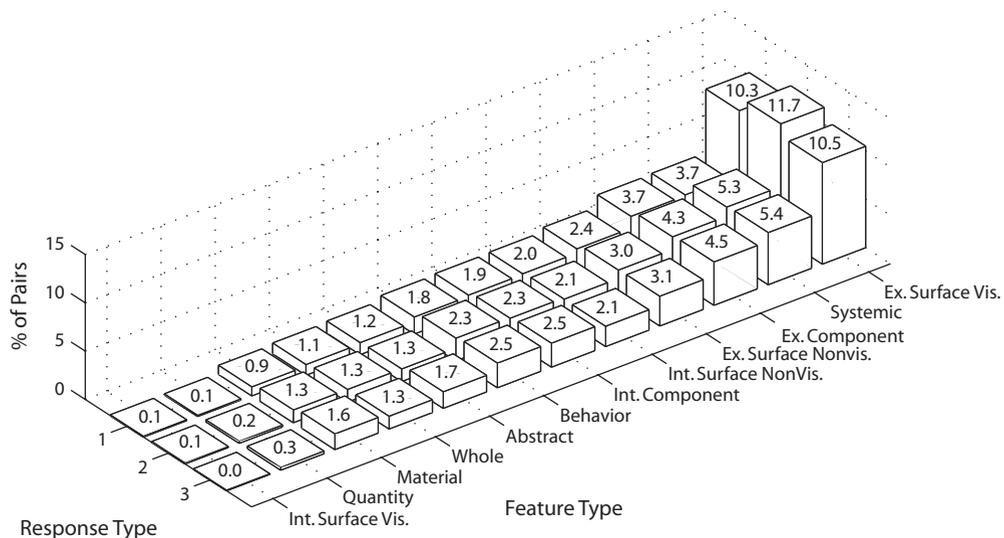


Figure 8. Bar chart with the percentages of occurrences of entity properties in the association responses.

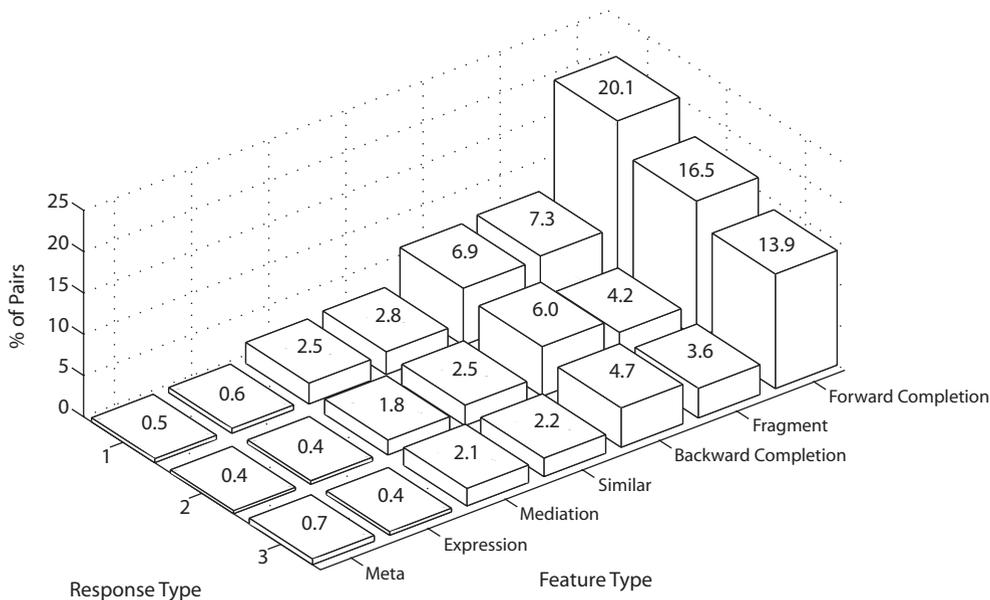


Figure 9. Bar chart with the percentages of occurrences of lexical properties in the association responses.

our task allowed the testing of a number of hypotheses regarding the response characteristics of associations that are given early or late.

First, the idea that association responses are predominantly paradigmatic was refined. In fact, the syntagmatic responses were the modal response type for adjectives and verbs but not for nouns. Furthermore, the paradigmatic response tendency for nouns showed a decrease in later association responses.

Second, it was shown that the secondary and tertiary responses add extra information, which makes a distributional account of meaning based on the aggregated re-

sponses a good candidate for a high-dimensional spatial model of meaning such as HAL (Burgess & Lund, 1997) or LSA (Landauer & Dumais, 1997). The continuous task led to a denser network representation as measured by the average distance between two arbitrary points in the network. The average distance between two concepts was four edges in an undirected network, whereas the corresponding network based on Nelson et al. (2004) needed an average of five edges to connect two arbitrary nodes. Furthermore, the large-scale analysis of the semantic network indicated that the connection pattern between nodes corresponds to a small-world structure. These results confirm

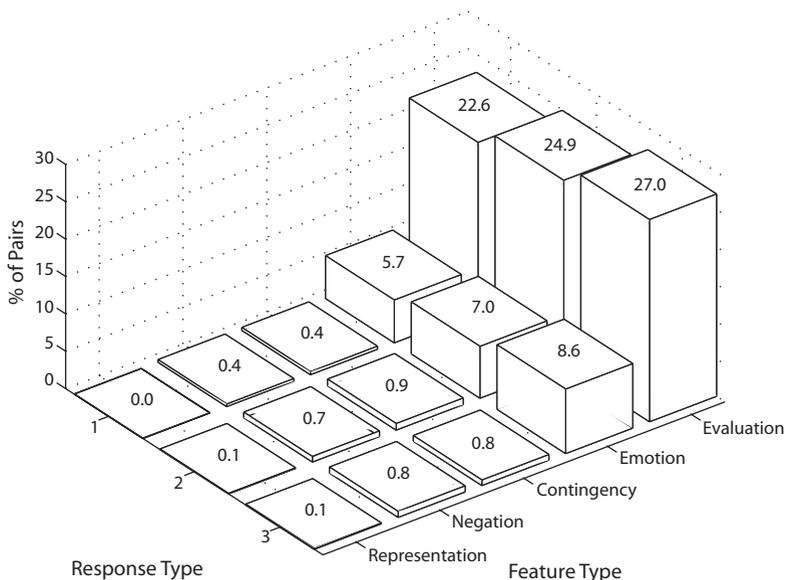


Figure 10. Bar chart with the percentages of occurrences of introspective properties in the association responses.

earlier findings by Steyvers and Tenenbaum (2005) based on the English data of Nelson et al. Finally, and perhaps most crucially, the continuous task showed a different distribution of semantic content types in the responses. This result fits the prediction of the model of word association proposed by de Groot (1980), according to which two processes determine the generation of word associations: (1) a fast process during which reactions are retrieved between directly connected words, without the need for retrieval of the meaning of the stimulus word; and (2) a relatively slow process during which no use is made of automatic connections between word associations, but which depends on a meaningful interpretation of the stimulus word.

The theory of de Groot (1980) shows a striking resemblance to the predictions of the LASS theory (Barsalou et al., in press). According to LASS, certain types of information become activated very quickly, and this process can be attributed to distributional properties in language, such as their co-occurrence. Other information becomes available after extended processing, but this aspect requires the situated simulation of conceptual properties of the cue. Despite the fact that LASS attributes certain types of semantic relations to different systems, perhaps instantiated in different neural substrates, it is hard to prove such claims (which localize the source of semantic relationships) on the basis of these data alone. However, at least one study (Simmons, Hamann, Harenski, Hu, & Barsalou, 2007) has used fMRI measurements during a property-generation task and an association task and found a correspondence between the activation of linguistic areas such as Broca's area for early responses and areas such as the precuneus, which are often associated with mental imagery for late responses. Other studies have used ERP measurements to investigate the time course of associations and have demonstrated that less frequent associations typically involve different neuroanatomical areas located in the right hemisphere than do frequent associations (Abdullaev & Posner, 1997). Evidently, there are still aspects that can differentiate first, second, and third responses, but, on the basis of the findings described above, researchers interested in using these association norms can make more informed decisions on whether to use only primary association response data (i.e., based on first associates only) or the aggregated data (i.e., based on first, second, and third associates). Furthermore, alternative theories to dual coding theories have also been proposed (e.g., the relational/distinctiveness processing theory—Marschark & Hunt, 1989). Although these accounts offer some predictions about language utility measures such as imageability or concreteness, they often do not make any predictions about the semantic composition of association responses, and further treatment would be beyond the scope of this article.

The study of the network topology confirmed the findings of the small-world network structure of Steyvers and Tenenbaum (2005), despite the fact that the Dutch association network was not as extensive as the network based on the Nelson et al. (2004) norms. Furthermore, as in Steyvers and Tenenbaum, the high-density nodes, called *hubs*, are often adjectives. Although both networks exhibit small-world properties, there are some deviations between the

two networks that are likely due to language-specific properties and to the choice of materials and procedures. The networks that are presented here might also deviate from the neurological substrates that instantiate these properties in the brain. Even though small-world properties have been found in many growing systems and there is some evidence for them in brain networks (Sporns, Chialvo, Kaiser, & Hilgetag, 2004), it is not clear if these findings can be generalized to all neural structures that encode meaning at the conceptual or semantic level.⁶ Allowing some speculation, a more functional view of an associative network with hardwired small-world property constraints could prove a useful quick-and-dirty heuristic for the retrieval of conceptual information. This might provide an interesting solution to the omnipresent frame problem in the study of conceptual representations. When representation of concepts is seen as the representation of similarity between concepts, it is unclear how to restrict the representations that are needed to represent this similarity. For example, it is likely that the concept of an apple is defined by its similarity to pears or other fruit members, but it is not clear how this representation can be bounded so as not to include cars or lightbulbs. If associations reflect the contingencies in the world through language, organized according to a certain type of network architecture such as a small-network architecture, this might provide the structural bounds for selecting information and presenting a bottleneck for the activation of conceptual information used in recognition or categorization. Furthermore, it has been suggested that these network architectures are dynamically advantageous because they are more synchronizable or error tolerant (Strogatz, 2001). Finally, the finding of power-law degree distributions in semantic networks is significant because it indicates that a small number of words appear as associates for a great variety of cues. This finding has theoretical implications for spatial models of meaning, since this kind of phenomenon is difficult to produce in spatial representations (Griffiths, Steyvers, & Tenenbaum, 2007; Steyvers & Tenenbaum, 2005).

Our study of cue centrality showed that, regardless of how this centrality is measured, the time at which a word is acquired (i.e., its AoA) is strongly related to the central position in the network. Previous research with neuropsychological patients has shown that AoA is the best predictor of performance in patients who exhibit a loss of knowledge about word meaning (Bell, Davies, Hermann, & Walters, 2000; Kay & Hanley, 1999; Ukita, Abe, & Yamada, 1999). In line with previous accounts by Brysbaert et al. (2000) and Steyvers and Tenenbaum (2005), these observations support a mechanistic explanation of how networks grow over time and how this affects later retrieval.

Our findings have also shown that cue centrality is not strongly related to the degree of concreteness or imageability, a result that does not support the findings by de Groot (1989) and others who have explained imageability or concreteness effects in terms network connectedness (e.g., Schwanenflugel & Shoben, 1983). In contrast to imageability, node centrality was related to word frequency and AoA. The magnitude of this relationship was similar for both measures of word availability. In order to disentangle

these effects, more specific tests are needed. One possibility would be to investigate differences among verbs, adjectives, and nouns, since a similar analysis of response availability in a continuous word association task indicated that the association frequency of verbs and adjectives is related to the imageability and AoA of these responses, but not to word frequency (De Deyne and Storms, 2008).

The semantic analysis of the word associations revealed that certain knowledge types are generated early in a continuous association task. In particular, the fact that taxonomic information (and especially superordinate information) becomes available more quickly than conceptual information (i.e., entity or situation properties) puts into question the involvement of conceptual information in tasks such as semantic categorization. This is in accordance with experimental evidence from speeded categorization tasks with items from natural categories, in which it was shown that speeded categorization involves both the retrieval of associations and similarity-based concept representations (Hampton, 1997).

The semantic taxonomy was expanded to allow the identification of certain cue–response relationships typically found in word associations such as mediated responses. If many of the responses are based on a mediated relationship between cue and association, this would show that the task is largely affected by chaining (Nelson et al., 2004). Although chaining might be considered undesirable for certain studies, our findings indicated that this is of little concern when only three responses per cue are elicited. When a large sample of cues and response relationships were coded, it was found that a maximum of 0.24% of the data could have been affected by chaining. The semantic taxonomy also functions as a guide to judge comparable models. For instance, when a high-dimensional semantic model based on concept features is compared with a similar model based on word associations, it is hard to interpret different performances of these models. What sets semantic representations based on text collocates such as LSA (Landauer & Dumais, 1997) and HAL (Burgess & Lund, 1997) apart from models such as the Word Association Space (Steyvers et al., 2004) or models based on concept properties such as the featural and unitary semantic space (FUSS; Vigliocco, Vinson, Lewis, & Garrett, 2004) is the distribution of certain types of semantic properties. Meaningful interpretations depend on an understanding of the meaning of the model content.

For instance, if word associations simply stem from the learner's sensitivity to the statistical properties of language, then this can be modeled by co-occurrence models such as LSA or HAL. However, when different kinds of semantic information are distinguished, a different view can emerge. For instance, responses referring to visual properties (e.g., color, shape) occurred quite frequently. It is at least questionable whether these types of properties are similarly represented as co-occurrences in language. The answer to this question could provide evidence for competing views, such as the dual coding theory or co-occurrence models that do not recruit conceptual representations. Further research on this topic is currently being undertaken.

The present semantic analysis has touched upon only a few aspects of the nature of word associations. We hope this might lead to a better appreciation of the nature of word associations, yet it must be acknowledged that the presented analyses do not provide complete coverage of the potential offered by the data set. For instance, it does not describe category-specific differences for the association responses. Such an analysis (based on semantic features) has previously been performed to study category-specific deficits (Cree & McRae, 2003) and could be expanded to incorporate word associations as well. Furthermore, the in-depth treatment of topics such as the semantic composition of word associations would be beyond the scope of this article. However, it would be interesting to further investigate how cue availability properties influence the semantic types that are generated. For instance, on the basis of these data it is possible to test the hypothesis that low-frequency cues elicit more superordinate responses than high-frequency instances.

AUTHOR NOTE

This work was supported by Grant OT/01/15 from the Leuven University Council and Grant G.02666.02 from the Belgian National Science Foundation, awarded to the second author. Correspondence concerning this article should be addressed to S. De Deyne, Departement Psychologie, Hogere Cognitie, Tiensestraat 102, 3000 Leuven, Belgium (e-mail: simon.dedeyne@psy.kuleuven.be).

REFERENCES

- ABDULLAEV, Y. G., & POSNER, M. I. (1997). Time course of activating brain areas in generating verbal associations. *Psychological Science*, *8*, 56-59.
- ALBERT, R., JEONG, H., & BARABÁSI, A.-L. (1999). Internet: Diameter of the World-Wide Web. *Nature*, *401*, 130-131.
- ALLPORT, D. A. (1985). Distributed memory, modular subsystems and dysphasia. In S. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia* (pp. 32-60). Edinburgh: Churchill Livingstone.
- AMARAL, L. A., SCALA, A., BARTHELEMY, M., & STANLEY, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences*, *97*, 11149-11152.
- BAAYEN, R., PIEPENBROCK, R., & VAN RIJN, H. (1993). *The CELEX lexical database* [CD-ROM]. Philadelphia: University of Pennsylvania, Linguistic Data Consortium.
- BARABÁSI, A.-L., & ALBERT, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509-512.
- BARSALOU, L. W. (1999). Perceptual symbol systems. *Behavioral & Brain Sciences*, *22*, 577-660.
- BARSALOU, L. W., SANTOS, A., SIMMONS, W., & WILSON, C. D. (in press). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg, & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning*. Oxford: Oxford University Press.
- BELL, B. D., DAVIES, K. G., HERMANN, B. P., & WALTERS, G. (2000). Confrontation naming after anterior temporal lobectomy is related to age of acquisition of the object names. *Neuropsychologia*, *38*, 83-92.
- BRYBAERT, M., VAN WIJNENDEAELE, I., & DE DEYNE, S. (2000). Age-of-acquisition effects in semantic processing tasks. *Acta Psychologica*, *104*, 215-226.
- BUCHANAN, L., WESTBURY, C., & BURGESS, C. (2001). Characterizing the neighborhood: Semantic neighborhood effects in lexical decision and naming. *Psychonomic Bulletin & Review*, *8*, 531-544.
- BURGESS, C., & LUND, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language & Cognitive Processes*, *12*, 177-210.
- CRAMER, P. (1965). Response entropy as a function of the affective quality of the stimulus. *Psychonomic Science*, *3*, 347-348.
- CRAMER, P. (1968). *Word association*. New York: Academic Press.

- CREE, G. S., & MCRÆ, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of *chipmunk*, *cherry*, *chisel*, *cheese*, and *cello* (and many other such concrete nouns). *Journal of Experimental Psychology: General*, **132**, 163-201.
- DE DEYNE, S., & STORMS, G. (2007). Age-of-acquisition differences in young and older adults affect latencies in lexical decision and semantic categorization. *Acta Psychologica*, **124**, 274-295.
- DE DEYNE, S., & STORMS, G. (2008). Word associations: Norms for 1,424 Dutch words in a continuous task. *Behavior Research Methods*, **40**, 198-205.
- DEESE, J. (1962). Form class and the determinants of association. *Journal of Verbal Learning & Verbal Behavior*, **1**, 79-84.
- DEESE, J. (1965). *The structure of associations in language and thought*. Baltimore: Johns Hopkins Press.
- DE GROOT, A. M. B. (1980). *Mondelinge woordassociatienormen: 100 woordassociaties op 460 Nederlandse zelfstandige naamwoorden* [Oral word association norms: 100 word associations to 460 Dutch nouns]. Lisse, The Netherlands: Swets & Zeitlinger.
- DE GROOT, A. M. B. (1989). Representational aspects of word imageability and word frequency as assessed through word association. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 824-845.
- FREEMAN, L. (1977). A set of measures of centrality based on betweenness. *Sociometry*, **40**, 35-41.
- GALBRAITH, R., & UNDERWOOD, B. (1973). Perceived frequency of concrete and abstract words. *Memory & Cognition*, **1**, 56-60.
- GHYSELINCK, M., CUSTERS, R., & BRYLSBAERT, M. (2003). Age-of-acquisition ratings for 2322 Dutch words from 49 different semantic categories. *Psychologica Belgica*, **43**, 181-214.
- GHYSELINCK, M., DE MOOR, W., & BRYLSBAERT, M. (2000). Age-of-acquisition ratings for 2816 Dutch four- and five-letter nouns. *Psychologica Belgica*, **40**, 77-98.
- GLANZER, M. (1962). Grammatical category: A rote learning and word association analysis. *Journal of Verbal Learning & Verbal Behavior*, **1**, 31-41.
- GLENBERG, A. (1997). What memory is for. *Behavioral & Brain Sciences*, **20**, 1-55.
- GRIFFITHS, T. L., STEYVERS, M., & TENENBAUM, J. B. (2007). Topics in semantic representation. *Psychological Review*, **114**, 211-244.
- HAMPTON, J. (1997). Associative and similarity-based processes in categorization decisions. *Memory & Cognition*, **25**, 625-640.
- JEONG, H., TOMBOR, B., ALBERT, R., OLTVAI, Z. N., & BARABÁSI, A.-L. (2000). The large-scale organization of metabolic networks. *Nature*, **407**, 651-654.
- JOORDENS, S., & BECKER, S. (1997). The long and short of semantic priming effects in lexical decision. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 1083-1105.
- KAY, J., & HANLEY, J. R. (1999). Person-specific knowledge and knowledge of biological categories. *Cognitive Neuropsychology*, **16**, 171-180.
- LAMBERT, W. (1955). Associational fluency as a function of stimulus abstractness. *Canadian Journal of Psychology*, **9**, 103-106.
- LANDAUER, T., & DUMAIS, S. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, **104**, 211-240.
- LILJEROS, F., EDLING, C. R., AMARAL, L. A. N., STANLEY, H. E., & ÅBERG, Y. (2001). The web of human sexual contacts. *Nature*, **411**, 907-908.
- LIVESAY, K., & BURGESS, C. (1999). The influence of verbal ability on mediated priming. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 314-319). Mahwah, NJ: Erlbaum.
- LOCKER, L., JR., SIMPSON, G. B., & YATES, M. (2003). Semantic neighborhood effects on the recognition of ambiguous words. *Memory & Cognition*, **31**, 505-515.
- LUPKER, S. (1984). Semantic priming without association: A second look. *Journal of Verbal Learning & Verbal Behavior*, **23**, 709-733.
- MARSCHARK, M., & HUNT, R. R. (1989). A reexamination of the role of imagery in learning and memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 710-720.
- MCRÆ, K., & CREE, G. S. (2002). Factors underlying category-specific semantic deficits. In E. M. E. Forde & G. W. Humphreys (Eds.), *Category-specificity in brain and mind* (pp. 211-249). Hove, U.K.: Psychology Press.
- NELSON, D. L., & MCEVOY, C. L. (2000). What is this thing called frequency? *Memory & Cognition*, **28**, 509-522.
- NELSON, D. L., MCEVOY, C. L., & SCHREIBER, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, **36**, 402-407.
- NELSON, D. L., ZHANG, N., & MCKINNEY, V. (2001). The ties that bind what is known to the recognition of what is new. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 1147-1159.
- NEWMAN, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, **45**, 167.
- PAIVIO, A. (1986). *Mental representations: A dual coding approach*. New York: Oxford University Press.
- POSTMAN, L. (1964). The acquisition and retention of consistent associative responses. *Journal of Experimental Psychology*, **67**, 183-190.
- REDNER, S. (1998). How popular is your paper? An empirical study of the citation distribution. *European Physical Journal*, **4B**, 131-134.
- RUTS, W., DE DEYNE, S., AMEEL, E., VANPAEMEL, W., VERBEEMEN, T., & STORMS, G. (2004). Dutch norm data for 13 semantic categories and 338 exemplars. *Behavior Research Methods, Instruments, & Computers*, **36**, 506-515.
- SCHWANENFLUGEL, P., & SHOBEEN, E. (1983). Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 82-102.
- SHELTON, J. R., & MARTIN, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 1191-1210.
- SIMMONS, W., HAMANN, S., HARENSKI, C., HU, X., & BARSALOU, L. (2007). *fMRI evidence for word association and situated simulation in conceptual processing*. Manuscript in preparation.
- SPORNS, O., CHIALVO, D. R., KAISER, M., & HILGETAG, C. C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, **8**, 418-425.
- 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.
- STEYVERS, M., & TENENBAUM, J. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cognitive Science*, **29**, 41-78.
- STROGATZ, S. H. (2001). Exploring complex networks. *Nature*, **410**, 268-276.
- UKITA, H., ABE, K., & YAMADA, J. (1999). Late acquired words in childhood are lost earlier in primary progressive aphasia. *Brain & Language*, **70**, 205-219.
- VAN LOON-VERVOORN, W. (1985). *Voorstelbaarheidswaarden van Nederlandse woorden: 4600 substantieven, 1000 verba, 500 adjectieven* [Imageability ratings of Dutch words: 4600 nouns, 1000 verbs, 500 adjectives]. Lisse, The Netherlands: Swets & Zeitlinger.
- VAN LOON-VERVOORN, W. (1989). *Eigenschappen van basiswoorden* [Properties of basic words]. Lisse, The Netherlands: Swets & Zeitlinger.
- VIGLIOCCO, G., VINSON, D. P., LEWIS, W., & GARRETT, M. F. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, **48**, 422-488.
- WATTENMAKER, W., & SHOBEEN, E. (1987). Context and the recallability of concrete and abstract sentences. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 140-150.
- WATTS, D. J., & STROGATZ, S. H. (1998). Collective dynamics of "small-world" networks. *Nature*, **393**, 440-442.
- WU, L., & BARSALOU, L. (2007). *Grounding concepts in perceptual simulation: I. Evidence from property generation*. Manuscript submitted for publication.

NOTES

1. Note that the gammas are comparable to those reported by Steyvers and Tenenbaum (2005) but are not completely identical. This is likely due to the small differences in procedure used to build the network, such as the choice of the cutoff for the distributions tail, which is needed to identify the power-law distribution.

2. Using the out-degree or the sum of out- and in-degrees would introduce a bias, since the out-degree depends on the task characteristics, such as the number of associations that were generated for each cue.

3. Through the use of logarithmically transformed bins, the skew in the regression function is reduced and can be expected to approach the real values better.

4. The formula that is used to calculate the clustering coefficients CC1, CC2 is the corrected version: CC1' and CC2'. This correction takes into account the number of edges for the node and normalizes the out-

come by multiplication by the degree divided by the maximum degree of the network. For nodes with a high degree (a large number of edges), the interconnections in Neighborhood 1 or 2 is likely to be smaller than those of nodes with a very low degree, and would have an artificially small clustering coefficient.

5. The coding scheme used here was based on a similar, but improved, scheme currently used by Larry Barsalou (personal communication, January 2006).

6. As noted by Steyvers and Tenenbaum (2005), this is especially the case for the scale-free structure.

APPENDIX

Wu and Barsalou's (2007) and McRae and Cree's (2002) Semantic Taxonomic Codings, Adapted With Additional Types for Associations

Entity Features: Properties of a concrete entity, either animate or inanimate. Besides being a single, self-contained object, an entity can be a coherent collection of objects (e.g., forest).

External component: A 3-D component of an entity that, at least to some extent, normally resides on its surface. Examples: coconut <has a shell>; tricycle <has pedals>

External surface feature (visual or nonvisual): An external feature of an entity that is not a component and that is perceived on or beyond the entity's surface, including shape, color, pattern, texture, size, touch, smell, and taste. An additional distinction is made between visual surface features (shape, color, texture, size) and nonvisual surface features (touch, smell, taste). Examples: apple <is red>; blender <is loud>

Internal component: A 3-D component of an entity that normally resides completely inside the closed surface of the entity. Examples: cherry <has a pit>; car <has an engine>

Internal surface feature (visual or nonvisual): An internal feature of an entity that is not a component, that is not normally perceived on the entity's exterior surface, and that is only perceived when the entity's interior surface is exposed, including color, pattern, texture, size, touch, smell, and taste. An additional distinction is made between visual surface features (shape, color, texture, size) and nonvisual surface features (touch, smell, taste). Examples: watermelon <is red>; honey <tastes sweet>

Behavior: A chronic behavior of an entity that is characteristic of its nature and that is described as a characteristic property of the entity, not as a specific intentional action in a situation. Examples: dog <barks>; clock <ticks>

Material: A specification of the materials or substance of which the entity is made. Examples: oak <made of wood>; sink <made of enamel>

Quantity: A numerosity, frequency, size, intensity, or typicality of an entity or its properties. Examples: giraffe <has a long neck>; slippers <come in pairs>

Associated abstract entity: An abstract property of the target entity not dependent on a particular situation. Examples: harp <angels>; teacher <democrat>

Systemic feature: A global systemic feature of an entity or its parts, including states, conditions, abilities, traits. Examples: dolphin <is intelligent>; car <is fast>

Larger whole: A whole to which an entity belongs. Examples: ant <lives in a colony>; basement <is part of a house>

Situation Features: Properties of a situation, where a situation typically includes one or more agents, at some place and time, engaging in an event, with one or more entities in various semantic roles. Examples: picnic, conversation, vacation, meal.

Function: A typical goal or role that an entity serves for an agent in a situation by virtue of its physical properties with respect to relevant actions; Examples: tomato <eaten>; bed <used for sleeping>

Action: A (nonintrospective) action that an agent (human or nonhuman) performs intentionally in a situation. Examples: strawberry <is picked>; shirt <wear>

Object: An inanimate object in a situation, except buildings. Examples: watermelon <plate>; cat <sofa>

Person: An individual person or multiple people in a situation. Examples: toy <children>; car <passenger>

Living thing: A living thing in a situation that is not a person, including other animals, plants, and body parts. Examples: sofa <cat>; park <grass>

Social organization: A social institution, a business, or a group of people or animals in a situation. Examples: picnic <family>; dog <pack>

Social artifact: A relatively abstract entity, sometimes partially physical (e.g., book) and sometimes completely conceptual, created in the context of sociocultural institutions. Examples: shark <a movie (about)>; invention <Nobel prize>

Building: A building in a situation. Examples: book <library>; candle <church>

Location: A place where an entity can be found or where people engage in an event or activity. Examples: zebra <Africa>; cupboard <kitchen>

APPENDIX (Continued)

Time: A time period associated with a situation or with one of its features. Examples: turkey <Thanksgiving>; cabin <vacation>

Event: A stand-alone event or activity in a situation in which the action is not foregrounded but is on a relatively equal par with the setting. Examples: car <trip>; church <wedding>

Manner: The manner in which an action or event is performed in a situation, or in which an entity is transformed—typically, the modification of an action in terms of its quantity, duration, style, etc. Examples: potato <cooked>; egg <omelet>

Physical state: A physical state of a situation or any of its components except entities whose states are coded as systemic features. Examples: mountains <dewy>; beer <hangover>

Quantity: A numerosity, frequency, intensity, or typicality of a situation or any of its properties except when the feature is an entity feature, whose quantitative aspects are coded with entity-quantity. Examples: party <many people>; sea <long trip>

Taxonomic Categories: Categories in the taxonomy to which a concept belongs.

Superordinate: A category one level above the target concept in a taxonomy or referring to a basic kind of thing in existence, including thing, substance, object, human, animal, plant, location, time, activity, event, action, state, thought, and emotion. Examples: deer <a mammal>; hammer <a tool>

Coordinate: Another category in the superordinate category to which a concept belongs. Examples: coyote <dog>; veil <hat>

Subordinate: A category one level below the target concept in a taxonomy. Examples: lettuce <romaine>; pants <bell-bottoms>

Individual: A specific instance of a concept. Examples: dog <Lassie>; doll <Barbie>

Synonym: A synonym of a concept. Examples: calf <baby cow>; sink <basin>

Antonym: For nouns, this property also indicates contrasting concepts. Examples: black <white>; fruit <vegetable>

Introspective Features: Properties of a subject's mental state as he or she views a situation, or properties of a character's mental state in a situation.

Affect/emotion: An affective or emotional state toward the situation or one of its components by either the subject or the participant. Examples: wasp <annoying>; bomb <frightening>

Evaluation: A positive or negative evaluation of a situation or one of its components by either the subject or a participant. Examples: homework <stupid>; gown <fancy>

Representation: A relatively static or stable representational state in the mind of a situational participant, including beliefs, goals, desires, ideas, perceptions, etc. Examples: tree <wanted to cut it down>; tree <I had a good VIEW of a bird in it>

Contingency: A contingency between two or more aspects of a situation, including conditionals and causals, such as *if*, *enable*, *cause*, *because*, *becomes*, *underlies*, *depends*, *requires*, etc.; correlations such as *correlated*, *uncorrelated*, *negatively correlated*, etc.; others, including possession and means. Examples: garlic <causes bad breath>; shirt <requires ironing>

Cognitive operation: An operation on a cognitive state, including comparison, retrieval, learning. Examples: buffalo <like a cow>; magazine <like a book>

Negation: An explicit mention of the absence of something, with absence requiring a mental state that represents the opposite. Examples: ostrich <cannot fly>; apple <not an orange>

Quantity: A numerosity, frequency, intensity, or typicality of an introspection or one of its properties. Examples: truth <a SET of beliefs>; buy <I was VERY angry at the saleswoman>

Lexical Features: Properties at the word level by virtue of orthographic similarity and completions, mediated responses through implicit common features or similar concepts, words used in common expressions, and meta-comments pertaining to the task and the stimulus (e.g., indications of word class),

Forward completion: The use of a word as a prefix to the response. Example: bomb <shelter>

Backward completion: The use of a word as a suffix to the response. Example: fish <sword>

Word fragment: A part of a compounded word. Example: jellyfish <jelly>

Orthographic similarity: Applying a minor modification to the orthographic composition of the word in such way that the added or removed letters compose a meaningful unit. Typically, this includes rhymes. Example: wine <whine>

Mediation: A response generated through an intermediate associate which might refer to an entity or a shared entity property or situation property. Examples: canary <banana>; whale <fireman>

Expression: Word pairs that occur in common expressions or sayings. Examples: color <outside the lines>; birds <bees>

Metacomment: A comment on the task or the characteristics of the word as a lexical entity. Examples: noon <palindrome>; papaya <huh>