

72 Semantic Cognition: Its Nature, Its Development, and Its Neural Basis

JAMES L. MCCLELLAND, TIMOTHY T. ROGERS, KARALYN PATTERSON, KATIA DILKINA,
AND MATTHEW LAMBON RALPH

ABSTRACT We consider parallels in the development of conceptual knowledge and its disintegration in a neurological disorder. During development, children's concepts become more differentiated. Developing children make striking errors, attributing properties to objects that they do not have and over-extending frequent category labels to other objects. Patients with semantic dementia (SD), a progressive condition affecting the anterior temporal lobes, show a loss of differentiation of concepts. Like developing children, they attribute properties to objects that they do not have and over-extend frequent category labels. Such patients may draw four legs on a duck, or call an elephant a horse. We present a theoretical framework and computational models that explain these parallel phenomena. We also consider the disintegration of lexical knowledge, which parallels the deterioration of semantics in SD patients. We present a model that simulates these findings within a single system that processes and represents both lexical and semantic information.

Interest in the nature of conceptual knowledge extends back at least to the ancient Greek philosophers. In recent years, there has been a wide range of different approaches to understanding the nature of conceptual knowledge, its development, and its neural basis. In most other work, however, these issues are not all treated together. Instead, workers in philosophy, adult experimental psychology, child development, and cognitive neuroscience have pursued related questions in relative ignorance of each other's efforts. Even within cognitive neuroscience, there has been until recently a relative separation between approaches taken by neuropsychologists, who study the effects of brain disease on cognition in patients, and researchers who study the neural basis of conceptual knowledge in

neurologically intact populations, using functional imaging and related methods.

We have sought to develop an integrative perspective on these matters. Our effort is facilitated by our theoretical framework, which lends itself to implementation in computational models that capture both the ability to learn gradually from experience, in order to model development, and the tendency to degrade in a graded fashion, in order to capture the partial nature of the deficits resulting from brain injury. We begin with an overview of the theoretical framework and its application to three developmental phenomena. We then show how the framework also addresses parallel phenomena that arise in the striking neuropsychological condition called semantic dementia. We review evidence that the disorder affects knowledge of words as well as knowledge of things, motivating an extension of our theory in which knowledge of things and words is fully integrated, contra most other approaches. A final section describes imaging and magnetic stimulation studies in normals that test predictions arising from the theory; this section considers evidence from disorders other than semantic dementia that indicates how the theory might be extended to address the flexible use of semantic knowledge in complex task situations.

The PDP framework and its application to development

The parallel distributed processing approach (Rumelhart, McClelland, & the PDP Research Group, 1986) provides the starting point for our theory of semantic cognition. The fundamental tenets of the approach are as follows:

- Cognitive activities emerge from the interactions of large numbers of simple processing units and are distributed over populations of such units both within and across brain areas.
- Active representations in this framework—the representation someone may have when, for example, he brings to

1 JAMES L. MCCLELLAND Stanford University, Stanford, California.
TIMOTHY T. ROGERS University of Wisconsin, Madison, Wisconsin.
KARALYN PATTERSON MRC Cognition and Brain Sciences Unit, Cambridge, United Kingdom.
KATIA DILKINA Stanford University, Stanford, California.
MATTHEW LAMBON RALPH University of Manchester, Manchester, United Kingdom.

mind a particular dog, who may be greeting him by jumping up and down, barking, and licking at his face—are likewise distributed, involving activity of many contributing units in each of many disparate brain areas that contain neurons representing the shape, color, odor, movement, and sounds made by the dog being imagined.

- The ability of one kind of information (say the sight of a particular dog, or its bark, or simply the word “dog”) to bring this information to mind depends on knowledge stored in the patterns of strengths or “weights” on the connections among the participating neurons.

- The patterns of connection weights are gradually acquired through experience. This process takes place over developmental time (i.e., years), gradually affecting the details of the representations we are able to bring to mind from particular inputs, and thus shaping gradual change in our behavior in cognitive tasks.

- The units and the connections are the substrate we use to understand not only normal cognitive functions and their development, but also the effects of brain damage and brain diseases on these functions. In particular, we assume that damage or disease results in the loss or disruption of the units and/or connections.

Within this theory, the effects of experience on connection weights explain many aspects of conceptual development; and the effects of damage and disease on the units and connections explain the disintegration of these functions in conceptual disorders such as semantic dementia.

It should be noted that our theory posits two complementary learning systems, only one of which—the neocortical one—is the focus here. We also propose that there is a second, fast-learning system based in the medial temporal lobes, allowing rapid learning of new, arbitrary information (McClelland, McNaughton, & O’Reilly, 1995). Normally, new semantic learning depends on both systems working together, but semantic knowledge is thought to gradually become independent of the fast-learning system, as evidenced by patients with profound amnesia for new information who nevertheless retain their semantic abilities (Squire, 1992).

SOME PHENOMENA IN CONCEPTUAL DEVELOPMENT As children develop, their conceptual knowledge gradually changes, in a way that appears to reflect aspects of experience. Three key features of this process are the following:

- Progressive differentiation of conceptual knowledge in development
- A tendency to overgeneralize names of frequently occurring objects
- A tendency to produce “illusory correlations” in attributing properties to objects

Differentiation in development has been explored in several contexts. In one investigation, Keil (1979) asked children what kinds of attributions or “predications” could apply to particular objects. For example, he asked children if it was “ok” or “silly” to say that something (say, a movie or a chair) “is sorry” or “is an hour long.” Using their responses, he constructed, for each of many children at each of several ages, a “predicability tree” like the ones shown in figure 72.1. Here we see trees for representative children at progressively older ages. Clearly, these trees indicate conceptual differentiation. As children grow older, they cease to lump together concepts that older children (and adults) pull apart. Similar conclusions arise from work on infants using nonverbal methods (e.g., Mandler & McDonough, 1993; Pauen, 2002).

Overgeneralization of frequently encountered names is also well documented. Such behavior is striking since it often represents an error that first emerges and then subsequently disappears—a classic “U-shaped” trend in development. As Mervis (1987) has investigated in detail, one common case of such overgeneralization is the extension of the word “dog” as a name for a wide variety of different animals, particularly other four-legged animals, by children in the early childhood years.

The developmental phenomenon of “illusory correlations” has been treated by many investigators as a sign that young children have acquired domain-specific causal theories that lead them to overapply properties attested in these theories to other objects (Keil, 1991). One such illusory correlation was observed by Gelman and Williams (1998). They showed pictures of objects to young children and asked if the objects could go up and down a hill by themselves. Children usually answered yes if the pictured object appeared to be an animal. When asked to explain their answers, children often attributed feet to the pictured animals, even in cases where no feet were in evidence in the pictures. While in some cases (e.g., an animal called an echidna, which appeared as a small furry ball) such attributions were justified, in other cases (e.g., a snake) they would clearly be “illusory correlations”—perceived correlations of movement properties with physical properties of objects that are often, but not always, valid. Such overattributions are not by any means restricted to young children, but also occur in adult cognition. For example, people in the United States have a tendency to perceive an innocuous object as a weapon when it is in the hands of an African American (Eberhart, Goff, Purdie, & Davies, 2004).

APPLICATION OF THE THEORY TO DEVELOPMENT: THE RUMELHART MODEL Our application of our theory to conceptual development grew out of earlier work by Hinton (1981, 1989) and Rumelhart (1990). Rumelhart wished to

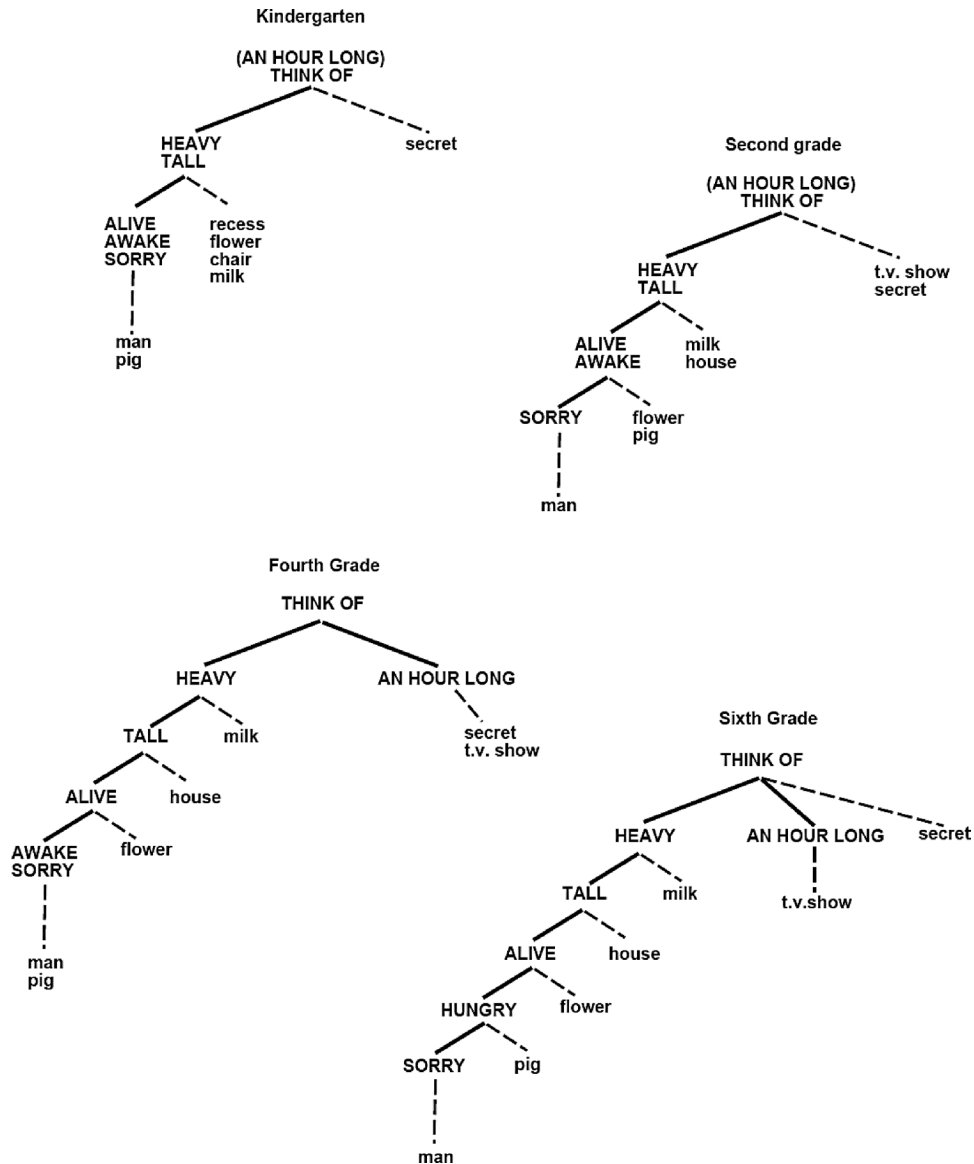


FIGURE 72.1 Examples of the predicability trees constructed by Keil (1979), from four individual children in different age groups. The trees indicate which of several predicate terms are accepted by the individual children as being applicable to various concepts; concepts accepting the same set of predicates are grouped

together at the same branch of the tree. (Reprinted from Rogers & McClelland, 2004, figure 1.3, p. 10, based on appendices A3, A17, A37, and A54 from Keil, 1979, pp. 181, 183, 185, and 187. Permission pending.)

articulate an alternative to prior ways of thinking about conceptual knowledge and chose the domain of living things as his example. Instead of storing knowledge explicitly as a network of linked propositions, as shown in figure 72.2, he proposed it might be stored in connections among simple processing units, as in the network shown in figure 72.3. In this network, called the *Rumelhart network*, the knowledge, say, that a robin can grow, can move, and can fly is not stored in explicit propositions, but instead arises from a pattern-completion process. We query the network to tell us what a robin can do by activating units for “robin” and “can” on

the input side of the network, and propagating activation forward through intermediate layers to the output layer. The network is simplified in many ways relative to our overall theory, but it provides a useful ground for explaining the patterns we have discussed that arise in development.

When it is first initialized, the network’s connection weights are both small and random, so that a query produces very neutral and undifferentiated patterns of activation at all levels of the network forward from the item and context inputs. However, the network is trained with repeated exposure to the information in figure 72.2. Each input consists of




FIGURE 72.2 A taxonomic hierarchy of the type used by Quillian (1968) in his propositional model of the organization of knowledge in memory. ISA links may be followed up the tree to infer proper-

ties not explicitly connected to items below. (Reprinted from Rogers & McClelland, 2004, figure 1.2, p. 6. Permission pending.)

an item (one of the eight items at the bottom of the figure) in each of four relational contexts (called “is” for appearance properties, “ISA” for the categories to which it belongs, “can” for things it can do, and “has” for specifying its parts). Activity propagates forward through the network and is then compared to the correct completion provided by the environment. Note that we do not envision this as an explicit instruction process, but simply a matter of anticipating future inputs from current inputs. It is the mismatch between what the network anticipates and what is provided by the environment that drives the adjustment of connections. In this case, target patterns are the correct set of completions of the item-context pair provided as input. We treat the model as a proxy for the learning children do both from explicit (verbal) propositions provided by others in their environment and from actual experiences with objects in different contexts. For example, the child watching a robin on a branch may not anticipate at first that it will fly away as a cat creeps up upon it. Witnessing that the bird does fly away provides a signal that does not match expectations, and it is the difference between the (null) expectation and the witnessed action that then drives connection-based learning.

The details of the learning process are described in Rogers and McClelland (2004). The process is described as “back-propagation of error,” and its biological plausibility has been much maligned, but it has repeatedly been shown how the necessary error information can be derived from temporal differences in activation in networks with bidirectional con-

nections (e.g., O’Reilly, 1996). What is important for our purposes is that all the connections in the network are affected by this process. Some of these connections serve to change how patterns of activation inside the system affect the output units. Others serve to influence how external inputs (on the left in the diagram) are internally represented. The changing structure in these representations is crucial for the patterns of change that we see in development.

Figure 72.4 presents the patterns of activation seen in the “representation” layer of the network in figure 72.3 at three points during learning. Each histogram bar in the left panel of the figure represents the activation of one of the units in the representation layer of the network for a particular item at a particular time. The set of eight such bars for a particular item at a particular time constitutes that item’s internal representation. We can see that initially each unit takes a middling activation value for all items, so that the patterns are not well differentiated. The small differences at this stage largely reflect the initial random noise in the connection weights. In contrast, at the last time point (right column of figure) it is clear that the patterns have become quite differentiated. Hierarchical clustering analysis (not shown) reveals what can also be seen by eye. The network treats the two different types of fish as very similar, but as quite distinct from the two types of birds; all these are very different from all of the plants. Among the plants, the two trees are very similar and so are the two flowers. The trees and flowers are somewhat differentiated from each other, though not as much so as the birds are from the fish.

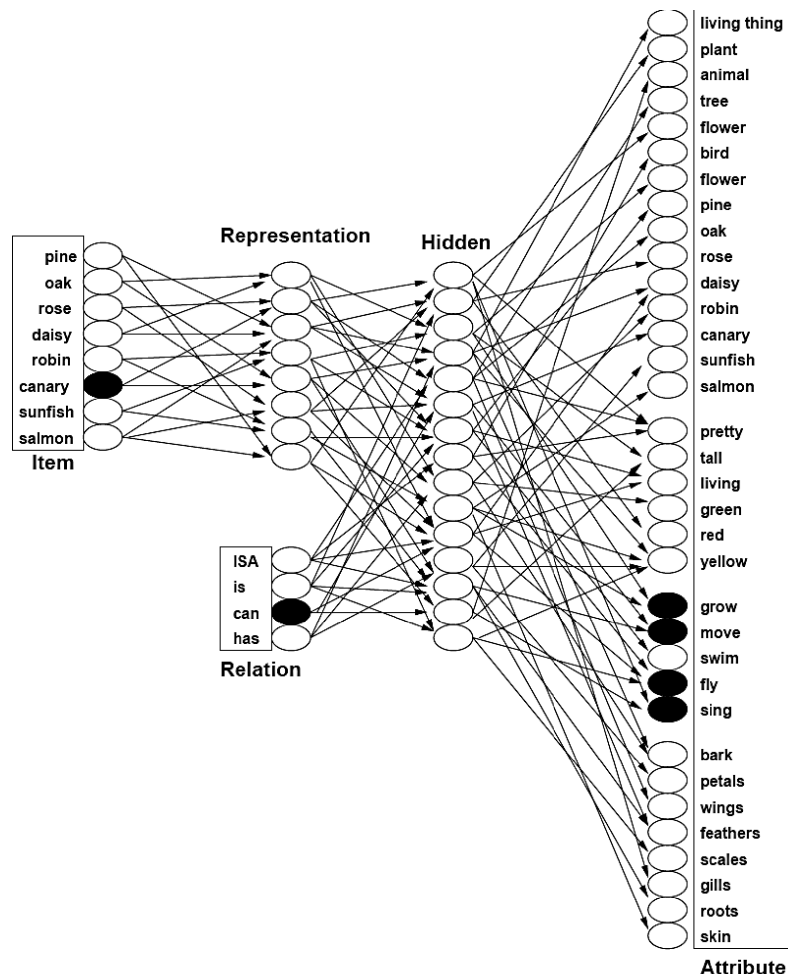


FIGURE 72.3 The connectionist model of semantic memory used in the developmental simulations, adapted from Rumelhart (1990; Rumelhart & Todd, 1993). The entire set of units used in the network is shown. Input units are shown on the left, and activation propagates from the left to the right. Where connections are indicated, every unit in the pool on the left is connected to every unit in the pool to the right. Each unit in the *Item* layer corresponds to an individual item in the environment. Each unit in the *Relation*

layer represents contextual constraints on the kind of information to be retrieved. Thus the input pair *canary can* corresponds to a situation in which the network is shown a picture of a canary and asked what it can do. The network is trained to turn on all the units that represent correct completions of the input query and to turn off all other units. In the example shown, the correct units to activate are *grow*, *move*, *fly*, and *sing*. (Reprinted from Rogers & McClelland, 2004, figure 2.2, p. 56. Permission pending.)

What is particularly interesting for our purposes is the fact that the representations undergo a progressive differentiation. Specifically, at the intermediate time point shown, the network has successfully differentiated the plants from the animals before it further differentiates these two superordinate categories into their particular intermediate-level types. The full trajectory of this process is illustrated in a two-dimensional projection in figure 72.5. An animated version of this figure best captures the progressive differentiation of the network's representations over time. The animation at <http://psychology.stanford.edu/~jlm/Presentations/Differentiation.mpg> illustrates that the progressive differentiation process is highly stagelike in character, as seen in many aspects of children's cognitive development (McClelland, 1989). In the animation, we can see that all the patterns are

initially undifferentiated; they first divide into the plants on the one hand and the animals on the other; then the mammals diverge from the birds; then the trees from the flowers; then the different individuals differentiate from one another.

OVERGENERALIZATION OF NAMES AND ILLUSORY CORRELATIONS The overgeneralization of frequent names and the presence of illusory correlations both arise from the progressive differentiation process. Simulations illustrating the transitory developmental emergence of both these effects are shown in figures 72.6 and 72.7. In figure 72.6 we see what happens in a simulation in which there were four trees, four flowers, four fish, four birds, and five four-legged animals. Among the latter, one—the dog—occurred ten times more frequently than the others, including the goat

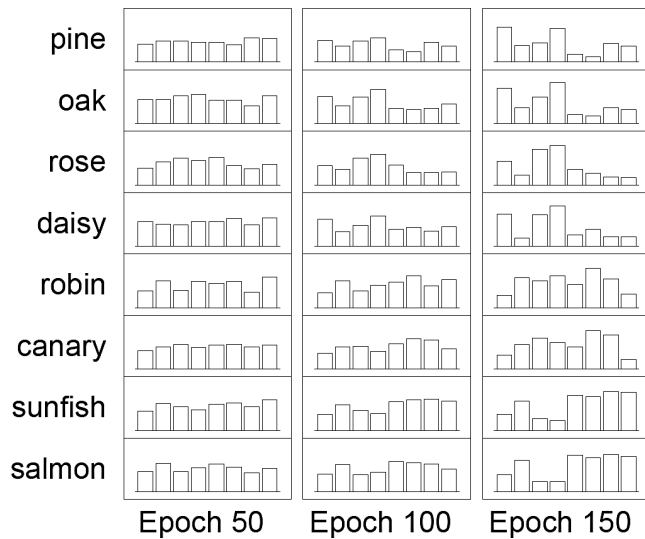


FIGURE 72.4 The process of differentiation of conceptual representations as seen in the Rumelhart model. Learned internal representations of eight items at three points during learning, using the network shown in figure 72.3. The height of each vertical bar indicates the degree of activation for one of the eight units in the network's *Representation* layer, in response to the activation of a single *Item* unit in the model's input. Early in learning (50 Epochs), the pattern of activation across these units is similar for all eight objects. After 100 epochs of training, patterns have begun to differentiate at the superordinate level (plants versus animals) but not at the intermediate level (trees versus flowers, birds versus fish). This further differentiation is apparent after 150 epochs, and continues down to the subordinate level as training continues. (Reprinted from Rogers & McClelland, 2004, figure 3.1, p. 86. Permission pending.)

(consistent with input children receive, Rogers & McClelland, 2004, chapter 5). In this case, the network had a tendency, at a certain point in its development, to activate the name “dog,” not only for the dog but also for the other four-legged animals, including the goat. This tendency arises at a time when the four-legged animals are differentiated from the other animals but are not yet well differentiated from each other, and falls away again as the different land animals pull apart. This tendency makes sense from an optimal inference point of view, if we consider the conditional probability of hearing the label “dog” when experiencing the internal representation shared by the land animals (this probability is $10/(10 + 1 + 1 + 1 + 1)$ or about 0.7). Thus when the network has only one shared internal representation for all land animals, its best guess is to call them all “dog.” Once the representations differentiate, however, the conditional probability of each label given each item's representation changes dramatically, leading the network no longer to overgeneralize.

In figure 72.7, we see the tendency for the network to activate the property “has leaves” for the pine tree. There is a phase, relatively early in the network's development, where

it attributes leaves to all things with a middling activation value (this value represents the proportion of all objects that have leaves). When the plants begin to differentiate from the animals, however, the network at first attributes leaves more strongly to all the plants, including the pine tree. Only as it comes to differentiate the pine from the other plants does it reverse this “illusory correlation.” Once again, this pattern makes sense from an optimal inference point of view. At the point in development where the representation of the pine tree is identical to the representation of the other plants, the conditional probability of “leaves” is very high (all the other plants have leaves). Again, once the representation of the pine differentiates from the representation of each of the other plants, the conditional probabilities change. Now, the network can learn the conditional probability of leaves for pine is actually 0.

SENSITIVITY TO COHERENT COVARIATION AND ITS DEPENDENCE 5

ON THE ARCHITECTURE OF THE NETWORK In Rogers and McClelland (2004) the reasons why these three phenomena occur in the network are extensively explored. The essential point of this analysis is the observation that the connection weights from the item units to the representation units—and in consequence, the patterns of activation assigned to each concept on the representation layer—are sensitive to the pattern of coherent covariation of properties across the items presented to the network. The fact that all the animals share one set of properties that none of the plants have, while the plants share another set of properties that none of the animals have, is responsible for the first wave of differentiation. The error signals reaching the representation layer (and therefore driving the connection weights) tend to push the representations in a similar direction for all the animals—a direction that is different from the direction in which the error signals push the representations of all the plants. The subsequent differentiation of the different types of animals is also driven by the fact that each type of animal shares a set of properties that none of the other types possess, and similarly for the differentiation of the different types of plants. Now, it is the connections forward from the representation layer to the output layer that determine what output the network generates when a given pattern is present on the representation units. During the phase of development when the network is essentially representing the dog and all the other land animals as the same, but different from the plants and from the other types of animals, the correct name response when this shared representation is present over the representation layer is usually “dog,” because the dog occurs more frequently than any of the other land animals. The error-correcting learning process thus pushes the weights forward from this representation layer to activate the name “dog” more often than any other name, thereby accounting for the tendency for all land animals to be called “dog” at




FIGURE 72.5 Trajectories of item representations (lines fanning from center of space) and their ultimate positions in the Rumelhart network's semantic representational space. Shaded regions around the final points indicate the approximate size of regions associated with subordinate, intermediate, and superordinate levels

as indicated. When an item has higher frequency, the region associated with that item and its properties is increased. (Reprinted from Rogers & McClelland, 2004, figure 3.9, p. 112. Permission pending.)

this stage of development. In the case of attributing leaves to the pine tree, the situation is similar. As long as the representation of the pine tree is similar to the representation of all other plants, the network tends to attribute “has leaves” to it as it does to the other plants. Thus both overgeneralization of frequent names and illusory correlations arise as a consequence of the sensitivity of the representations in the network to coherent covariation.

It is important to see that the network's tendency to be sensitive to coherent covariation, and thus to exhibit differentiation, name overextension, and illusory correlations, is a feature of its architecture (figure 72.8). In the extreme, if each item-context pair projected to its own distinct representation unit, which in turn projected forward to the appropriate properties for the object (as in figure 72.8B), there would be no sensitivity to coherent covariation at all. The error signals driving learning of each property of each item would be completely segregated from those relevant to every other property. In our actual architecture (figure 72.8A), something very different happens. Because the error signals for each property of each object in each context are projected on the same set of representation units, and because different concepts share these representation units, what is learned about an object in one context tends to be shared

both across objects and across contexts. In short, a key observation from our theory is this:

Sensitivity to coherent covariation, a tendency central to explaining many aspects of conceptual development, requires the use of a shared representation layer mediating all aspects of conceptual knowledge of all different kinds of things.

The particular architecture of the Rumelhart network is only one architecture with this property; in the rest of this article we will be considering networks with slightly different architectures that still make use of a single shared representation mediating all kinds of knowledge of all kinds of things.

Disintegration of conceptual knowledge in semantic dementia

Although the network discussed so far is quite abstract, the simulations do suggest an important hypothesis about the architecture of the cortical semantic network: that there must exist some place in the network where all different kinds of information *converge*, so that different items and events, regardless of the modality of input or the particular semantic domain, get processed through the same set of neurons and synapses. It is precisely this convergence of




FIGURE 72.6 Activation of all the name units when the model is probed for its knowledge of basic names, in a network trained with dog patterns eight times as frequently as other mammal patterns in the environment. Early in learning, the network tends to inappropriately activate the name “dog,” especially for related objects. (Reprinted from Rogers & McClelland, 2004, figure 5.1, p. 214. Permission pending.)

information that promotes sensitivity to coherent covariation in this network. This same convergence can explain the striking pattern of semantic deficits seen in patients with a rare neurological disorder, semantic dementia.

CHARACTERISTICS OF THE DISORDER Semantic dementia (SD) is a progressive deterioration of conceptual knowledge in the context of otherwise greatly preserved cognitive function (Snowden, Goulding, & Neary, 1989; Hodges, Patterson, Oxbury, & Funnell, 1992). Patients show serious deficits in any task requiring them to access knowledge about any type of thing from any form of input (Bozeat, Lambon Ralph, Patterson Garrard, & Hodges, 2000). Thus patients are impaired at both comprehending and producing speech, recognizing words and line drawings of common objects, indicating the correct color for black-and-white drawings of familiar items, choosing which of a set of items makes a particular sound, matching objects on the basis of shared function or use, demonstrating the use of everyday objects, and even recognizing common odors. These impairments affect knowledge for all different kinds of concepts—living and nonliving, abstract and concrete, verbs and nouns—and are apparent regardless of the modality of reception or expression tapped by a particular test. Despite these serious

disabilities, patients with SD perform normally or near-normally in tests of basic perception, episodic and working memory, executive function, problem solving, and attention; and, apart from word-finding difficulties arising from their conceptual deficits, they produce fluent and grammatical speech. Such patients thus appear to exhibit a pure and progressive cross-modal and domain-general impairment of semantic or conceptual knowledge (Patterson, Nestor, & Rogers, 2007; Lambon Ralph & Patterson, 2008).

As notable as the syndrome itself is the remarkable anatomical specificity of the cortical atrophy observed in the disease, which without exception affects anterolateral regions of the temporal lobes. Bilateral degeneration is the norm, though pathology is usually asymmetrical and patients with left-predominant atrophy present at about double the rate of right-predominant cases. The typical pattern on structural MRI is well-defined atrophy of both anterior temporal lobes that is maximal at the temporal pole and on the adjacent rostral-inferior surface (Hodges & Patterson, 2007; Patterson et al., 2007).

The pattern of semantic dysfunction together with the anatomical specificity of the atrophy provide strong evidence that the cortical semantic network adheres to the “convergence principle” suggested by the modeling work discussed earlier. Specifically, the findings from SD suggest that semantic knowledge for all kinds of concepts, across all modalities of reception and expression, depends upon a relatively circumscribed region of the anterior temporal lobes. Perhaps, then, these regions play a functional role similar to that of the *representation* layer of the Rumelhart network: by processing information about all kinds of items in many different situations and contexts, perhaps these regions form learned internal representations of inputs that capture the semantic similarity among items.

MODELING SEMANTIC DEMENTIA Rogers, Lambon Ralph, Garrard, and colleagues (2004) investigated this hypothesis using a PDP model based on the architecture shown in figure 72.9. Here, different kinds of sensory, motor, and linguistic information are represented in different pools of units, with each pool dedicated to a particular kind of information. These *surface* representations receive direct input from the corresponding sensory systems so that, whenever a given stimulus is encountered, the units that code its directly observed properties are activated. Like many other researchers, we believe the different surface representations to be subserved by different brain areas, organized predominantly by modality, and situated near the sensory channels from which they receive input (Martin & Chao, 2001). Also in common with other researchers, the function of the semantic system in this framework is to mediate learned associations among the various different surface representations—so that, for instance, when a line




FIGURE 72.7 The activation of the “has leaves” and “can sing” output units across the first 5000 epochs of training, when the network is probed with the inputs *pine has* and *canary can*, respectively. At epoch 1500, the network has been trained 150 times to turn off the “has leaves” unit in response to the input *pine has*, and

to turn on the unit “can sing” in response to the input *canary can*. Nevertheless, the network still activates the “has leaves” unit for the pine tree, and fails to activate the “can sing” unit for the canary. (Reprinted from Rogers & McClelland, 2004, figure 6.6, p. 254. Permission pending.)

drawing of a pencil is observed, representations of associated attributes in other modalities, such as the appropriate color (yellow), name (“pencil”), and action (writing), will also become activated in the appropriate modality-specific brain areas.

Our framework differs from some others in proposing that the associations between all different forms of surface representations are mediated by a central “hub” in the anterior temporal lobes. The hub itself receives no direct input from sensory systems, but receives connections from and sends connections to the surface representations that code representations of particular sensory, motor, and linguistic attributes. The hub, illustrated in figure 72.9, is similar to the *item representation* layer in the Rumelhart network in that the representations there are a consequence of the learning process that shapes the weights projecting into and out from the layer. The surface representations in figure 72.9 are analogous to the *attribute* output units in the Rumelhart model, in that they explicitly encode observable characteristics of items in the environment. Because processing is recurrent in this model—activation flows both from the surface representations to the hub and from the hub back to the surface—there is no need to have separate “input” and “output” layers. Instead, any given input can be specified as a distributed pattern of activation across corresponding units

in the surface layers of the model. These inputs will propagate activation up to the hub of the anterior temporal lobes, which will then feed activation back to the surface representations to activate other properties of the item that have not been directly observed.

This framework also provides a natural paradigm for learning. The top-down activation of surface attributes may be viewed as the generation of an implicit expectation about the item’s unobserved properties. If these expectations are contradicted by a subsequent observation, the discrepancy can be used as an error signal to drive weight changes throughout the network, so that the system comes to generate increasingly accurate expectations.

Rogers, Lambon Ralph, Garrard, and colleagues (2004) used a model implementing these ideas to assess whether the theory could account for patterns of semantic impairment observed in SD. The model consisted of a *Visual* layer in which each unit coded a visually apparent property of an object; a *Verbal* layer, in which each unit coded a predicate (e.g., “has wings”) that might appear in a verbal description of an object, including names and other descriptors; and a *Semantic* layer that mediated interactions between these. Visual perception of an object was simulated by directly activating the object’s properties in the Visual layer and allowing activation to propagate through Semantic units to

6

Y

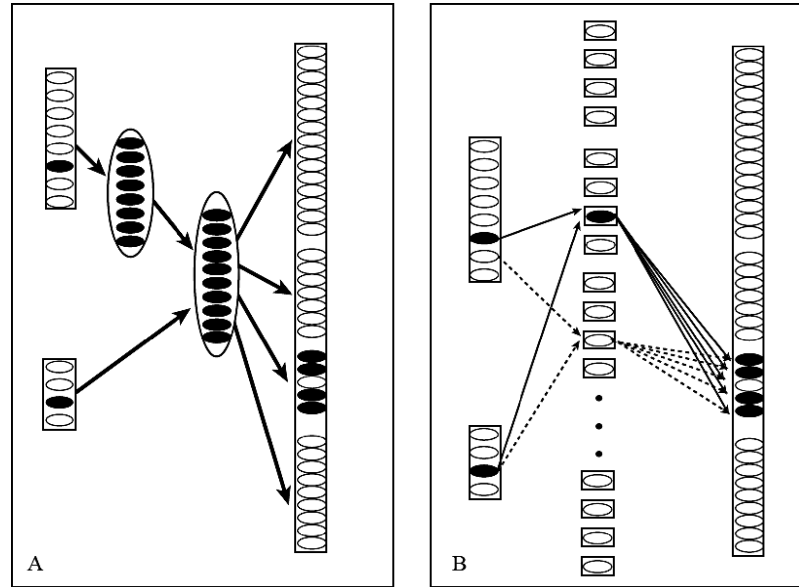


FIGURE 72.8 Two alternative feedforward architectures for mapping from localist *Item* and *Context* inputs to sets of output properties. Thick arrows in *A* indicate full connectivity among units in the sending layer to those in the receiving layer; thin arrows in *B* indicate individual connections. The shading of units indicates how activation spreads forward given the input *canary can*. (*A*) The Rumelhart network architecture, in which all items first map to a single context-independent hidden layer and then converge with context inputs in a second hidden layer. The first layer of hidden units receives error signals that are filtered through the second,

convergent representation; hence the architecture constrains the network to find context-independent representations of individual items that are sensitive to coherent covariation of properties across different contexts. (*B*) Separate localist representations for every possible conjunction of item and context. Only connections for *canary can* (solid arrows) and *salmon can* (dotted arrows) are shown. In this case the model will not generalize and will not be sensitive to coherent covariation. (Adapted from Rogers & McClelland, 2004, figure 9.1, p. 357. Permission pending.)




FIGURE 72.9 Architecture of the model used to simulate semantic dementia. (Reprinted from Rogers, Lambon Ralph, Garrard et al., 2004, figure 1, p. 207. Permission pending.)

Verbal units. Presentation of an object name was simulated by directly activating the single unit representing the name in the Verbal layer; and presentation of a verbal description of an object was simulated by activating the subset of units representing the presented predicates. The model was then trained with backpropagation to produce the correct visual pattern, verbal pattern, or name when provided with one of these representations as input. Representations on the

Semantic layer were not specified but, as in the Rumelhart network, emerged as a consequence of learning. Once trained, the model permitted simulation of the most basic tasks used to assess semantic memory in patients with SD, including visual object naming, drawing-to-name, delayed-copy drawing, word-to-picture matching, sorting words and pictures, and so on. The patterns used to train the model were constructed to capture important aspects of similarity

structure apparent in both verbal attribute-listing studies and drawings of common objects. In both cases, superordinate category structure was clearly apparent—items from the same superordinate category (e.g., animals, manmade objects, and plants) tended to share many verbal descriptors and also to have similar visible parts in their drawings. There was also some more specific structure, especially among the set of animals; for example, different birds were more similar to each other than they were to other types of animals.

After training, the model represented each individual item—whether accessed by means of a single name, a verbal description, or a visual pattern—with its own domain-general pattern of activation across Semantic units. Just as with the Rumelhart network, these patterns captured similarity relations, with semantically related items represented by similar patterns of activation.

The deficit in semantic dementia was simulated by removing an increasing proportion of the connection weights projecting into or out from the intermediating hidden layer; the model was then tested on analogues of the semantic tasks used with patients. Comparable results are obtained by deleting individual units rather than individual connections.

The model naturally captures both the cross-modal and cross-category nature of the semantic impairment in SD. Semantic task performance was compromised both for visual semantic tasks like delayed-copy drawing—where the participant must draw a previously viewed item from memory after a short delay—and for cross-modal tasks like object naming; and the magnitude of impairment was roughly equivalent for animals and manmade objects (Lambon Ralph, Lowe, & Rogers, 2007). These aspects of model performance follow from the fact that it respects the convergence constraint identified in the previous section. When these weights were lesioned, performance for all tasks and semantic domains was affected.

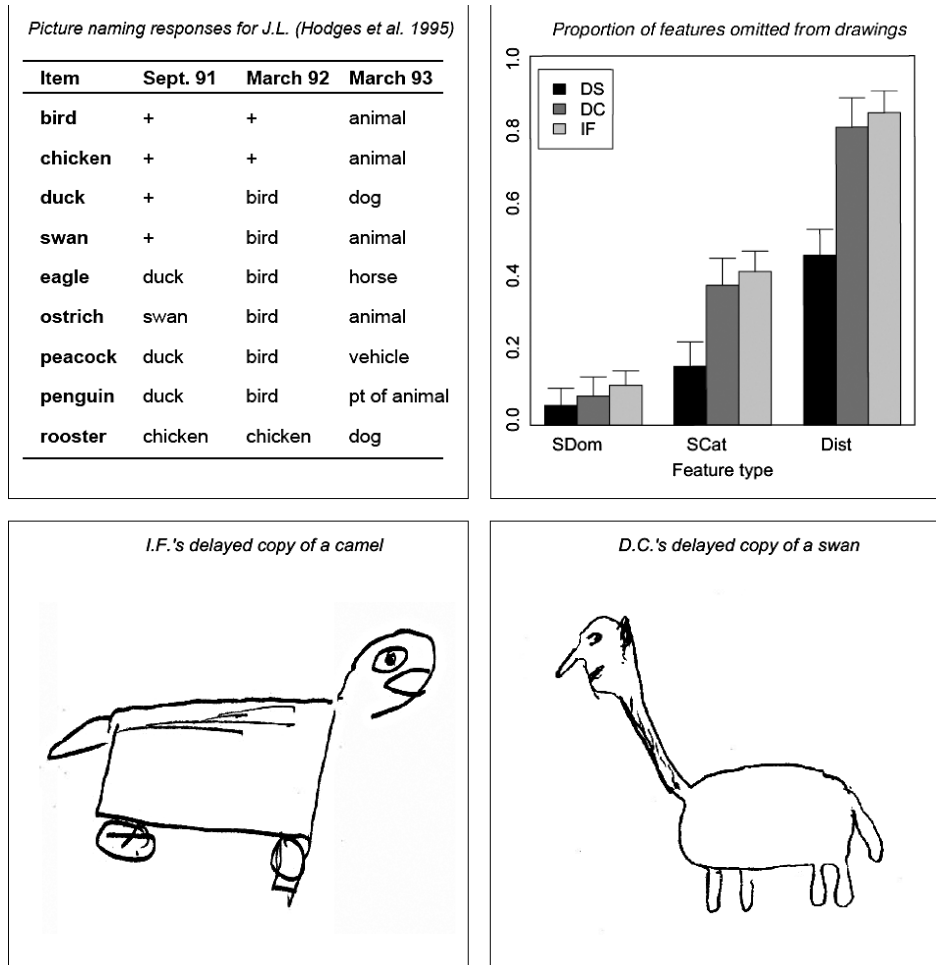
PARALLELS BETWEEN DEVELOPMENT AND DISINTEGRATION

The model also offers a clear explanation of what for us is one of the key features of semantic dementia: loss of differentiating details about particular concepts together with spared knowledge of more general information. This loss of differentiating detail was first documented in the initial report of semantic dementia by Warrington (1975). Warrington also pointed out the parallelism between this finding and the lack of differentiation in early stages of conceptual development. Warrington demonstrated that knowledge of an item's properties that characterize broad semantic categories—for instance, the fact that tigers have fur—is much more robust than knowledge of item-specific properties, such as the fact that tigers have stripes. Knowledge about properties that characterize very specific classes—for instance, particular breeds of dog—is much more vulnerable to early impairment,

as is knowledge about less frequent and less prototypical items. So, when categorizing familiar objects, patients with very mild impairment will usually fail at naming objects at subordinate levels such as “robin” or “BMW,” but can seem unimpaired at more general levels such as “bird” or “car”; and even very semantically impaired individuals can succeed as well as controls at very general categories such as “animal” or “vehicle” (Rogers & Patterson, 2007).

In the model described by Rogers, Lambon Ralph, Garrard, and colleagues (2004), this erosion of knowledge about the individuating details of specific concepts arises as a consequence of the similarity structure of the distributed representation acquired in the semantic layer. To see this point, consider how the healthy model retrieves a fact specific to a particular subordinate concept, such as that a robin has a red breast. Even though Rogers and colleagues used the architecture shown in figure 72.9, the situation is still well reflected in figure 72.5, which shows the situation arising in the simpler training environment used in the Rumelhart network. What we see in this figure is a graphical depiction of the fact that the models learn to treat the robin as quite similar to all the other birds (in this case, the only other bird is the canary); yet none of these other birds has a red breast like the robin. Therefore, to correctly activate the “red breast” units in either the Visual or Verbal layer, the Hidden layer must instantiate the pattern of activation corresponding to the robin almost exactly—if this pattern is just a little different, then it may become more similar to another individual bird that does not have a red breast, and the system will fail to activate the correct property in the periphery. Thus relatively small distortions to the correct representation will prevent the model from strongly activating properties that are unique to very specific categories.

Now consider a property that the robin shares with other birds, like “has wings.” In this case, the model need not instantiate precisely the right representation to retrieve the property, since it is common to all of the birds and hence will be activated by all of the patterns that are somewhat similar to the correct pattern. If the robin representation is distorted so that it more closely resembles the canary representation, this distortion will not disrupt activation of the “has wings” units in the Visual and Verbal layers, because the canary has wings just like the robin (as do all of the other birds). Thus even with a relatively severe distortion to the representation the system can still generate the correct outputs for category-typical properties. The same argument suggests why still more general properties—like the fact that animals have eyes—are even more robust. If the robin representation is, as a consequence of brain damage, so degraded that it becomes less distinguishable from the various mammals as well as from the other birds, the system will still continue to correctly activate properties held in common between birds and mammals.



20 FIGURE 72.10 Evidence of conceptual disintegration in semantic dementia. (A) Naming responses given by patient JL to pictures of birds (drawn from a set of line drawings for which control subjects consistently provide the name given in the left column) at three times during the progression of his illness (+ indicates correct responses). (B) Proportion of features of different types omitted from drawings by three other semantic dementia patients. Patients were shown a picture of the object including all the tested properties and were asked to copy the picture from memory after a 10-second delay. All copied the picture accurately while it remained in view, but had difficulty in reproducing the distinctive but not the domain-general properties of the pictured objects after a delay. SDom: Properties shared by typical members of the general domain (e.g.,

21 eyes, shared by animals) of the test item. SCat: Properties shared by typical members of the superordinate category (e.g., wings, shared by birds). Dist: Distinctive features of the test item itself (e.g., stripes, distinctive attribute of tiger). (C) Delayed copy of a camel; no hump is evident. (D) Delayed copy of a swan. A long neck is present, indicating some preserved representation of specific information, but there are four legs, illustrating the tendency these patients have to fill in properties that are generally present in items within the overall domain (animals) even if not present in the specific item (swan) or even its immediate superordinate (bird). (Redrawn from McClelland & Rogers, 2003, figure 2, p. 312. Panel A excerpted from the appendix of Hodges, Graham, & Patterson, 1995, p. 490. Permission pending.)

The gradual loss of idiosyncratic detail coupled with the preservation of information shared among members of a category creates progressive dedifferentiation of semantic knowledge, paralleling the progressive differentiation seen in development. Hand in hand with this progressive dedifferentiation are two other phenomena that parallel those seen in development: (1) overgeneralization of names of frequently occurring objects and (2) illusory correlations, or the attribution of category- or domain-general properties to objects that lack these properties. Evidence of these aspects of SD is shown in figure 72.10. In figure 72.10A, we present

picture-naming data from patient JL at different stages of his progressive deterioration (Hodges, Graham, & Patterson, 1995). Here we see that, as his impairment becomes progressively worse, he shows an increasing tendency to overapply names of the more common animals (e.g., duck) to less common animals (eagle, peacock). In figure 72.10C and 72.10D, we present delayed copies made by two patients of a swan and a camel. In the latter case, the differentiating detail of the camel's hump is lost; more strikingly, in the former, a property typical of the broad class of animals—that of having four legs—is added to the swan, making it far

more like other animals. These phenomena are not idiosyncratic to particular objects or patients (figure 72.10B). There is a general tendency to produce names of more frequent category coordinates in object naming (Woolams, Cooper-Pye, Hodges, & Patterson, 2008) and a general tendency to both omit differentiating details and to incorrectly incorporate domain-general properties in patient's delayed copying (Bozeat et al., 2003).

The reasons for overgeneralization of frequent names and illusory correlations can again be understood with reference to figure 72.5. As illustrated in the figure, the points in representational space associated with specific items are a function of the item's frequency. A very small region surrounding a very specific point is associated with the idiosyncratic properties of individual objects, including their names, whereas a much larger region is associated with the properties of frequent objects (such as a dog) or the shared properties of many objects (such as having four legs). As a result, any distortion of the representation of a particular, relatively uncommon item will tend to result in the network's representation landing in a part of the space associated with more common objects and more typical properties. Indeed, the model shown in figure 72.9 was able to simulate closely the proportions of category coordinate naming errors seen in SD, as well as the proportion of item-specific omissions and category-general overgeneralization errors made by such patients (Rogers, Lambon Ralph, & Garrard, 2004).

In summary, just as the Rumelhart model explains the progressive differentiation of conceptual knowledge over development, the model of Rogers, Lambon Ralph, Garrard, and colleagues (2004) explains the apparent reversion of this process in SD: the gradual erosion of knowledge about the details that individuate concepts, beginning with very specific concepts and progressing to more and more general concepts. All these phenomena arise from the same general principles: that semantic knowledge is acquired through domain-general mechanisms, in a system that learns mappings among various different "kinds" of sensory, motor, and linguistic information and that stores these mappings within a convergent architecture in which all kinds of information for all kinds of concepts are processed through the same neurons and synapses.

Nonsemantic deficits in semantic dementia

The previous section demonstrates how the essential principle of convergence from widely distributed modality-specific brain regions can help us to understand the *semantic* deficits observed in semantic dementia (SD). This section addresses a *further* set of deficits observed in SD. These deficits are important because the abilities affected are ones that have been thought by many to operate *without* reference to seman-

tic knowledge. If SD constitutes a semantic impairment, why should these abilities be affected? One possibility is that the nonsemantic impairments are simply additional independent deficits that arise because of abnormalities in other, nonsemantic regions. However, these additional deficits both (1) occur consistently with the semantic deficit and (2) are similar in nature to the semantic deficit. From these observations, we have argued that they are a part of the core deficit itself (e.g., Patterson et al., 2006). We concentrate here on "nonsemantic" SD deficits in four tasks using words, though similar deficits occur with other kinds of stimuli. Nearly all SD patients had abnormal performance on each of the four tests: (1) lexical decision, in which the participant must judge whether each of a series of letter strings is a real word or must choose the real word when items are presented in pairs; (2) oral reading of single printed words; (3) written spelling of single spoken words; and (4) oral production of the past tense from present tense (stem) forms of verbs (Benedet, Patterson, Gomez-Pastor, & de la Rocha, 2006; Funnell, 1996; Graham, Patterson, & Hodges, 2000; Hodges et al., 1995; Patterson & Hodges, 1992; Patterson, Lambon Ralph, Hodges, & McClelland, 2001; Rogers, Lambon Ralph, Hodges, & Patterson, 2004; Saffran, 2003; Ward, Stott, & Parkin, 2000). The basis for performing three of these tasks (reading, spelling, and past tense formation) is traditionally considered to be a joint function of a system of rules and a system of lexical entries; the rules and lexical entries are often considered separate from each other and also separate from semantic knowledge (e.g., Caramazza, 1997; Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001; Levelt, 1989; Pinker, 2001). Lexical decision is often thought to depend only on lexical entries, although the presence of typicality effects in this task might be taken to suggest that some other system of knowledge, perhaps a system of rules, would be needed here, too.

Within the PDP framework, performance on all these tasks is thought to rely on a single integrated processing system that contains neither rules nor lexical entries but, like the semantic system, is sensitive both to properties that items share with each other and to idiosyncratic, item-specific information (Rumelhart & McClelland, 1986; Plaut, McClelland, Seidenberg, & Patterson, 1996). Words, like objects, tend to have properties that they share with others. Tigers have fur like other animals, but they also have their idiosyncratic stripes. Similarly, the word *pint* has correspondences it shares with many other words in the pronunciation of most of its letters, but it has an idiosyncratic correspondence in the pronunciation of the vowel; and the irregular verb *keep* forms its past tense like regular words in most respects. In regular verbs, a /t/ would be added to the stem (cf. *bake-baked*, pronounced /be:kt/). The same is true with *keep-kept*, but in addition there is an idiosyncratic vowel adjustment.

Semantic dementia patients have difficulty in all of these tasks in ways that parallel the deficits they show in semantic tasks. As their semantic disorder progresses, they make progressively more errors on words that are of low typicality, especially when they are also of low frequency. Also, the nature of the patients' errors mirror precisely the "illusory correlations" discussed earlier with reference to knowledge of the properties of objects. Semantic dementia patients apply to atypical words the correspondences that the words would have if they were more typical. Asked to read the word *sew*, they often say "sue" using the typical correspondence in *crew*, *few*, *new*, and so on; asked to spell "cough," they write *coff*, using the typical spelling for /f/ as in *off*, *scoff*, *fluff*, *cuff*, and so on. Asked to put the sentence "Every day I fight with my brother" in to the past tense, they will often say, "Yesterday I fought with my brother." A typicality effect is also seen in lexical decision. Asked which of the two letter strings *seize* and *seese* is a real word, severe SD patients actually prefer the incorrect but more typical spelling *seese*. The effect is strictly analogous to an effect seen in an object-decision task (Rogers, Lambon Ralph, Hodges, 2004). Asked to choose between a real elephant with large floppy ears of the kind that one only sees on elephants and an otherwise identical elephant with smaller, more typical ears taken from a monkey, severe SD patients will tend to choose the pseudo-elephant over the real one.

A SINGLE SYSTEM FOR KNOWLEDGE OF BOTH OBJECTS AND WORDS Clearly, semantic dementia patients show a conjunction of semantic and lexical deficits. These deficits not only correlate (Graham et al., 2000; Patterson et al.,

2006; Woollams, Lambon Ralph, Plaut, & Patterson, 2007) but also exhibit similar characteristics. In semantic tasks, the patients tend to lose knowledge of specific and idiosyncratic properties of objects while retaining and overextending knowledge of general and typical properties. Similarly, in lexical tasks such as word reading, word spelling, and lexical decision, they tend to lose knowledge of atypical items and mappings while retaining and overextending knowledge of the typical (as regularization errors). The strong correlation of these deficits and the similarity of their nature suggest that the tasks all depend on the same set of processing structures that also underlie semantic processing. On this view, it is damage to these structures that produces both the semantic deficits and the parallel deficits in lexical tasks seen in SD patients.

Dilkina, McClelland, and Plaut (2008) developed a connectionist model implementing just such a single-system approach to semantic and lexical processing and showed that it could reproduce the convergent pattern of semantic and lexical deficits seen in SD. The simulation focused primarily on one "semantic" task—picture naming—and one "lexical" task, word reading. The model builds on the one used by Rogers, Lambon Ralph, Garrard, and colleagues (2004). In place of individual units to represent printed or spoken words, Dilkina and colleagues used patterns of activation over orthographic (letter) and phonological (sound) units for the spellings and sounds of words, respectively (figure 72.11). Like Rogers and colleagues, the model uses a single cross-modal level of representation that integrates all types of information about both words and objects and thus corresponds to the convergent semantic representations of




FIGURE 72.11 Network architecture used to simulate a single integrative system for semantic and lexical processing. (Reprinted with permission from Dilkina, McClelland, & Plaut, 2008, figure 3, p. 143. Permission pending.)

the earlier model. The model also builds on earlier work (Plaut et al., 1996; Seidenberg & McClelland, 1989) by including a “direct” route between spelling and sound, in addition to bidirectional connections between both spelling and sound and the semantic layer. Although the direct route tends to specialize in capturing typical spelling-to-sound correspondences whereas the pathway through the integrative layer tends to specialize in idiosyncratic word-specific information, the partitioning is not absolute, and neither pathway corresponds to a strict rule system or a strictly lexical system.

The model was trained to map among four surface representations: (1) visual representations of what entities (objects and animals) look like; (2) action representations of how one may interact with these entities; (3) phonological representations of their names; and (4) orthographic representations of these names. The visual and action representations were binary vectors based on probabilistic category prototypes similar to those used in Rogers, Lambon Ralph, Garrard, and colleagues (2004). The name representations were simple onset-vowel-coda patterns that approximate English spelling-sound consistencies. After the network was trained to map from either visual or orthographic input to produce all four surface patterns, the semantic layer was progressively damaged to simulate semantic dementia.

The model exhibited the overall characteristic performance of SD patients—strong frequency effects in both tasks, and a frequency-by-typicality interaction in reading, as well as a high correlation between naming and reading of irregular words. In addition, the model was applied to the specific pattern of deficit seen in five patients tested with the same set of materials, including the case of an SD patient (patient EM, Blazeley, Coltheart, & Casey, 2005) who showed spared reading of low-frequency exception words in spite of a fairly profound impairment in standard semantic tasks. The model was able to fit the specific pattern of reading and naming data observed in all five patients, including patient EM.

Patients like EM have been used to argue against the single-system account, and thus it is important to understand how the model was able to address this patient, while at the same time addressing the other four, more typical, cases. The basis for this approach lies in incorporating the assumption that there are both premorbid and postmorbid *individual differences* that can contribute to the detailed pattern of performance seen in individual SD patients. Dilkina and colleagues (2008) focused on three such factors—premorbid experience with reading, premorbid capacity of the neural substrate mapping visual word form to phonological word form, and the spatial distribution of the lesion. Each of these factors was motivated by previous literature strongly suggesting that people indeed vary along these dimensions. The three factors were independently manipulated in the model.

Each of them significantly and independently affected the relative robustness of naming and reading. Notably, the model was able to successfully fit the SD dissociation case EM by manipulations that made reading relatively more robust to damage than naming. Notably, an experience manipulation alone was sufficient. Even though the model posits that naming and reading involve a single underlying system, greater premorbid experience with reading makes this system’s reading performance more robust under damage. In such cases, while naming declines quickly, the decline in reading may be delayed. Similar effects occur with other manipulations that increase reading robustness or that distribute the lesion toward connections from visual more than orthographic input. The account predicts that, as patients like EM progress in their illness, the deficit will eventually affect reading in all cases. Where data are available to test this prediction, it has held up to date (Woolams et al., 2007).

In summary, Dilkina and colleagues (2008) provide a theoretical and a computational account of how semantic and lexical deficits arise within a single system, and how they may appear to dissociate in some cases. While this work does not rule out a separate-system account, it shows that it is not necessary to postulate two systems to explain a handful of dissociation cases. Moreover, a single-system approach seems much more suitable in light of the large body of evidence showing a highly consistent SD profile where semantic and lexical abilities decline together, where there is a distinct frequency-by-typicality interaction in both domains, and where compromised performance results in homologous types of errors.

Roles of the ATL and other brain areas in semantic cognition

As noted throughout this chapter, our approach to understanding the nature and neural basis of conceptual knowledge has deliberately spanned a wide range of approaches and sources of data. In the preceding sections we reviewed a selection of these, including our overarching theoretical framework for considering conceptual knowledge and how this has been implemented in various PDP models addressing two core sources of empirical data—namely, the development of concepts in children and the structured degradation of concepts observed in semantic dementia. In this final section we review other sources of evidence bearing on our hypothesis that the anterior temporal lobes (ATLs) provide a hub over which conceptual knowledge is represented. First we review convergent evidence for the notion that the ATL is a critical part of the wider brain network that supports semantic cognition. Then we broaden the theoretical canvass to include the role of other brain regions. This allows us to consider a wider range of brain mechanisms that contribute

11

to semantic cognition, construed broadly as the task-modulated use of semantic knowledge to guide behavior.

CONVERGENT EVIDENCE FOR THE ROLE OF THE ATL IN SEMANTIC MEMORY There is considerable debate about the putative role of different brain regions in tasks requiring use of semantic knowledge (Hickok & Poeppel, 2007; Martin, 2007; Patterson et al., 2007; Wise, 2003). As noted earlier, in SD a selective semantic impairment is paired with relatively circumscribed atrophy of the anterior, inferolateral temporal lobes, bilaterally. Thus, as already articulated, the simplest and most obvious hypothesis is that the ATL areas are critical for semantic memory (Lambon Ralph & Patterson, 2008; Patterson et al.; Rogers, Lambon Ralph, Garrard et al., 2004). Given that SD is a neurodegenerative condition, there is no absolute boundary to the damage, and there is always the possibility that subthreshold damage or dysfunction due to invading pathology occurs elsewhere and that it is this more subtle, widespread damage that is the root of the patients' semantic impairment (Martin, 2007). It is critically important, therefore, to derive convergent evidence about the putative role of ATL regions in conceptual knowledge.

Converging evidence comes in three forms: other patient groups, functional neuroimaging, and transcranial magnetic stimulation (TMS). Other neurological conditions do produce semantic impairment when damage affects the same bilateral temporal lobe regions as semantic dementia. These include Alzheimer's disease (Hodges, Salmon, & Butters, 1992) and herpes simplex virus encephalitis (Lambon Ralph et al., 2007; Noppeney et al., 2007), although the more widespread brain damage associated with these neurological diseases leads to additional cognitive and memory impairments (Lambon Ralph & Patterson, 2008). The functional neuroimaging literature provides a slightly complex picture. If one primarily looks at the fMRI literature, there is a distinct lack of evidence for our hypothesis: fMRI studies of semantic memory or comprehension rarely find activation in anterior temporal lobe regions (Devlin et al., 2000; Garavan, Ross, Li, & Stein, 2000). While there may be important task design issues in some of these studies, the failure to find anterior temporal lobe activation reflects, at least in part, fMRI signal loss and distortion that is particularly pronounced in orbitofrontal cortex and the inferior and polar aspects of the temporal lobes (Devlin et al.; Wise, 2003). Functional neuroimaging that utilizes PET does detect semantically related activation in the anterior temporal lobes, even when the same experiment conducted in fMRI does not (Devlin et al.). Likewise, semantic-related processing in the ATL has been observed in normal participants by using MEG, irrespective of whether the stimulus is presented in the auditory or visual modality (Marinkovic et al., 2003), matching findings from early PET-based studies

based on pictorial or verbal input (Vandenberghe, Price, Wise, Josephs, & Frackowiak, 1996).

Of course, the areas differentially activated in imaging studies do not imply a *necessary* role (Price & Friston, 2002). Given the potential doubt over the neuropsychological data, we have recently initiated a new line of investigation that uses offline, repetitive transcranial magnetic stimulation (rTMS) to probe the role of ATL in neurologically intact participants (Lambon Ralph, Pobric, & Jefferies, 2008; Pobric, Jefferies, & Lambon Ralph, 2007). By using timed versions of the semantic tasks used in SD studies, we have been able to compare the pattern observed in the patients with that seen in normal participants post rTMS. The results closely mirror the characteristics of semantic dementia. For example, ATL rTMS produces a temporary slowing of responses on semantically related tasks (e.g., synonym judgment) but not other cognitive tasks matched for overall difficulty (e.g., number judgment). The same stimulation also affects expressive tasks (picture naming but not number reading is slowed). Intriguingly, like the SD patients, a greater effect was observed for identifying concepts at a specific level (e.g., *golden retriever*) than at a basic level (e.g., *dog*). The relative role of the left versus the right ATL can also be probed using this rTMS method. In one study we compared left versus right ATL rTMS on the same synonym judgment task. A comparable slowing of semantic decision times was observed, indicating that both left and right ATL support semantic memory (Lambon Ralph et al.). This pattern was replicated in a further study on a test of semantic association (e.g., between pyramids and palm trees; Bozeat et al., 2000; Howard & Patterson, 1992). In these assessments, stimuli are presented either as pictures or as written words. Repetitive TMS to either left or right ATL produced equivalent slowing on both the verbal and nonverbal versions of the task (Pobric, Jefferies, & Lambon Ralph, submitted). 12

These convergent results are all in keeping with our hypothesis that the ATL lobes jointly provide an amodal hub for semantic knowledge. However, there are still some puzzling data. Perhaps the most striking results come from patients with ATL resection for intractable epilepsy, who are not clinically associated with a postsurgical semantic impairment, at least not to the same degree as SD patients (Hermann, Davies, Foley, & Bell, 1999). Future studies of semantic memory that directly compare SD and TLE-resection patients are required to understand if the TLE data are truly inconsistent with our hypothesis. Most of the literature on the sequelae of temporal lobe resection is focused upon episodic memory impairment and anomia (which might itself reflect subtle semantic impairment; Lambon Ralph, McClelland, Patterson, Galton, & Hodges, 2001), and semantic memory is rarely formally tested (Giovagnoli, Erbetta, Villani, & Avanzini, 2005). Where semantic perfor-

mance has been assessed, studies have found subtle multimodal impairments both in unoperated TLE patients (Giovagnoli et al., 2005) and in patients after temporal lobe resection (Wilkins & Moscovitch, 1978). Furthermore, temporal lobe resection is a unilateral procedure, but SD patients have bilateral temporal lobe atrophy; it may be that bilateral damage is required to produce significant semantic impairment. It must also be noted that localization of function is complicated in these patients because long-standing epilepsy might lead to changes in neural organization. Indeed, recent imaging studies have shown that white matter connectivity and neurotransmitter function are significantly altered in this condition (Hammers et al., 2003; Powell et al., 2007). In addition, there might be some postsurgical reorganization that is less likely in neurodegenerative conditions when the brain is subjected to constant brain injury (Welbourne & Lambon Ralph, 2005). Consistent with this hypothesis, Wilkins and Moscovitch (1978) found a negative correlation between the severity of semantic impairment and time postsurgery.

BEYOND THE ATL: THE ROLE OF OTHER BRAIN REGIONS IN SEMANTIC COGNITION As articulated previously, our theoretical framework proposes one hypothesis about the way in which semantic knowledge is acquired through development and how it breaks down to produce the multimodal semantic impairments observed in semantic dementia and other ATL-focused neurological diseases. However, full-fledged semantic cognition—defined here as the adequate use of semantic knowledge to guide complex behavior—requires not only the ability to activate stored information from all modalities. It also requires the ability to shape or regulate the activation of task- and time-relevant information in order to produce flexible and appropriate behavior. Some kind of regulatory process is critical: we store a wealth of information about the meanings of words and objects, but frequently only a subset of this knowledge is required for a task—indeed, other aspects of knowledge may actually be inappropriate and unhelpful. As an example, consider the radically different uses of the same knife in make a sandwich: these can include piercing a package, slicing bread, meat, or cheese; scooping and spreading mustard or mayonnaise on the sandwich; and so on. Specific aspects of the knife's properties (and ways of holding and manipulating it) must be brought to the fore, one by one, while the most commonly listed property of cutting has to be inhibited in many of these activities. Indeed, in the case of scooping, the canonical function of the knife has to be disregarded altogether and replaced by a substituted function in place of another object (spoon). In sum, in addition to the acquisition and activation of conceptual knowledge, the ability to regulate and shape is critical to any complete account of semantic cognition.

The distinction between semantic representations and control processes that regulate processes acting on these representations helps to resolve a puzzle highlighted by a comparison of different, semantically impaired patient groups (i.e., patients who fail both verbal and nonverbal semantic tasks). Patients with ATL damage are not the only ones to exhibit poor semantic performance across different modalities. Indeed, it is possible to find a subset of aphasic patients who have multimodal semantic impairments (Chertkow, Bub, Deaudon, & Whitehead, 1997; Jefferies & Lambon Ralph, 2006) arising from temporoparietal or prefrontal damage rather than ATL damage. We refer to this pattern as *semantic aphasia* (SA; Jefferies, Patterson, & Lambon Ralph, 2008). By directly comparing semantic aphasia and semantic dementia, we have been able to demonstrate that each group's failure on semantic tasks is qualitatively different and that they should not be considered as the same type of impairment. We have hypothesized that the patient groups reflect the two primary ingredients in semantic cognition: semantic dementia reflects a degradation of the core conceptual knowledge, whereas semantic aphasia arises from a deficit in the regulation of semantic cognition. We find that SD patients are highly consistent across different semantic tasks: patients who retained knowledge of an item in one task were typically able to demonstrate this knowledge in all other tasks. In contrast, SA patients show significant correlations/consistency only between different versions of the same semantic task (e.g., judgments of semantic association for words and pictures). Unlike SD, the SA patients' ability to retrieve information is inconsistent when tested across tasks with different semantic control demands (e.g., judgments of semantic association versus word-picture matching). Moreover, SA patients' ability to make semantic judgments can be predicted by how readily the relevant semantic dimension can be discerned and competitors rejected. For such patients, cues or constraints provided by the examiner can boost their performance considerably. The patients' errors in picture naming provide further. The SD patients make frequent coordinate and superordinate semantic errors (such as saying "dog" or "animal" for goat). The SA patients also make associative errors (e.g., producing the response "nuts" for squirrel); these responses are virtually never seen in SD. These errors suggest that the SA patients retain a considerable amount of knowledge about unnamed targets (in order to be able to generate such errors) and suggest that their difficulty lies in directing activation toward the correct name and away from irrelevant associations.

These patient studies provide a direct convergence with fMRI studies of semantic processing in normal participants. These studies consistently implicate prefrontal cortex and the temporoparietal junction in tasks requiring controlled semantic processing—for example, when a particular aspect of meaning must be selected or when there is strong

13

competition from alternative responses (Thompson-Schill, Desposito, Aguirre, & Farah, 1997; Wagner, Pare-Blagoev, Clark, & Poldrack, 2001). It is possible that the relevant control or shaping processes may underpin executive/attentional functions more generally, as these same regions are commonly activated in a variety of tasks requiring cognitive control (Garavan et al., 2000; Peers et al., 2005). In keeping with this hypothesis, the SA patients tend to fail executive/attentional tasks even when they do not involve semantic information (Jefferies & Lambon Ralph, 2006).

Conclusion

This article has reviewed research spanning a wide range of research approaches. Behavioral investigations of developing children and neuropsychological patients, computational modeling investigations, and investigations of brain activity in healthy human volunteers using noninvasive imaging methods and TMS have all been used to support an overall account of the nature of semantic knowledge, its development and disintegration, and its instantiation in networks of interconnected areas of the brain. The approach has had some success in linking research from all these different methods under a common theoretical framework based on the principles of parallel distributed processing and in providing the stimulus for a considerable body of ongoing research. More work needs to be done to flesh out the theory and to better understand how activation of semantic and other forms of knowledge thought to depend on the anterior temporal lobes is influenced by activations in other brain areas.

REFERENCES

- BENEDET, M., PATTERSON, K., GOMEZ-PASTOR, I., & DE LA ROCHA, M. L. G. (2006). "Non-semantic" aspects of language in semantic dementia: As normal as they're said to be? *Neurocase*, *12*, 15–26.
- BLAZELY, A., COLTHEART, M., & CASEY, B. J. (2005). Semantic impairment with and without surface dyslexia: Implications for models of reading. *Cogn. Neuropsychol.*, *22*, 695–717.
- BOZEAT, S., LAMBON RALPH, M. A., GRAHAM, K. S., PATTERSON, K., WILKIN, H., ROWLAND, J., et al. (2003). A duck with four legs: Investigating the structure of conceptual knowledge using picture drawing in semantic dementia. *Cogn. Neuropsychol.*, *20*, 27–47.
- BOZEAT, S., LAMBON RALPH, M. A., PATTERSON, K., GARRARD, P., & HODGES, J. R. (2000). Non-verbal semantic impairment in semantic dementia. *Neuropsychologia*, *38*, 1207–1215.
- CARAMAZZA, A. (1997). How many levels of processing are there? *Cogn. Neuropsychol.*, *14*(1), 177–208.
- CHERTKOW, H., BUB, D., DEAUDON, C., & WHITEHEAD, V. (1997). On the status of object concepts in aphasia. *Brain Lang.*, *58*, 203–232.
- 14 CIPOLOTTI, L., & WARRINGTON, E. K. (1995). Semantic memory and reading abilities: A case report. *J. Int. Neuropsychol. Soc.*, *1*, 104–110.
- COLTHEART, M., RASTLE, K., PERRY, C., LANGDON, R. J., & ZIEGLER, J. C. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychol. Rev.*, *108*, 204–256.
- DEVLIN, J. T., RUSSELL, R. P., DAVIS, M. H., PRICE, C. J., WILSON, J., MOSS, H. E., et al. (2000). Susceptibility-induced loss of signal: Comparing PET and fMRI on a semantic task. *Neuroimage*, *11*, 589–600.
- DILKINA, K., McCLELLAND, J. L., & PLAUT, D. C. (2008). A single-system account of semantic and lexical deficits in five semantic dementia patients. *Cogn. Neuropsychol.*, *25*(2), 136–164.
- EBERHARDT, J. L., GOFF, P. A., PURDIE, V. J., & DAVIES, P. G. (2004). Seeing black: Race, crime, and visual processing. *J. Pers. Soc. Psychol.*, *87*, 876–963.
- FUNNELL, E. (1996). Response biases in oral reading: An account of the co-occurrence of surface dyslexia and semantic dementia. *Q. J. Exp. Psychol. A*, *49*(2), 417–446.
- GARAVAN, H., ROSS, T. J., LI, S. J., & STEIN, E. A. (2000). A parametric manipulation of central executive functioning. *Cereb. Cortex*, *10*, 585–592.
- GELMAN, R., & WILLIAMS, E. M. (1998). Enabling constraints for cognitive development and learning: A domain specific epigenetic theory. In D. Kuhn & R. Siegler (Eds.), *Handbook of child psychology*, Vol. 2: *Cognition, perception and development* (5th ed., pp. 575–630). New York: John Wiley and Sons.
- GIOVAGNOLI, A. R., ERBETTA, A., VILLANI, F., & AVANZINI, G. (2005). Semantic memory in partial epilepsy: Verbal and non-verbal deficits and neuroanatomical relationships. *Neuropsychologia*, *43*, 1482–1492.
- GRAHAM, N. L., PATTERSON, K., & HODGES, J. R. (2000). The impact of semantic memory impairment on spelling: Evidence from semantic dementia. *Neuropsychologia*, *38*, 143–163.
- HAMMERS, A., KOEPP, M. J., RICHARDSON, M. P., HURLEMANN, R., BROOKS, D. J., & DUNCAN, J. S. (2003). Grey and white matter flumazenil binding in neocortical epilepsy with normal MRI. A PET study of 44 patients. *Brain*, *126*, 1300–1318.
- HERMANN, B., DAVIES, K., FOLEY, K., & BELL, B. (1999). Visual confrontation naming outcome after standard left anterior temporal lobectomy with sparing versus resection of the superior temporal gyrus: A randomized prospective clinical trial. *Epilepsia*, *40*, 1070–1076.
- HICKOK, G., & POEPEL, D. (2007). Opinion—The cortical organization of speech processing. *Nat. Rev. Neurosci.*, *8*, 393–402.
- HINTON, G. E. (1981). Implementing semantic networks in parallel hardware. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 161–187). Hillsdale, NJ: Erlbaum.
- HODGES, J. R., GRAHAM, N., & PATTERSON, K. (1995). Charting the progression of semantic dementia: Implications for the organisation of semantic memory. *Memory*, *3*, 463–495.
- HODGES, J. R., & PATTERSON, K. (2007). Semantic dementia: A unique clinicopathological syndrome. *Lancet Neurol.*, *6*, 1004–1014.
- HODGES, J. R., PATTERSON, K., OXBURY, S., & FUNNELL, E. (1992). Semantic dementia: Progressive fluent aphasia with temporal lobe atrophy. *Brain*, *115*, 1783–1806.
- HODGES, J. R., SALMON, D. P., & BUTTERS, N. (1992). Semantic memory impairment in Alzheimer's disease: Failure of access or degraded knowledge? *Neuropsychologia*, *30*, 301–314.
- HOWARD, D., & PATTERSON, K. (1992). *The Pyramids and Palm Trees Test: A test of semantic access from words and pictures*. Bury St. Edmunds, UK: Thames Valley Test Company.
- JEFFERIES, E., & LAMBON RALPH, M. A. (2006). Semantic impairment in stroke aphasia vs. semantic dementia: A case-series comparison. *Brain*, *129*, 2132–2147.

- JEFFERIES, E., PATTERSON, K., & LAMBON RALPH, M. A. (2008). Deficits of knowledge vs. executive control in semantic cognition: Insights from cued naming. *Neuropsychologia*, *46*, 649–658.
- KEIL, F. C. (1979). *Semantic and conceptual development: An ontological perspective*. Cambridge, MA: Harvard University Press.
- KEIL, F. (1991). The emergence of theoretical beliefs as constraints on concepts. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind: Essays on biology and cognition*. Hillsdale, NJ: Erlbaum.
- LAMBON RALPH, M. A., LOWE, C., & ROGERS, T. T. (2007). Neural basis of category-specific semantic deficits for living things: Evidence from semantic dementia, HSVE and a neural network model. *Brain*, *130*, 1127–1137.
- LAMBON RALPH, M. A., McCLELLAND, J. L., PATTERSON, K., GALTON, C. J., & HODGES, J. R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *J. Cogn. Neurosci.*, *13*, 341–356.
- LAMBON RALPH, M. A., & PATTERSON, K. (2008). Generalisation and differentiation in semantic memory: Insights from semantic dementia. *Ann. NY Acad. Sci.*, *1124*, 61–76.
- LAMBON RALPH, M. A., POBRIC, G., & JEFFERIES, E. (2008). Conceptual knowledge is underpinned by the temporal pole bilaterally: Novel data from rTMS. *Cereb. Cortex*.
- 15 LEVELT, W. J. M. (1989). *Speaking, from intention to articulation*. Cambridge, MA: MIT Press.
- MANDLER, J. M., & McDONOUGH, L. (1993). Concept formation in infancy. *Cogn. Dev.*, *8*, 291–318.
- MARINKOVIC, K., DHOND, R. P., DALE, A. M., GLESSNER, M., CARR, V., & HALGREN, E. (2003). Spatiotemporal dynamics of modality-specific and supramodal word processing. *Neuron*, *38*, 487–497.
- MARTIN, A. (2007). The representation of object concepts in the brain. *Annu. Rev. Psychol.*, *58*, 25–45.
- MARTIN, A., & CHAO, L. L. (2001). Semantic memory in the brain: Structure and processes. *Curr. Opin. Neurobiol.*, *11*, 194–201.
- McCLELLAND, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 8–45). New York: Oxford University Press.
- McCLELLAND, J. L., McNAUGHTON, B. L., & O'REILLY, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.*, *102*, 419–457.
- McCLELLAND, J. L., & ROGERS, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nat. Rev. Neurosci.*, *4*, 310–322.
- MERVIS, C. B. (1987). Child basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge, UK: Cambridge University Press.
- NOPPENEY, U., PATTERSON, K., TYLER, L. K., MOSS, H., STAMATAKIS, E. A., BRIGHT, P., et al. (2007). Temporal lobe lesions and semantic impairment: A comparison of herpes simplex virus encephalitis and semantic dementia. *Brain*, *130*, 1138–1147.
- O'REILLY, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Comput.*, *8*, 895–938.
- PATTERSON, K., & HODGES, J. R. (1992). Deterioration of word meaning: Implications for reading. *Neuropsychologia*, *30*(12), 1025–1040.
- PATTERSON, K., LAMBON RALPH, M. A., HODGES, J. R., & McCLELLAND, J. L. (2001). Deficits in irregular past-tense verb morphology associated with degraded semantic knowledge. *Neuropsychologia*, *39*, 709–724.
- PATTERSON, K., LAMBON RALPH, M. A., JEFFERIES, E., WOOLLAMS, A., JONES, R., HODGES, J. R., & ROGERS, T. T. (2006). “Presemantic” cognition in semantic dementia: Six deficits in search of an explanation. *J. Cogn. Neurosci.*, *18*(2), 169–183.
- PATTERSON, K., NESTOR, P. J., & ROGERS, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat. Rev. Neurosci.*, *8*, 976–987.
- PAUEN, S. (2002). The global-to-basic shift in infants’ categorical thinking: First evidence from a longitudinal study. *Int. J. Behav. Dev.*, *26*(6), 492–499.
- PEERS, P. V., LUDWIG, C. J. H., RORDEN, C., CUSACK, R., BONFIGLIOLI, C., BUNDESEN, C., et al. (2005). Attentional functions of parietal and frontal cortex. *Cereb. Cortex*, *15*, 1469–1484.
- PLAUT, D. C., McCLELLAND, J. L., SEIDENBERG, M. S., & PATTERSON, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychol. Rev.*, *103*, 56–115.
- POBRIC, G., JEFFERIES, E., & LAMBON RALPH, M. A. (submitted). Non-verbal semantic memory in the anterior temporal lobes: TMS evidence. 16
- POBRIC, G. G., JEFFERIES, E., & LAMBON RALPH, M. A. (2007). Anterior temporal lobes mediate semantic representation: Mimicking semantic dementia by using rTMS in normal participants. *Proc. Natl. Acad. Sci. USA*, *104*, 20137–20141.
- POWELL, H. W. R., PARKER, G. J. M., ALEXANDER, D. C., SYMMS, M. R., BOULBY, P. A., WHEELER-KINGSHOTT, C. A. M., et al. (2007). Abnormalities of language networks in temporal lobe epilepsy. *Neuroimage*, *36*, 209–221.
- PRICE, C. J., & FRISTON, K. J. (2002). Degeneracy and cognitive anatomy. *Trends Cogn. Sci.*, *6*, 416–421.
- QUILLIAN, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 227–270). Cambridge, MA: MIT Press.
- ROGERS, T. T., LAMBON RALPH, M. A., GARRARD, P., BOZEAT, S., McCLELLAND, J. L., HODGES, J. R., et al. (2004). The structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychol. Rev.*, *111*, 205–235.
- ROGERS, T. T., LAMBON RALPH, M. A., HODGES, J. R., & PATTERSON, K. (2004). Natural selection: The impact of semantic impairment on lexical and object decision. *Cogn. Neuropsychol.*, *21*(2/3/4), 331–352.
- ROGERS, T. T., & McCLELLAND, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- ROGERS T. T. & PATTERSON, K. (2007). Object categorization: Reversals and explanations of the basic-level advantage. *J. Exp. Psychol. Gen.*, *136*, 451–469.
- RUMELHART, D. E. (1990). Brain style computation: Learning and generalization. In S. F. Zornetzer, J. L. Davis, & C. Lau (Eds.), *An introduction to neural and electronic networks* (pp. 405–420). San Diego: Academic Press.
- RUMELHART, D. E., & McCLELLAND, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland & D. E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 216–271). Cambridge, MA: MIT Press. 17

- RUMELHART, D. E., MCCLELLAND, J. L., and THE PDP RESEARCH GROUP. (1986). *Parallel distributed processing: Explorations in the micro-structure of cognition* (Vols. 1 & 2). Cambridge, MA: MIT Press.
- RUMELHART, D. E., & TODD, P. M. (1993). Learning and connectionist representations. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 3–30). Cambridge, MA: MIT Press.
- [18] SAFFRAN, E. M. (2000). Word retrieval and its disorders. *Cogn. Neuropsychol.*, *16*, 777–790.
- SAFFRAN, E. M., COSLETT, H. B., MARIN, N., & BORONAT, C. B. (2003). Access to knowledge from pictures but not words in a patient with progressive fluent aphasia. *Lang. Cogn. Process.*, *18*(5/6), 725–757.
- [19] SCHWARTZ, M. F., SAFFRAN, E. M., & MARIN, O. S. M. (1980). Fractionating the reading process in dementia: Evidence for word-specific print-to-sound associations. In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia*, London: Routledge and Kegan Paul.
- SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychol. Rev.*, *96*, 523–568.
- SNOWDEN, J. S., GOULDING, P. J., & NEARY, D. (1989). Semantic dementia: A form of circumscribed temporal atrophy. *Behav. Neurol.*, *2*, 167–182.
- SQUIRE, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.*, *99*, 195–231.
- THOMPSON-SCHILL, S. L., DESPOSITO, M., AGUIRRE, G. K., & FARAH, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *Proc. Natl. Acad. Sci. USA*, *94*, 14792–14797.
- VANDENBERGHE, R., PRICE, C., WISE, R., JOSEPHS, O., & FRACKOWIAK, R. S. J. (1996). Functional-anatomy of a common semantic system for words and pictures. *Nature*, *383*, 254–256.
- WAGNER, A. D., PARE-BLAGOEV, E. J., CLARK, J., & POLDRACK, R. A. (2001). Recovering meaning: Left prefrontal cortex guides controlled semantic retrieval. *Neuron*, *31*, 329–338.
- WARD, J., STOTT, R., & PARKIN, A. J. (2000). The role of semantics in reading and spelling: Evidence for the “summation hypothesis.” *Neuropsychologia*, *38*, 1643–1653.
- WARRINGTON, E. K. (1975). Selective impairment of semantic memory. *Q. J. Exp. Psychol.*, *27*, 635–657.
- WELBOURNE, S. R., & LAMBON RALPH, M. A. (2005). Subtracting subtractivity? A connectionist account of recovery in single word reading following brain damage. *Cog. Affect. Behav. Neurosci.*, *5*, 77–92.
- WILKINS, A., & MOSCOVITCH, M. (1978). Selective impairment of semantic memory after temporal lobectomy. *Neuropsychologia*, *16*, 73–79.
- WISE, R. (2003). Language systems in normal and aphasic human subjects: Functional imaging studies and inferences from animal studies. *Br. Med. Bull.*, *65*, 95–119.
- WOOLLAMS, A. M., LAMBON RALPH, M. A., PLAUT, D. C., & PATTERSON, K. (2007). SD-squared: On the association between semantic dementia and surface dyslexia. *Psychol. Rev.*, *114*(2), 316–339.
- WOOLLAMS, A. M., COOPER-PYE, E., HODGES, J. R., & PATTERSON, K. (2008). Anomia: A doubly typical signature of semantic dementia. *Neuropsychologia*, *46*, 2503–2514.

AUTHOR QUERY FORM

Dear Author

During the preparation of your manuscript for publication, the questions listed below have arisen. Please attend to these matters and return this form with your proof.

Many thanks for your assistance.

Query References	Query	Remarks
1.	Au: Please add departments.	
2.	Au: We don't have Hinton '89	
3.	Au: OK to lowercase "is," "can," and "has," as in the figure? Or should they	
4.	Au: Slight edits here	
5.	Au: Not to be picky, but the needles of a pine are leaves. They're just not broad leaves.	
6.	Au: OK as edited? for ATL?	
7.	Au: This is the correct Rogers . . . '04, isn't it?	
8.	Au: "and 72.10D" OK? (see p. 9-201)	
9.	Au: We don't have Saffran '03. We have Saffran '00 and Saffran et al. '03.	
10.	Au: We don't have Pinker '01	
11.	Au: OK as edited?	
12.	Au: Pls update Pobric . . . submitted if possible	
13.	Au: An omission? ". . . provide further evidence," perhaps?	
14.	Au: Cipolotti & Warrington was not cited. OK to delete?	
15.	Au: Lambon Ralph, Pobric . . . : Pls give vol. & pages.	
16.	Au: Pls update Pobic . . . submitted if possible.	
17.	Au: Rumelhart & McClelland: "(1986)" per p. 9-161. OK?	
18.	Au: Saffran '00, Saffran . . . '03; and Schwartz . . . '80 weren't cited. (maybe one or more should have been)	
19.	Au: Seidenberg . . . : Pls give pages for chapter in book.	
20.	Au: Letters are used in text. We added them to figure & caption. OK?	
21.	Au: OK to delete "115"?	