

Object Categorization: Reversals and Explanations of the Basic-Level Advantage

Timothy T. Rogers

University of Wisconsin—Madison

Karalyn Patterson

Medical Research Council Cognition and Brain Sciences Unit

People are generally faster and more accurate to name or categorize objects at the basic level (e.g., *dog*) relative to more general (*animal*) or specific (*collie*) levels, an effect replicated in Experiment 1 for categorization of object pictures. To some, this pattern suggests a dual-process mechanism, in which objects first activate basic-level categories directly and later engage more general or specific categories through the spread of activation in a processing hierarchy. This account is, however, challenged by data from Experiment 2 showing that neuropsychological patients with impairments of conceptual knowledge categorize more accurately at superordinate levels than at the basic level—suggesting that knowledge about an object’s general nature does not depend on prior basic-level categorization. The authors consider how a parallel distributed processing theory of conceptual knowledge can reconcile the apparent discrepancy. This theory predicts that if healthy individuals are encouraged to make rapid categorization responses, the usual basic > general advantage should also reverse, a prediction tested and confirmed in Experiment 3. Implications for theories of visual object recognition are discussed.

Keywords: semantic memory, categorization, object recognition, basic level, semantic dementia

This article addresses a paradox in the study of human conceptual knowledge that arises when two separate literatures are considered side-by-side. The first literature concerns the *basic-level advantage*: Rosch et al.’s (1976) discovery that, of all the various categories to which a given item could belong (e.g., “collie,” “dog,” “canine,” “mammal,” “animal,” “pet”), some appear to be more readily accessible to the human mind than others. Among other things, Rosch and colleagues (Mervis & Rosch, 1981; Rosch, 1978; Rosch et al., 1976) showed that people are faster to categorize most items at an intermediate level of specificity (such as “dog”) than at more general (i.e., “animal”) and more specific (i.e., “collie”) levels; that in free speech, people prefer to use such names over more general or specific labels; and that the first nouns children typically acquire are also at this level. Such basic-level effects are among the best known and most robust phenomena in the study of human categorization (Johnson & Mervis, 1997; Murphy & Brownell, 1985; Murphy & Lassaline, 1997; Tanaka & Taylor, 1991; Tversky & Hemenway, 1984).

The second literature concerns the gradual erosion of conceptual knowledge in a progressive neurological condition called *semantic*

dementia (SD; Garrard & Hodges, 2000; Hodges, Garrard, & Patterson, 1998; Snowden, Goulding, & Neary, 1989). Patients with SD exhibit a progressive and eventually profound deterioration of knowledge about the meanings of words and objects, together with a remarkable sparing of many other cognitive abilities until late in the course of the disease (Patterson & Hodges, 2000). The syndrome is associated with circumscribed atrophy of the anterior temporal cortex (Garrard & Hodges, 2000; Mummery et al., 2000; Nestor, Fryer, & Hodges, 2006). As Warrington (1975) noted in one of the first experimental investigations of SD, knowledge about general (i.e., superordinate) concepts appears to be considerably less vulnerable in the disorder than knowledge about more specific concepts. When sorting words or pictures, SD patients are more accurate when asked to sort on the basis of general criteria (e.g., “living” vs. “manmade”) than on basic-level criteria (e.g., “car” vs. “boat”; Hodges, Graham, & Patterson, 1995; Rogers, Lambon Ralph, Garrard, et al., 2004). When SD patients are asked to name pictures of objects, as the disease progresses, an object like a duck that typically yields subordinate-level names in normal individuals (i.e., *duck*) typically gives way to a basic-level response such as *bird*, which in turn typically gives way to a more general response such as *animal* (Hodges et al., 1995; Saffran & Schwartz, 1994; Schwartz, Marin, & Saffran, 1979; Warrington, 1975). In drawing tasks, patients frequently omit properties specific to a particular concept (e.g., the udder of a cow, the hump on a camel), but rarely do they omit properties common to the category superordinate (e.g., the eyes of the camel, the mouth of the cow; Bozeat et al., 2003). These and a host of other similar phenomena documented in SD (Adlam, Patterson, Rogers, Salmond, & Hodges, 2006; Bozeat, Lambon Ralph, Patterson, Garrard, & Hodges, 2000; Hodges, Spatt, & Patterson, 1999; Rogers, Lambon Ralph, Hodges, & Patterson, 2003, 2004) suggest that, as the conceptual knowledge system deteriorates, what is retained is increasingly restricted to the general and typical.

Timothy T. Rogers, Department of Psychology, University of Wisconsin—Madison; Karalyn Patterson, Cognition and Brain Sciences Unit, Medical Research Council.

We thank Matthew Lambon Ralph, Elizabeth Jefferies, and Roy Jones, who arranged for and conducted the testing of patients in Bath and Manchester, and we also thank Greg Murphy for his critical feedback that contributed substantially to this article. This work was supported by an Interdisciplinary Behavioral Science Centers for Mental Health Grant from the National Institute of Mental Health (MH64445) and by a grant from Research into Ageing and Dunhill Medical Trust.

Correspondence concerning this article should be addressed to Timothy T. Rogers, Department of Psychology, University of Wisconsin—Madison, 1202 West Johnson Street, Madison, WI 53706. E-mail: trogers@wisc.edu

Both sets of phenomena have been explained with reference to spreading-activation models derived from Collins and Quillian's (1969; Collins & Loftus, 1975; Quillian, 1968) influential theory, an effort that nicely illustrates the paradox. In spreading-activation theories, conceptual knowledge is stored as a system of propositions organized hierarchically in memory. Concepts—mental representations of categories—occupy nodes in the hierarchy, which are linked together by class-inclusion propositions called *ISA links*. For instance, the knowledge that a canary is a kind of bird is represented by connecting the node for *canary* to the node for *bird* with a class-inclusion (i.e., ISA) link; knowledge that birds are animals is stored by connecting the *bird* node to the *animal* node, and so on. Other facts are stored as predicates attached to the various nodes—for instance, to store the information that a canary can sing, the model attaches the predicate *can sing* to the *canary* node; and to store the information that all animals can move, the predicate *can move* is attached to the *animal* node. To make inferences about the properties of a given concept such as *canary*, the model first retrieves all of the predicates stored directly with the corresponding node (e.g. *can sing*); but activation then spreads upward along the ISA links so that the predicates attached to more inclusive concepts also get attributed to the probe concept. For *canary*, activation first spreads to the *bird* node, supporting the inference that the canary can fly, and then up to the *animal* node, supporting the inference that the canary can move.

In one influential account of basic-level effects, Jolicoeur, Gluck, and Kosslyn (1984) proposed that certain nodes within a Quillian-like processing hierarchy serve as “entry points” for probing the semantic network. Visual stimuli are first classified into one of these entry-level categories by means of a perceptual processing mechanism so that any information stored directly with the corresponding entry-level node becomes available earliest in processing. Additional information about the stimulus becomes available later, as activation spreads upward from the entry point toward more inclusive concepts or downward toward more specific concepts. Basic-level effects are observed for typical category members because the basic-level category nodes serve as the entry-point for such items: For instance, a visual stimulus such as a robin first activates the *bird* node, providing rapid access to the name “bird” and other typical bird properties (i.e., *has wings* and *can fly*). Retrieval of properties that the robin shares with all animals takes longer as it requires a search of the semantic network upward from the entry point. Retrieval of properties idiosyncratic to the robin takes longer either because nodes below the entry point must be searched or because more specific classification requires more extensive visual processing (Collin & McMullen, 2005; Murphy & Smith, 1982). Jolicoeur et al. (1984) further suggested that atypical category members fail to show a basic-level advantage because their entry-points are specific rather than basic. For example, the entry level for a picture of a penguin would be the node corresponding to *penguin* rather than the *bird* node that serves as the entry point for more typical birds.

Warrington (1975) also used the Collins and Quillian (1969) framework to explain how patients with SD retain knowledge of superordinate category information even as they lose knowledge of more specific concepts/features. She suggested that, in the healthy system, the hierarchy is probed from the top down—so that, to retrieve information about a canary, activation begins with a general category like *animal*, flows downward to *bird*, and finally to

canary. In SD, the disease process degrades the processing hierarchy from the bottom up—first destroying very specific nodes, then intermediate nodes, and only affecting the most general nodes in the most severe phases of the disease. This suggestion neatly explained her initial observations of a fine-to-coarse order of knowledge disintegration in SD and could be extended to account for the now comprehensive documentation of such effects in SD across a very broad variety of tasks.

The paradox, then, is this: To explain basic-level advantages, the entry-level theory proposes that visual stimuli are first classified at an intermediate level (for typical items) or a subordinate level (for atypical items), and that retrieval of more superordinate information requires the spread of activation through a processing hierarchy. Yet patients with SD appear to retain superordinate-level knowledge at a stage when they fail to retrieve intermediate- and specific-level information—suggesting that access to such superordinate information does not require prior activation of more specific entry-level concepts. To explain the latter result, Warrington (1975) proposed that the taxonomy is probed from the top down—but if this were the case, it is not clear why the intact semantic system typically retrieves information most rapidly at the basic level.

The contradictory findings are also puzzling from the perspective of the other prominent theory of basic-level effects, the *differentiation theory* proposed by Murphy and colleagues (Murphy, 1991, 2002; Murphy & Brownell, 1985; Murphy & Lassaline, 1997). This view contends that basic categories are not “privileged” in processing but that basic-level advantages arise as a consequence of their structural properties in a system that activates category representations at all levels of specificity in parallel. Specifically, basic-level categories are both *distinct* and *informative*. More general categories can be quite distinct from one another—for instance, the category “furniture” shares little in common with the category “animal”—but these categories are not very informative: Given the information that something is a piece of furniture, there are relatively few inferences one can make about its other properties (Rosch, Mervis, et al., 1976). More specific categories are very informative—when told something is a rocking chair, one can infer many of its properties—but they are not very distinctive: Rocking chairs share many properties with other kinds of chairs and so may be difficult to distinguish from them. Basic-level categories are both distinct and informative and so provide an especially useful means of “carving the world at its joints.” These differing structural properties may influence the speed with which information is retrieved across different levels. For instance, the superordinate level may be comparatively slow to activate because it provides a rather poor match to any particular stimulus; the subordinate level may provide a slow response because it requires resolution of competition amongst many similar alternatives; but the basic level, as in the Goldilocks story, will be “just right” (i.e., fastest because of its informative and distinctive properties).

The differentiation theory is appealing for several reasons: It can explain why basic-level advantages are not observed for atypical category members (Murphy & Brownell, 1985) and why, for some concepts (e.g., *chicken*), participants may be faster to categorize at both subordinate and superordinate levels than at the basic level (Rips, Shoben, & Smith, 1973). Furthermore its basic tenets have been tested in artificial category-learning experiments, in which the structural properties of the to-be-learned categories were ma-

nipulated by the experimenters, with results generally supporting the theory (Mervis & Crisafi, 1982; Murphy, 1991; Murphy & Smith, 1982). Perhaps most importantly, the differentiation theory explains why some categories are more rapidly accessed than others and links this explanation to Rosch's (1978) original insight that basic categories capture the most information with the least redundancy. Despite these virtues, observations from SD still pose a puzzle for the differentiation theory: If basic categories maximize information and distinctiveness, and if these characteristics make basic-level categories easier to activate in the healthy system, then why is basic-level information more vulnerable to semantic impairment than the less-informative superordinate category representations?

In this article we attempt to resolve this seeming contradiction, and in so doing will offer an explanation of the basic-level advantage that builds upon the differentiation theory but extends it in important respects. Our explanation is based upon a parallel distributed processing (PDP) theory of conceptual knowledge, described in previous work (McClelland & Rogers, 2003; Rogers et al., 2004; Rogers & McClelland, 2004). Rogers and McClelland (2004) showed how the key factors that give rise to basic-level effects under the differentiation theory—*informativeness* and *distinctiveness*—also produce basic-level advantages in a simple model implementation of the theory. In this sense, the PDP account is very similar in spirit to the differentiation theory. The particular mechanisms that give rise to this effect, however, differ somewhat from those alluded to in the differentiation theory; and here we consider whether these mechanisms provide a means of resolving the apparent paradox.

We begin by establishing experimentally that there really is a paradox. Although the extensive literature on SD strongly suggests that knowledge about broad categories is less vulnerable than knowledge about more specific categories, no study has previously tested whether the basic-level advantage observed in healthy adults is preserved, diminished, eliminated, or reversed in patients with SD. In Part 1, we replicate Rosch's (1976) finding of a basic-level advantage for category verification in a sample of healthy older adults. We then test category verification for the same materials in a case series of patients with SD and investigate how the basic-level effect changes with increasingly severe semantic impairment. In Part 2, we consider an explanation of basic-level effects arising from the PDP theory of conceptual knowledge, which suggests how such effects can be reconciled with the pattern of semantic impairment typically observed in SD (Rogers & McClelland, 2004). We show that this explanation makes counterintuitive predictions about the time course of normal processing in the category-verification task. These predictions differ from those offered by the entry-level theory of the basic-level advantage but are consistent with a differentiation-like explanation, couched within a PDP framework. In Part 3, we test the predictions in a new behavioral paradigm. Following these experiments, we consider implications of these findings for hypotheses about the interaction between mechanisms of visual categorization and semantic retrieval and processing.

Part 1: The Impact of Category Level on Normal and Semantically Impaired Category Verification

Experiment 1: Category Verification in Healthy Adults

The aim of the first experiment was to replicate the common finding of a basic-level advantage for category verification in a set

of stimuli for later use in patient testing. In this task, participants viewed a category name such as "bird," "robin," or "animal" on a computer screen, followed by a color photograph of a common object. Their task was to indicate by buttonpress, as quickly and as accurately as possible, whether the item in the photograph belonged to the category named. Category labels were either specific names (e.g., "robin," "yacht"), basic-level names (e.g., "bird," "boat"), or superordinate names (e.g., "animal," "vehicle"), and we measured the speed and accuracy of responses at these three different levels of specificity.

Method

Participants. Participants ($N = 28$) were 14 men and 14 women between the ages of 55 and 75 ($M = 67$) selected from the volunteer participant pool at the Medical Research Council (MRC) Cognition and Brain Sciences Unit.

Stimuli. Stimuli were 72 color photographs of animals and vehicles. They included a range of animals and vehicles for classification at the general level; a set of dogs, birds, cars, and boats for classification at the basic level; and three different photographs for each of eight specific categories: Pekinese, Labrador, kingfisher, robin, yacht, ferry, BMW, and Morris (a familiar make of car in the United Kingdom). Each photograph appeared once as a target and once as a distractor, yielding 144 trials total. In the specific conditions, distractors were always items from the same basic-level category as the target; for instance, if the category name was "yacht," the distractor was a photograph of a different kind of boat. In the basic-level condition, distractors were always from the same superordinate category as the target: For the category name "boat," the distractor was always another kind of vehicle. For the general condition, distractors were selected from a different superordinate category than the target (e.g., for "vehicle," the distractor could be an animal or plant).

Procedure. Words and pictures were presented to participants on a computer monitor running the freely available DMDX experimental-psychology software package. On each trial, participants viewed a fixation point followed by a category name presented at fixation for 1 s. The word was followed by a blank screen for 500 ms, and this in turn was followed by a matching or nonmatching color photograph that remained onscreen until a response was detected. Responses were recorded by buttonpress on a computer mouse; participants were instructed to press the left mouse button if they believed the item in the photograph was a member of the category named and the right button otherwise. All trials were ordered randomly for each participant. Reaction times and accuracy were recorded by DMDX.

Results

We used two measures to investigate the basic-level advantage in this task: reaction time for correct responses and d' as a measure of accuracy for discriminating targets from distractors in each condition. As shown in Figure 1, both measures showed a basic-level advantage. Repeated-measures analyses of variance (ANOVAs) with category level (specific, basic, or general) treated as a three-level within-subjects factor showed a main effect of category level on RT, $F(2, 54) = 5.0, p < .02, \eta_p^2 = 0.16$, and accuracy, $F(2, 54) = 4.0, p < .04, \eta_p^2 = 0.24$. Planned-




Figure 1. Mean reaction times (left) and discrimination accuracy measured as d' (right) in general, basic, and specific conditions of Experiment 1. Error bars indicate standard error of the mean in each condition.

comparisons paired-sample t tests showed that participants were significantly faster and more accurate in the basic-level than in the general level ($p < .005$ for speed, $p < .009$ for accuracy) and were significantly faster in the basic than in the specific condition ($p < .03$), with a trend toward greater accuracy in this contrast as well ($p < .08$). There were no reliable differences in speed or accuracy between specific and general conditions.

In sum, the results, not surprisingly, replicate a very robust result: Healthy participants are faster and more accurate to categorize visually presented objects at the basic level than at more general and more specific levels. For these particular stimuli, participants were equally fast and equally accurate in classifying at the general and specific levels, indicating that these conditions were matched for overall difficulty.

Experiment 2: Category Verification in SD

The aim of Experiment 2 was to determine whether the advantage for categorization at the basic level observed in healthy older adults in Experiment 1 persists in SD, or whether—as expected on the basis of previous studies of conceptual degradation in this condition—disease progression in fact engenders a reversal of the normal basic > general advantage. Eight patients with SD viewed the same set of word–photograph pairs used in the previous experiment and were asked to decide for each whether the photograph matched the word. To elicit the best performance possible, we did not subject the patients to any time pressure, and reaction times (RTs) were not taken. We then measured the accuracy of their decisions for specific, basic-level, and general category names.

Method

Participants. Participants were 8 patients diagnosed with SD according to both cognitive and neuroradiological criteria as outlined by Hodges, Patterson, Oxbury, and Funnell (1992), and who were identified through memory-disorder clinics in the United Kingdom. Basic neuropsychological data collected at approximately the same time period as the experimental testing session are shown in Table 1. Because of differences in the routine test

batteries administered in Cambridge, Bath, and Manchester, scores on some tests were not available within the same year of testing for patients BS and EK; in these cases, we have recorded scores taken approximately 2 years before assessment on the category-verification task. The table also lists means and standard deviations taken from approximately age-matched healthy controls, collected in previous experiments from volunteers in the MRC participant pool.

Scores on the Mini-Mental State Examination provide a general measure of cognitive status; all patients showed impairment on this measure. The Rey–Osterrieth Complex Figure Test (ROCF; Osterrieth, 1944; Rey, 1941) provides a basic measure of visuospatial abilities, as do the two subtests of the Visual Object and Space Perception (VOSP; Warrington & James, 1991) battery. All patients scored normally on these tasks, with the single exception of BS who fell just below the cutoff on the VOSP cube-analysis task. Delayed copy of the ROCF figure provides a measure of episodic memory; all but two of the more severe cases (BS and IB) performed normally on this task. Forward and backward digit span provide simple assessments of working memory capabilities; all patients were normal on these measures. Finally, the Wisconsin Card Sorting Test (Heaton, Grant, & Berg, 1995) provides an assessment of executive function. Three patients (GE, BS, and EK) were impaired relative to controls on this task; however, all three performed above the published cutoff for impairment in this age range (50–74, where scores exceeding 1 are considered within the normal range). In sum, all patients performed generally well on these cognitive measures, with reliable impairment observed in just 2 patients on the episodic memory task and in just 1 patient in one of three tests of visual perception.

In contrast to this generally good cognitive profile, all patients showed moderate to severe impairments on three common tests of semantic memory: picture naming for 64 line drawings of common objects; 10 alternative forced-choice word-to-picture matching for the same 64 items; and the picture version of the Pyramids and Palm Trees Test of semantic association. As indicated in the table, healthy controls typically perform near ceiling on these measures, and all 8 patients were reliably impaired relative to controls on all three semantic measures. The general neuropsychological data are

Table 1
Basic Neuropsychological Data on the 8 Patients Tested in Experiment 2

	Controls		Patients							
	<i>M</i>	<i>SD</i>	JM	DD	DV	GE	PF	BS	IB	EK
Demographics										
Sex			M	M	M	M	F	M	M	F
Age			67	72	65	50	67	71	57	63
General neuropsychology										
MMSE/30	28.8	0.9	24	22	23	–	24	25 ^a	23	26 ^a
Visual perception										
Rey copy/36	34.2	1.6	36	34	35	35	32	32	32	27
VOSP–dot count/10	9.9	0.3	10	10	10	10	10	10	10	9 ^a
VOSP–cube analysis/10	9.3	1.5	9	10	10	10	10	5	10	10 ^a
Memory										
Rey delayed copy/36	18.3	5.2	17	–	13	26	9	0	4	–
Forward digit span	7.1	0.9	7	7	7	7	8	6	8	5
Backward digit span	5.4	1.4	7	7	3.5	4	6	4	5	3
Executive function										
WCST/6	5.9	0.4	6	6	6	4	6	2	5	3
Semantics										
Naming/64	62.3	1.6	30	37	10.5	13	10	4	5	2
Word–Picture matching/64	63.8	0.4	61	55	47	32	29	27	19	15
PPT pics/52	51.2	1.4	46	45	30.5	34	27	33 ^a	33	39 ^a

Note. Dashes indicate that the patient was never tested. MMSE = Mini-Mental State Examination; VOSP = Visual Object and Space Perception Battery; WCST = Wisconsin Card Sort Task; PPT pics = Pyramids and Palm Trees Test, picture version.

^aIndicates scores taken more than 1 year prior to testing on the category verification task.

thus consistent with the view that SD represents a relatively pure impairment of semantic memory.

Materials. The same word–photograph pairs from the previous experiment were used. Each photograph was printed on a single sheet of A4 paper with the corresponding word printed above it. The 144 stimuli were arranged in pseudorandom order (ensuring that the same photograph did not appear twice within 5 successive trials) and bound. All patients viewed the stimuli in the same pseudorandom order.

Procedure. Participants sat across the table from the experimenter. For each trial, they were shown the word and picture while the experimenter read the word aloud, then he or she asked

whether the word matched the picture by saying, for instance, “See this? Is this a dog?” Patients were encouraged to respond “yes” or “no.” Responses were coded by the experimenter on a sheet of paper not visible to the patient.

Results

Figure 2 shows the d' scores for discriminating targets from distractors across specific, basic, and general conditions. The left side shows d' scores in each condition plotted against the severity of semantic impairment as assessed by scores on the word–picture matching test in Table 1. The data were smoothed by averaging the

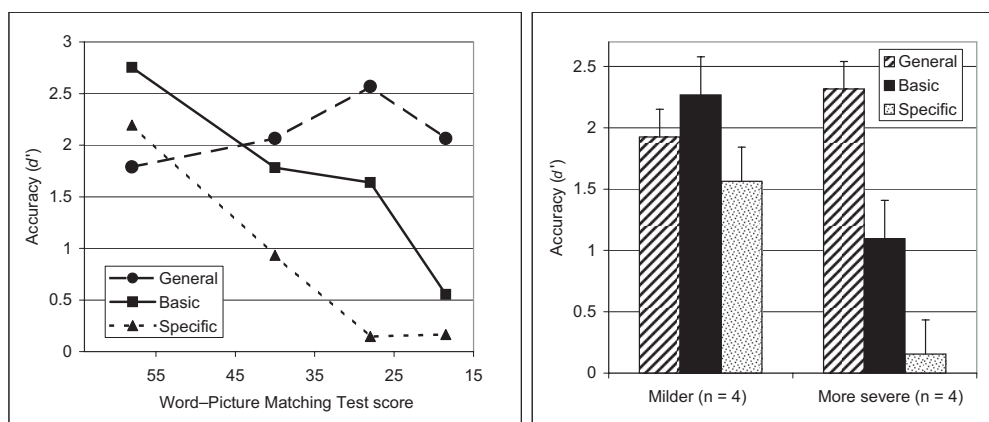


Figure 2. Left panel: Mean discrimination accuracy for semantic dementia patients in each condition plotted against the severity of semantic impairment as measured by the Word–Picture Matching Test (WPM). Each point represents the average of 2 patients with similar scores on the WPM test. Right panel: Mean discrimination accuracy for 4 milder and 4 more severely impaired patients. Error bars show standard error of the means.

scores for pairs of patients with comparably severe semantic impairments—thus each point in the plot reflects the average d' and word–picture matching scores for 2 patients. It is apparent that performance declined with the magnitude of the semantic impairment but that the slope of this decline differed for specific, basic, and general conditions. Whereas the milder patients showed an advantage for basic-level classification, performance declined for basic and specific conditions, but not for the general condition, so that more severe patients showed a substantial advantage for general classification. Only the 2 participants with the mildest impairments showed numerically better performance at the basic level than at the general level.

The right side of the figure presents the mean data for the 4 mildest and 4 more severe patients. By inspection, a small but clear basic-level advantage in the mild group changes to a substantial advantage for general classification in the more severe cohort. This observation was substantiated by a repeated-measures ANOVA in which category level was treated as a within-subjects factor and extent of semantic impairment (mild or severe) was treated as a between-subjects factor. The results showed significant main effects of both severity, $F(1, 6) = 7.2, p < .04, \eta_p^2 = 0.55$, and category level, $F(2, 12) = 14.4, p < .001, \eta_p^2 = 0.71$, with Bonferroni corrected post hoc contrasts, indicating that the latter effect results from worse performance overall for the specific relative to the general condition ($p < .01$). Most interesting, however, was the significant interaction between severity and category level, $F(2, 12) = 8.4, p < .006, \eta_p^2 = 0.58$. Planned paired-samples t tests revealed that performance was significantly better in the general than in the basic condition ($p < .02$) for the severe cohort but not for the milder cohort ($p = ns$).

Discussion

Experiment 2 showed that the standard advantage for classifying visual stimuli at the basic level gives way to a general-level advantage as conceptual knowledge degrades. The pattern cannot be explained with reference to the overall difficulty of the various conditions because it is the typically more difficult condition (general classification) that is better preserved in the patients and the typically easier condition (basic-level classification) that deteriorates more rapidly in SD. Moreover, the two conditions that, from the healthy control data, seem matched for difficulty—general and specific conditions—show a dramatic difference in the more severely impaired group, with general-level performance near normal and performance at the specific level essentially at chance.

These data suggest, in concert with the bulk of the literature, that patients with SD retain fairly good knowledge of superordinate concepts even when entry-level concepts have degraded; this suggests in turn that retrieval of superordinate information does not depend upon prior activation of entry-level concepts. There are, however, at least two alternative explanations of the patient data consistent with the entry-level view. First, patients might simply be misclassifying the stimuli at the entry level, thus producing errors at the basic level but not the superordinate level. For instance, perhaps the patients have lost the entry level concepts “goat,” “pig,” and so on, but retain the concept “dog”—so that they incorrectly judge the former items to be dogs (thus producing errors in the basic-level condition) and then conclude that they are

animals (thus producing correct responses at the superordinate level). Second, it is possible that the entry level simply “moves up” the hierarchy as the disease progresses—so that superordinate category representations previously accessible only through a spreading-activation mechanism come to be directly activated by the visual stimulus (i.e., when the *dog* node degrades, the *dog* stimulus directly activates the *animal* node in the network).

Despite these caveats, the data are sufficiently compelling to warrant consideration of how alternative explanations of the basic-level advantage might explain the apparent preservation of general category knowledge in SD. In the next section, we describe an account that derives from the parallel distributed processing theory of conceptual knowledge (McClelland & Rogers, 2003; Rogers, Lambon Ralph, Garrard, et al., 2004; Rogers & McClelland, 2004). We see that, under this account, basic-level advantages arise from the same factors emphasized by the differentiation theory—informativeness and distinctiveness—but that the PDP theory makes counterintuitive predictions about the time course of processing in the category-verification task. These predictions are then tested in Part 3.

Part 2: A PDP Account of Basic-Level Effects

Several different PDP models of semantic memory have appeared in the literature (Cree, McRae, & McNorgan, 1999; Devlin, Gonnerman, Andersen, & Seidenberg, 1998; Farah & McClelland, 1991; Gots & Plaut, 2002; Hinton, 1981, 1989; Hinton & Shallice, 1991; Humphreys, Lamote, & Lloyd-Jones, 1995; Lambon Ralph, McClelland, Patterson, Galton, & Hodges, 2001; McRae, Sa, & Seidenberg, 1997; Plaut, 1999, 2002; Rumelhart & Todd, 1993; Schyns, 1991). The models differ in many of their details but derive from a similar theoretical approach to semantic cognition (see Rogers & McClelland, 2004). In this section, we provide a brief overview of the theory and consider how it addresses both the fine-to-coarse deterioration of semantic knowledge in SD as well as basic-level phenomena. In so doing, we refer to previous work described by Rogers, Lambon Ralph, Garrard, et al. (2004) and Rogers and McClelland (2004), but the principles that give rise to the phenomena of interest are general to PDP models that derive from the basic theoretical framework, and we believe the same effects would arise in many of the other model implementations cited above.

The PDP Theory

Our theory is premised on the notion that knowledge about the meanings of words and objects emerges from the interactive activation of perceptual, motor, and linguistic representations across different modalities of reception and expression. The word *rose*, when spoken, may call to mind an object with a particular shape, color, texture, smell, and so on, and these associations are necessary for apprehension of the word’s meaning. Correspondingly, when shown a black-and-white photograph of a rose, we can call to mind the appropriate name and other statements about it (“it’s pretty,” “it has thorns”); its characteristic color, smell, and so on; and when walking on a dark night, the smell of roses might provoke a similar constellation of associations across other modalities.

The idea that meanings emerge from such associations has a long and storied history in cognitive science that we do not consider further here, except to note that the basic idea has received considerable empirical support from functional neuroimaging and neuropsychological studies over the past two decades (Barsalou, Simmons, Barbey, & Wilson, 2003; Chao, Haxby, & Martin, 1999; H. Damasio et al., 1996; Kellenbach, Brett, & Patterson, 2001; Pulvermueller, Harle, & Hummel, 2001; Saffran, 2000; Thompson-Schill, 2003; Warrington & Shallice, 1984). Indeed, every theory we are aware of that addresses the neural underpinnings of semantic knowledge depends, to a greater or lesser extent, on this premise (though there is considerable debate as to whether conceptual knowledge inheres solely in such associations or whether other kinds of representations and processes are also required).¹ In common with these other views, we propose that different kinds of sensory-motor information are coded in separate, neuroanatomically distinct cortical regions and that semantic or conceptual knowledge inheres in the learned associations among the modality-specific representations coded in these different regions.

We further propose (consistent with the PDP approach to cognition) that, within each region, different representations are instantiated as different patterns of activity distributed across the same neurons. For instance, different shapes are coded with different patterns of activity across the neurons in the “shape” regions, different patterns of motion are coded with different patterns of activity across the “motion” region, and so on. These distributed representations capture modality-specific similarity structure—so that, for instance, two objects with similar shapes but different colors (e.g., a lightbulb and a pear) would be coded with similar patterns in the “shape” regions and different patterns in the “color” regions, whereas the reverse would be true for two objects with different shapes but a similar color (e.g., a strawberry and a fire truck).

Finally, in contrast to some theorists (e.g., A. R. Damasio, 1989) but consistent with others (e.g., Humphreys et al., 1995), we propose that the associations between the various different modality-specific representations are not predominantly mediated by separate pathways between the corresponding cortical regions but by a shared set of amodal semantic representations. We believe these representations to be coded in the anterior temporal cortex—the focus of the neuropathology observed in SD. Like the modality-specific surface representations, the semantic representations in our theory are instantiated as distributed patterns of activity across a set of neurons, with different patterns corresponding to different concepts. Whereas the surface representations capture modality-specific similarity structure, the semantic representations capture conceptual similarity structure. That is, items that are “the same kind of thing” will be represented as similar, even if they differ greatly in shape (e.g., a flamingo and a hummingbird), color (e.g., a lemon and a lime) and so on; whereas items that are “different kinds of things” will be represented with dissimilar patterns, even if they share a similar shape (e.g., a pear and a lightbulb), color (e.g., a strawberry and a fire truck), and so on.

On this view, the semantic representations act as a kind of “hub,” allowing the instantiation of a representation in one modality (e.g., shape) to provoke the instantiation of associated words, actions, and perceptual characteristics in other modalities.

The basic idea is schematized in Figure 3. Here, the appearance of a visual stimulus such as a line drawing first engages a representation of object shape, which then engages a semantic representation in the hub, which in turn engages representations of associated attributes in other sensory, motor, and language regions. These can then feed back to influence the representation coded in the hub.

This general idea—that perceptual similarity structure is coded in more peripheral modality-specific regions, whereas conceptual similarity structure is captured by more central semantic representations that serve to mediate cross-modal mappings—is common to most of the various PDP models of semantic memory cited previously. The hypothesis that the anterior temporal lobes function as a cross-modal “hub” is supported by neuroanatomical observations (Gloor, 1997), functional neuroimaging data (H. Damasio et al., 1996; Gauthier, Anderson, Tarr, Skudlarski, & Gore, 1997; Gorno-Tempini, Wenman, Price, Rudge, & Cipolotti, 2001; Nakamura et al., 2001; Rogers et al., 2006; Scott, Leff, & Wise, 2003), and computational arguments (Rogers, Lambon Ralph, & Garrard, et al., 2004; Rogers & McClelland, 2004), in addition to neuropsychological observations from SD (Adlam, Patterson, Rogers, Salmond, & Hodges, 2006; Bozeat et al., 2000; Bozeat, Lambon Ralph, Patterson, & Hodges, 2002; Patterson et al., 2006; Rogers, Lambon Ralph, Hodges, & Patterson, 2004).

In our theory, and in contrast to most other related views, the semantic representations do not themselves encode any explicit or directly interpretable content—they do not consist of propositions, mental images, semantic features, amodal lexical symbols, or the like. The content of semantic memory is grounded in the perceptual, motor, and language representations that negotiate our interactions with the environment, as suggested by Barsalou et al. (2003) and others (Glenberg & Robertson, 2000; Pulvermueller et al., 2001). Representations within the hub subservise a particular functional role: They mediate the interactions amongst these perceptual, motor, and linguistic representations—so that, for instance, when a word form is activated via perceptual mechanisms, this promotes a pattern of activation in the hub, which then feeds back to modality-specific areas to instantiate representations of the sensory/motor properties of the item denoted by the word. We refer to these intermediating representations as *semantic* because this functional role is critical to semantic capabilities. Because semantically related items are represented with similar patterns of activity in the hub, these representations provide the basis for semantic induction and generalization. For instance, if the system learns that a pear has seeds inside, this new information will tend to generalize to other kinds of fruit because they are represented with similar patterns of activity in the hub; but it will not generalize to the lightbulb which, despite its similar shape, is represented with a very different pattern in the hub. Semantic similarity relations are not captured by any individual modality in the semantic network; hence, without the common intermediating representations at the hub—if, for instance, representations in different modalities projected directly to one another—there would be no basis for semantic generalization and induction (see Rogers &

¹ Rogers and McClelland (2004) discussed in considerable detail how many of the criticisms of so-called associationist theories can be addressed by the theoretical framework defined in this article.

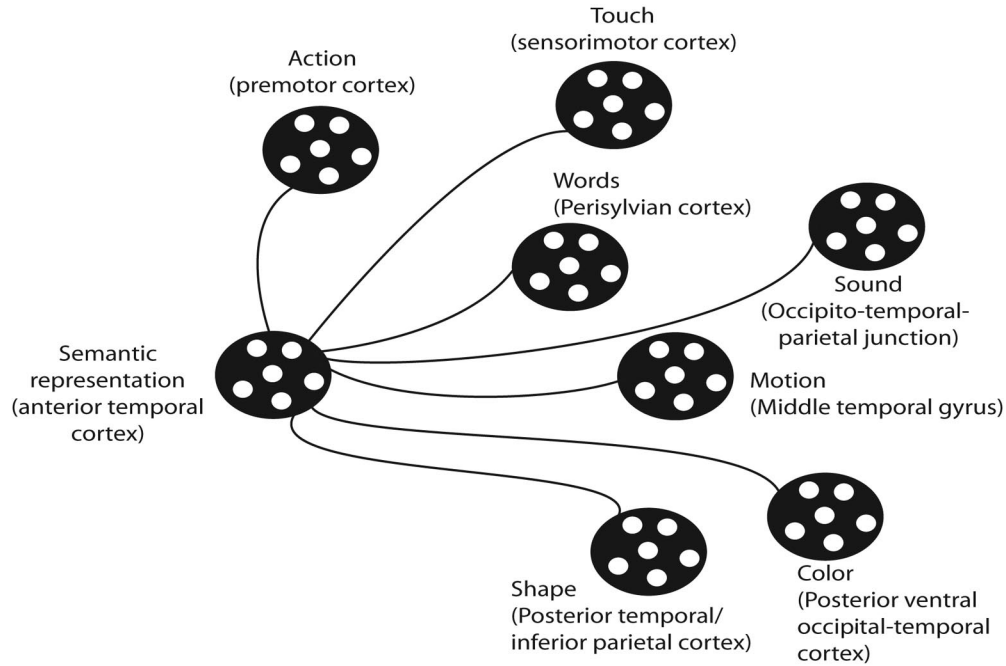


Figure 3. Schematic of the theoretical framework. Each black circle represents a separate region of cortex, and the lines indicate how different regions are connected in the cortical semantic network under the theory. Each region outside the “hub” encodes a different kind of perceptual, motor, or linguistic representation. Within each region, representations are instantiated as distributed patterns of activity across neuron-like processing units. The various “surface” (i.e., sensory, motor, and language) representations communicate with one another through a common set of distributed representations coded in the anterior temporal cortex. Functional and neuroanatomical labeling of these regions, though meant to be suggestive, is based on that proposed by McClelland and Rogers (2003) on the basis of functional imaging and neuropsychological studies.

McClelland, 2004, chapter 9 for discussion of this point). In this sense, our theory is also similar in spirit to other proposals emphasizing that the main function of semantic memory is to support inferences about the unobserved properties of objects and events (or objects and events referred to in speech) from partial information (Anderson, 1991; Corter & Gluck, 1992).

If the intermediating semantic representations are the only ones that capture conceptual similarity structure, where does such structure come from? Our theory suggests that the structure emerges from a domain-general learning process, as the hub gradually learns to complete appropriate mappings between surface representations in different modalities. Specifically, the similarity structure in the hub comes to mirror the similarities apparent across all of the different sensory, motor, and language representations to which it is connected—so that items that have many different kinds of properties in common (similar shapes, movements, functions, colors, and so on) come to be represented as similar (Rogers & McClelland, 2004; Rumelhart, Hinton, & Williams, 1986; Rumelhart & Todd, 1993). These cross-modal similarities can differ substantially from the similarities captured by any single modality alone—for instance, although the pear and the lightbulb have similar shapes, they differ greatly in their color, their texture, the uses to which they are put, the way they are described in language, the contexts in which they are typically found, and so on. Bananas and pears, though they have different shapes and colors, are similar in these other respects. The representations that inter-

mediate mappings across all modalities come to represent the pear and banana as similar to one another, and as different from the lightbulb, because they are tuned by the structure apparent across all the different surface representations to which they are connected (see Rogers, Lambon Ralph, & Garrard, et al., 2004).

The theory thus contends, as did Rosch (Rosch, 1978; Rosch et al., 1976; Rosch, Simpson, & Miller, 1976), that the conceptual similarity relations that govern generalization and induction are shaped by the perceived structure of the world across all different modalities of expression and reception. We further assume, as did Rosch (Rosch, 1978; Rosch et al., 1976), and consistent with the differentiation theory of basic-level effects (Murphy & Brownell, 1985), that exemplars of a basic-level category such as “dog” tend to share many attributes across modalities and share considerably fewer attributes with exemplars of other basic-level categories (e.g., “bird”)—so that, for instance, different individual dogs are all represented with quite similar patterns of activation in the hub, whereas other kinds of animals (pigs, goats, birds and so on) are represented with somewhat different patterns, and nonanimals are represented with dramatically different patterns. In other words, basic-level categories correspond to relatively tight and widely separated—that is, distinct and informative—clusters of distributed representations in the network hub. Subordinate concepts then correspond to smaller and less well-separated clusters within the basic-level cluster and so have many near neighbors from different subordinate groups—they are informative but not distinctive. Su-

perordinate concepts correspond to more inclusive but sparser clusters in the representation space—they are distinctive but not as informative. This notion is schematized in Figure 4.

A PDP Account of Basic-Level Effects

In this framework, similar to the differentiation theory, basic-level effects arise as a direct consequence of the similarity structure of the representations coded in the hub. The mechanism by which this structure results in a basic-level advantage differs somewhat from that proposed by the differentiation theory: It depends upon the pattern of generalization fostered by these representations as the network learns to name (Rogers & McClelland, 2004). To see this, consider how a network like the one in Figure 3, with semantic representations like those schematized in Figure 4, might behave when learning to name various animals at subordinate, basic, and general levels. When taught that a certain bird is a *canary*, this response will tend to generalize to items with similar representations, namely other birds. When the network later learns that another particular bird is called *robin*, and not *canary*, the earlier learning must be reversed to some extent—the weights must change to lower the activation of *canary* and raise the

activation of *robin*, and this change in turn will tend to generalize to all the birds. That is, when items with similar representations have different names (as in subordinate naming), similarity-based generalization produces interference that slows name acquisition. In contrast, when learning to call the canary a *bird*, the response generalizes strongly to all the other birds, speeding learning of the basic-level term for all of these; and because nonbirds are all represented as distinct from the canary, the response will not tend to generalize to these other items so that little interference is produced. Thus, the similarity structure of the representations produces little interference when learning basic-level names and much more interference when learning more specific names.

Next, consider learning to call the canary an *animal*. In this case, the response will again tend to generalize strongly to the other birds, which is appropriate because all other birds are also animals. The response will not generalize very strongly, however, to the other animals because they are represented as quite distinct from the canary. So, for every time the word *animal* is encountered during learning, it will benefit only a small fraction of the items to which the name applies (namely, all of the items from the same basic category). In contrast, every time the basic-level term *bird* is

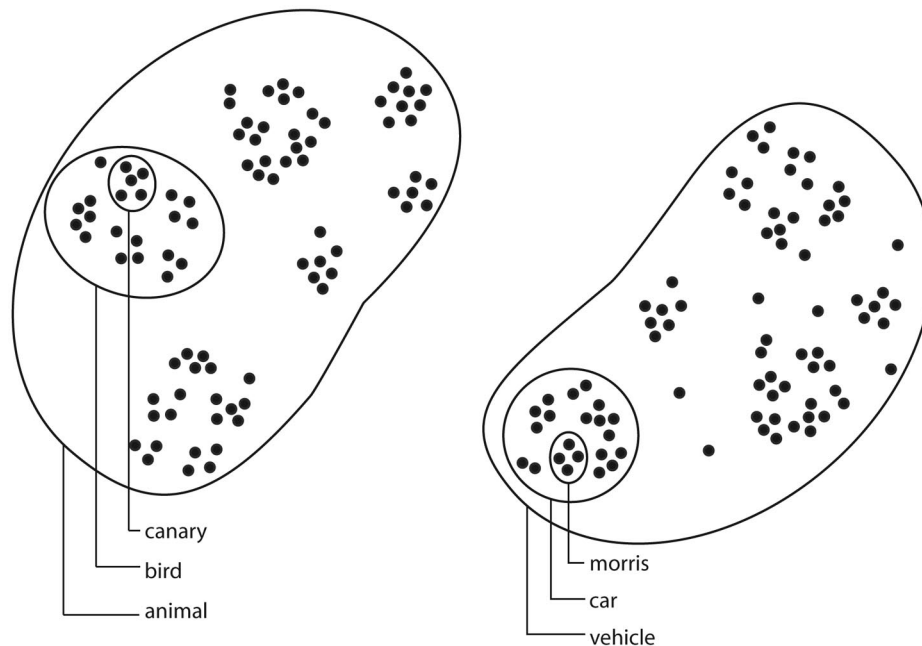


Figure 4. Schematic representation of the similarities among distributed semantic representations for various animals and vehicles postulated by the theory. In the parallel distributed processing theory, semantic representations are instantiated as patterns of activity across neuron-like processing units in the anterior temporal-lobe “hub”. Any particular representation can thus be construed as a point in a high-dimensional space. In the figure, each point indicates the representation of a particular item (e.g., a particular individual canary), and the proximity between points indicates the degree of similarity between the corresponding representations. Various individual canaries are very similar to one another; birds in general are quite similar to one another and somewhat segregated from other animals; but animals are, overall, somewhat more similar to one another than they are to vehicles. The contours and labels indicate how names at different levels of specificity map onto this representation space. Specific names apply to items with similar representations (i.e., all the canaries) but do not apply to close neighbors (i.e., other birds). General names apply across a broader swath of the representation space and so encompass items with quite different representations. Basic-level terms map most closely onto clusters in the representation space.

used for an individual item, it generalizes strongly to most or all of the other items for which it is an appropriate response because these are all represented as similar. As a result, if the terms *bird* and *animal* appear equally often in the environment, the name *bird* will be learned more quickly. Similarity-based generalization thus promotes more rapid learning of basic-level names for words of equal frequency.

Of course, these influences will vary depending upon the exact similarity structure of the semantic representations and the frequency of the labels at different levels of specificity. The general principles are as follows: (a) the more frequent the label, the more quickly and strongly it will become activated, all else being equal; and (b) these frequency effects interact with the similarity structure of the semantic representations, so that (again, all else being equal) names are more slowly acquired and more difficult to activate when they apply across sets of items with very different representations, or there are items with different names that have very similar representations. Exemplars of basic-level categories are represented as similar to one another and as distinct from other items, and so basic-level names get the most benefit and the least interference from similarity-based generalization.

Accounting for the Preservation of Superordinate Information in SD

If basic-level terms are, for the reasons described above, easier to learn and faster to activate, why should more superordinate information be more robust to brain damage? The answer suggested by our theory (see Rogers, Lambon Ralph, & Garrard, et al., 2004) becomes apparent by again considering the similarity structure of the semantic representations shown in Figure 4. Recall that each point in this illustration corresponds to a pattern of activity across neurons in the anterior temporal cortex, with proximal points corresponding to similar patterns of activity. When the anterior temporal cortex deteriorates in SD, the pattern that arises there in response to some stimulus is necessarily distorted because the disease has destroyed some of the neurons that coded the healthy pattern. Small distortions to an item's representation will disturb the ability to activate, elsewhere in the semantic network, very specific properties, whereas the ability to activate more general properties remains somewhat more intact.

To see this, consider what must happen in the healthy network when it names birds at the subordinate level. To correctly activate the name *canary* in verbal-output parts of the network, the hub must instantiate the *canary* pattern almost exactly. Because the name does not apply to all of the other birds, and because these all have quite similar representations, then very small deviations from the *canary* pattern will lead the network to stop producing the name *canary* as output. In contrast, a name like *bird* applies to the canary and to all the other birds—so that the network need not get the pattern for *canary* exactly right in order to produce this name. If, as a consequence of disease, the pattern for *canary* gets distorted so that it becomes more like the pattern for *crow* or *robin*, the network will still be able to produce the name *bird*, because this response can be generated from any bird-like pattern. So, basic-level naming can tolerate small distortions to the intermediating representations; and by the same reasoning, naming at still more general levels can tolerate still greater distortions to the representations. For instance, the name *animal* applies not just to

the birds but to the various land animals and fish as well—so even if the representation of canary is distorted so much that it falls closer to the representation of dog than to the various bird representations, the network will still be able to correctly produce the response *animal*.

In summary, under the PDP theory, basic-level advantages arise because names tend to generalize across items with similar representations, so that names that apply consistently across a set of tightly clustered and well-separated representations are easiest to learn; but the robustness of naming at different levels of specificity also depends upon the breadth of the representations spanned by the name once it has been acquired. In a feed-forward computer model based on the PDP theory, Rogers and McClelland (2004) showed that these influences combine to produce a basic-level advantage for naming and category verification but that naming at the basic level declines more rapidly as internal representations are increasingly distorted so that the basic- over general-level advantage is eventually reversed, as we have shown it to be in SD. In a fully recurrent model similar to that shown in Figure 3, Rogers, Lambon Ralph, & Garrard, et al. (2004) showed that such effects can explain a broad range of observations in SD.

Part 3: A Counterintuitive Prediction

The PDP theory makes a prediction about the time course of activation for names at different levels of specificity that directly contradicts the entry-level theory and that initially seems counterintuitive from the perspective of the differentiation theory. The prediction of the entry-level theory is straightforward. The theory states that people are fastest to name and verify category membership at the basic level (for typical items) because basic category representations are the first to be activated by a visual stimulus—so participants should always be most accurate to verify category membership at the basic level, no matter how quickly they are required to respond. Similarly, the aim of the differentiation theory is to explain why people are often fastest to categorize at the basic level, and so there is no reason to expect, under this theory, that the basic-level advantage should be qualitatively altered under speeded-response conditions. In contrast, the PDP theory predicts that, if participants are pressured to respond more quickly than they ordinarily would, the basic-level advantage for category verification should give way to a basic-level disadvantage relative to more superordinate categorization.

To see this, consider what happens in a model like that shown in Figure 3 when it is presented with a visual stimulus such as a picture of a canary. Prior to the onset of the stimulus, the hub representation is presumably in some neutral state. When the stimulus appears, visual input begins to change this state, moving it toward the correct representation (i.e., the pattern corresponding to *canary*). The key observation is that, because the *canary* representation is nested within the broader region of space to which the name *bird* applies, the network will begin to pass through the “bird” region before it reaches the *canary* representation; and as it enters the “bird” region, it will begin to activate the name *bird* in the output. Put differently, before the model's internal representation has established the particular pattern that represents the canary (from which it can activate all of the canary's properties in other parts of the network), it passes through a state that is generally similar to all of the individual birds that it knows about. Because

the name *bird* is true of all these items, the network will begin to activate the “bird” label from this incompletely-specified representation, before it has tuned the representation to the precise pattern corresponding to *canary*. The same argument applies to the name *animal*: Because the region to which the name *bird* applies is completely encompassed by the broader region to which *animal* applies, the model’s internal state will begin to pass through the “animal” region first. This in turn means that, counterintuitively, the name *animal* ought to begin to activate before the name *bird*.

Why then is there a basic-level advantage for naming or category verification at all? The reason is that, although *animal* begins to activate sooner than *bird*, the name *bird* gets fully activated much more quickly, due to the effects of similarity-based generalization discussed previously. Put differently, superordinate names begin to activate first because they apply broadly across a wide range of semantically related items; but their time course of full activation (where *full* means sufficient to provide a confident response) will be slower for the same reason—the broadly spaced representations to which they apply do not foster much generalization one to another so the name is slower to activate. Because basic names apply across a narrower region of the space, their activation starts later but accelerates more rapidly.

This idea is schematized in Figure 5. The left panel shows the movement of the hub’s internal state, beginning at a neutral point and then moving successively through regions to which the various names apply—first “animal,” then “bird,” and finally “canary.” The right panel shows the corresponding predicted time course of activation for these names. Assuming that the time to verify category membership ordinarily (i.e., without deadline pressure) depends upon one of the name outputs exceeding a threshold, then if the threshold for responding is set relatively high (as shown by the dashed horizontal line in the Figure 5), basic-level responses will be the first to meet it, producing a basic-level advantage for naming and category verification. What happens when the system is forced to respond at increasingly shorter latencies before any name has exceeded threshold? Under the assumption that the system produces a “best guess” proportional to the activation of the name unit, the basic- over general-level advantage should first vanish and then be reversed—so that, at very fast latencies, people should be more accurate for superordinate relative to basic categorization.

There are some hints in the literature to support this prediction. For instance, Van Rullen and Thorpe (2001) elicited remarkably fast (200–400 ms) and accurate responses in a task requiring




Figure 5. Schematic representation showing the basis for the parallel distributed processing theory’s prediction about the time course of activation for names at different levels of specificity. When a visual stimulus appears, the model’s semantic representation state begins to move from some neutral point toward the appropriate representation (e.g., a particular canary), as illustrated in the left panel. The right panel shows the predicted activation of different names over this time span. To reach the end state, the system’s internal representation first begins to pass through the region of the space to which the general name applies so that the general name begins to activate soonest; but because general-name learning does not benefit greatly from similarity-based generalization, the label is slow to activate. Some time later, the representation reaches the area in which the basic name applies, and the corresponding name begins to activate. Because basic-level clusters promote similarity-based generalization and minimize interference, the name activates more rapidly. The specific name does not begin to activate until the internal state is very close to the correct representation, so this is the last name to activate.

participants to detect whether a visually complex scene contained a means of transport (in one condition) or an animal (in another). The stimuli were presented for only 20 ms (though without backward masking), and participants showed better-than-chance responding as early as 200–250 ms—suggesting that superordinate-level information is available extremely rapidly in this kind of detection paradigm. In a similar vein, Large, Kiss, and McMullen (2004) used an oddball-detection paradigm to investigate the speed of visual classification at different levels of specificity. Participants monitored a stream of visual images for a particular target, with targets occurring rarely relative to distractors. “Oddball” targets have been associated with a negative event-related potential (ERP) deflection at an early poststimulus latency; hence, the authors were able to investigate the differentiation of targets from distractors in both the ERP signal and in behavioral data. The interesting finding was that correct responses to targets were faster in conditions requiring a superordinate discrimination (e.g., detecting animals, including 30 beagles and 20 other animals, from among a stream of manmade objects) than in conditions requiring a basic-level discrimination (e.g., detecting dogs, including the same 30 beagles and 20 other dogs from among a stream of animals). Consistent with this behavioral finding, analysis of the ERP showed an early peak discriminating superordinate responses from basic and subordinate responses, with subordinate responses differentiating from basic responses later in the wave form.

Such findings are suggestive because they appear to indicate that information sufficient to support categorization at the superordinate level is available earlier in processing than is the information required to differentiate items at the basic level, as predicted by the PDP theory. They are not conclusive, however, because the detection tasks that elicit these effects differ substantially from the tasks that reliably elicit basic-level advantages (e.g., naming, category verification, property verification, and so on). It is therefore not clear whether the surprising pattern of results arises because basic-level information is not, contrary to the received view, available earliest in processing, even in tasks that reliably elicit basic-level advantages, or whether the rapid visual detection tasks used in this research exert sufficiently different demands on visual, semantic, and executive processes that the usual basic-level advantage does not obtain.

To the contrary, the PDP theory predicts that, under speeded response conditions, a superordinate-over-basic advantage ought to be observed in the very same task and with the very same stimuli that, under nonspeeded conditions, elicit a basic-level advantage. This is the prediction tested in Experiment 3. The same participants from Experiment 1 completed the same category-verification task, with the same materials. The only difference was that the participants were required to make their responses in time with a deadline so that we could tightly control the amount of time they had to process the visual stimulus. We then investigated the impact of different response deadlines on accuracy and speed of superordinate-, basic-, and specific-level categorization.

Method

Participants and materials. The same participants from Experiment 1 participated in this experiment in the same session. All participants completed Experiment 3 after Experiment 1. The same stimuli from Experiment 1 were used in Experiment 3.

Procedure. On each trial, as before, participants first viewed a subordinate, basic, or general category name, followed by a matching or nonmatching color photograph and were asked to decide whether the photograph matched the name. In contrast to the previous experiment, participants were required to time their responses very precisely. To impose the response deadline, we used a tempo-matching procedure adapted from Kello (2004). After the presentation of the category name, participants heard four regularly-spaced beeps, and they were instructed to time their response as closely as possible with the fourth beep (i.e., *beep, beep, beep, response*) so that the fourth beep corresponds to a response deadline. In contrast to the Kello (2004) paradigm, the timing of the beeps was identical in every trial in our experiment (1 s from onset to onset), so participants were able to get the “rhythm” of responding. We then varied the onset of the picture with respect to the deadline. The experimental design is illustrated in Figure 6. Because the timing of the beeps was always the same, participants had no cue as to what the latency would be between the appearance of the picture and the response deadline from trial to trial.

All stimuli and conditions from Experiment 1 were repeated three times at three different deadlines—slow, medium, and fast—so that response deadline is a within-subjects factor. The slowest deadline was set separately for each participant to match his or her mean RT in the basic-level condition of Experiment 1. That is, if Participant A had a mean RT of 970 ms for correct basic-level trials in Experiment 1, then in the slow deadline condition for that participant, the picture always appeared 970 ms before the deadline. The medium deadline was set separately for each participant to be 200 ms shorter than the slow condition, and the fast deadline was set to be 200 ms shorter than the medium condition. So, the participant with a slow (i.e., “normal”) deadline of 970 ms would have 770 ms to process the picture in medium trials and 570 ms in fast trials. Trials were not blocked and were presented in a different random order for every participant so that the deadline was unpredictable from trial to trial. Participants received positive feedback if they responded within 100 ms of the deadline and negative feedback otherwise.

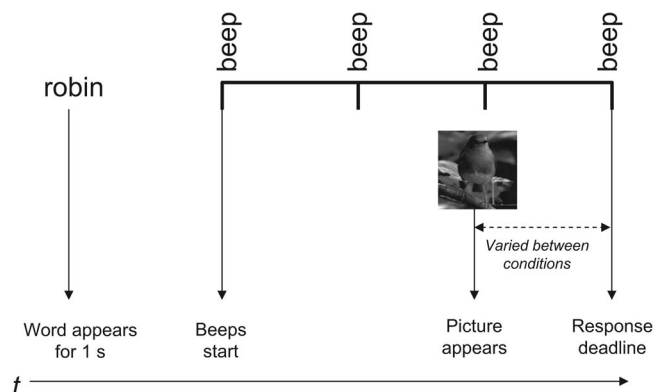


Figure 6. Schematic of the trial design. At the beginning of each trial the word appears, followed by four regularly spaced beeps. The response deadline is timed with the fourth beep. The onset of the picture with respect to the last beep is varied across slow, medium, and fast trials; t = time.

Results

Analysis 1. Responses were considered to have met the response criteria if they fell within 100 ms of the deadline because 100 ms marks the halfway point between deadlines in the slow, medium, and fast conditions. Note that, although this cut-off point seems fairly strict, the tempo-matching procedure allows participants to anticipate the upcoming deadline—so that it is considerably easier to meet the criterion when using this method compared with other methods (such as when using a tone to signal the deadline). In the first analysis, responses that did not meet the deadline were discarded (an average of 21% of trials across participants). Accuracy was then tabulated across the remaining trials for decisions at each deadline for each specificity condition (specific, basic, and general). As previously noted, we used d' as a measure of accuracy for discriminating between targets and distractors.

The top left panel in Figure 7 shows the mean RT for correct trials in specific, basic, and general conditions, with no response deadline (i.e., Experiment 1) and for the three deadline conditions

in Experiment 3. Participants clearly responded much more quickly for all conditions in Experiment 3 than for those in Experiment 1, a finding that may partially reflect familiarity with the stimuli in addition to the imposition of the response deadline. It is important to note, though, that the tempo-matching procedure worked: Participants had the slowest latencies in the slow condition, somewhat faster latencies in the medium condition, and the fastest latencies in the fast condition. A repeated-measures ANOVA predicting RT from response deadline (fast, medium, slow) and category level (general, basic, specific) showed a main effect of deadline, $F(2, 50) = 13.1, p < .001, \eta_p^2 = 0.83$, with Bonferroni corrected post hoc contrasts indicating that participants were slower in the slow condition compared with the medium condition ($p < .001$) and slower in the medium condition compared with the fast condition ($p < .001$). There was also a main effect of category level, $F(2, 50) = 8.5, p < .001, \eta_p^2 = 0.25$; Bonferroni corrected post hoc contrasts showed that this effect was driven by faster RTs to specific relative to general categorization ($p < .004$), but that there was no significant difference in RT




Figure 7. Top row: Reaction times (left) and discrimination accuracy (right) for categorization at different levels of specificity in Experiment 1 and in the three deadline conditions of Experiment 3, considering all trials that met deadline. Bottom row: The same data, discarding from all conditions any stimulus that failed to meet deadline in any condition. Exp1 = no deadline data from Experiment 1.

between basic classification and either specific or general classification. Most important, there was no significant interaction between deadline and category level, $F(4, 100) = 0.25, p = ns$ —indicating that the RTs at different levels of categorization, relative to one another, did not change across different deadlines.

The top right panel of Figure 7 shows discrimination accuracy (d') when there was no response deadline (Experiment 1) and at the three response deadlines in Experiment 3. For basic and specific conditions, accuracy declined as participants were offered less time to process the stimulus, but accuracy for general or superordinate classification was relatively constant—so that, at the fastest deadline, performance was better for the general than for the basic condition. A repeated-measures ANOVA predicting accuracy from deadline condition (slow, medium, fast, and the no-deadline condition from Experiment 1) showed (a) a main effect of deadline, with accuracy worse overall for faster than slower deadlines, $F(3, 78) = 13.1, p < .001, \eta_p^2 = 0.34$; but (b) no main effect of specificity, $F(2, 52) = 1.8, p = ns$; and (c) an interaction between deadline and specificity, $F(6, 156) = 5.3, p < .001, \eta_p^2 = 0.17$. A planned-comparisons paired-samples t test revealed that, at the fastest deadline, accuracy was significantly greater for general than for basic categorization ($p < .001$).

Note that, although the overall faster RTs in the tempo-matching task may partially reflect repetition priming or other influences of stimulus familiarity, such effects cannot explain the reversal of the basic-level over general-level advantage at the shortest response deadline. Because the deadline conditions (short, medium, or long) were ordered randomly, there is no reason to expect that items in the short condition were any more or less primed, on average, than items in the medium and long conditions. Furthermore, it is not clear how or why repetition priming would benefit general categorization any more than basic or specific categorization.

Analysis 2. Though the results of the first analysis match the predictions of the PDP theory, they may not be conclusive because of the decision to discard responses that did not meet the deadline in each condition. In the stimuli we used, it is possible that some items produce a reliable basic-level advantage, whereas others do not, and that RTs for such items rarely met the fastest deadline, so that most such trials were discarded in the fast condition, but not in the other conditions. In this case, the apparent superiority of general over basic-level classification might arise simply because many of the stimuli that carried the basic-level advantage in Experiment 1 were thrown out. To investigate this possibility, we conducted a second analysis of the data. For each participant, we discarded from all deadline conditions any stimulus item that failed to meet the deadline in any condition. For instance, if Participant A missed the deadline for the robin in the fast condition, the robin was removed from her data for all four conditions (i.e., the three deadline conditions of Experiment 3 and the no-deadline data from Experiment 1). No other stimuli were discarded. This ensured that, for each participant, exactly the same stimuli entered the analysis in all deadline conditions—although there were considerably fewer data points considered for each participant. Specifically, 53% of trials were discarded for each participant. We then investigated accuracy exactly as in Analysis 1.

The bottom two panels of Figure 7 show that Analysis 2 yielded the same pattern as Analysis 1: A basic-level advantage in the no-deadline conditions from Experiment 1 gave way to a general-over-basic advantage at the fastest deadline. A repeated-measures

ANOVA again showed a significant main effect of deadline indicating that on average accuracy declined with shorter deadlines, $F(3, 75) = 3.7, p < .02, \eta_p^2 = 0.13$. There was also a main effect of category level in this analysis, $F(2, 50) = 3.3, p < .05, \eta_p^2 = 0.12$; Bonferroni corrected post hoc contrasts suggested that this effect was driven by somewhat better performance in general relative to specific categorization across deadline conditions ($p < .08$, corrected). Most important, the interaction between deadline and category level was significant, $F(6, 150) = 3.5, p < .004, \eta_p^2 = 0.12$. Planned-comparisons t tests showed that, in the free-response data from Experiment 1, the basic-level over general-level advantage remained significant for these trials ($p < .008$) but that the effect reversed to give a significant advantage to general over basic-level classification at the fastest deadline ($p < .02$). Because exactly the same items appear in all deadline conditions for each participant in this analysis, this reversal cannot reflect the removal of items that “carry” the basic-level effect. The effect is particularly striking considering the reduction in power resulting from the pruning of half of the trials.

General Discussion

We began by noting a seeming contradiction apparent when studies of conceptual knowledge in healthy and neuropsychological populations are considered together: Whereas healthy individuals are faster and more accurate to categorize most stimuli at the basic level, patients with semantic dementia show better categorization at more general levels. Experiments 1 and 2 illustrated the contradiction: The same stimuli that produce a basic-level advantage for category verification in healthy older adults and patients at the mildest stage of SD give rise to a substantial superordinate-over-basic advantage in more moderate or severe SD patients. These findings seem especially puzzling from the perspective of “entry-level” theories of basic-level effects, which propose that basic-level category representations constitute the entry point into a semantic processing hierarchy so that access to more general information depends upon prior activation of more specific representations.

To resolve the apparent contradiction, we suggest an alternative explanation of basic-level effects, which builds upon the differentiation theory advocated by Murphy (Murphy, 2002; Murphy & Brownell, 1985). Like the differentiation theory, our theory suggests that basic-level effects arise, not because basic-level categories are privileged in processing, but because the exemplars of such categories are quite similar to one another and quite distinct from the exemplars of other categories. The particular mechanism by which such structural properties produce basic-level effects are, however, somewhat different in the two theories.

Specifically, the PDP theory suggests that basic-level effects arise as a consequence of similarity-based generalization in a semantic system, wherein amodal conceptual representations serve to mediate learned associations among various perceptual, motor, and linguistic representations. Instances of basic-level categories are represented as similar to one another and distinct from other items in this cross-modal “hub,” precisely because they share many properties with one another and are somewhat distinct from other items. Basic-level effects arise as a direct consequence of this structure: Because items from the same basic category are represented as similar to one another and distinct from other items,

learning generalizes strongly within such categories but only weakly between them. This pattern of generalization promotes rapid and strong learning of basic-level names—when learning that a canary is a kind of bird, for instance, the name *bird* will tend to generalize strongly to items with similar representations (i.e., other birds) but not to items with distal representations (e.g., other kinds of animals). More specific names, in contrast, suffer from similarity-based interference, whereas more general names get less benefit from similarity-based generalization.

Rogers and McClelland (2004) demonstrated that these influences can give rise to basic-level advantages in naming and category verification in a simple model implementation of the theory. Though not the focus of this article, it is perhaps worth noting that the authors also showed how the PDP theory can account for a range of other basic-level phenomena, including the primacy of the basic level in lexical acquisition (Mervis, 1987a), “child-basic” category phenomena (Mervis, 1987b), the influence of typicality on basic-level effects (Jolicoeur et al., 1984; Murphy & Brownell, 1985), and how the basic-level advantage changes with expertise (Johnson & Mervis, 1997; Tanaka & Taylor, 1991). As in the current work, the accounts of these phenomena offered by the PDP theory are similar in spirit to those offered by the differentiation theory, as instantiated within a PDP framework.

The couching of a differentiation-like theory within a PDP framework allows us to see why facilitation of basic-level naming in a healthy system is no guarantee of robustness in the face of semantic impairment. Because basic-level names apply across a comparatively narrow span of concepts, whereas more general names apply across a much broader span, relatively small distortions of semantic representations will compromise basic naming or categorization but will have less impact on these same tasks at a more general level.

This understanding of impaired categorization in SD supports a prediction about speeded categorization in healthy individuals. Individuals responding under strict time pressure are somewhat like patients with SD, in that they must generate a decision about category membership from a poorly specified internal semantic representation that does not correspond to the learned representation of a specific known item. In the case of SD, the correct internal representation is poorly specified because the cortical region that codes the representation has degraded with disease—effectively distorting the intermediating representations. In the case of speeded responding, the correct representation is poorly specified because the system has not had sufficient time to settle to the correct state. In both cases, some information can still be recovered—namely, properties held in common by a broad range of semantically related items. Because such properties are activated, at least weakly, whenever the system finds itself “in the right ballpark,” then they can be recovered even though the system is unable to find, or has not yet found, the right learned representation. Thus, the PDP theory predicted that individuals pressed to respond very rapidly should show a pattern of performance similar to SD—more accurate responding for general-level categorization than for basic-level categorization—in the very same task and for the very same stimuli that elicit a basic-level advantage in non-speeded responding. This prediction, supported by the results of Experiment 3, directly contradicts the entry-level theory of basic-level effects, and it illustrates how a differentiation-like explana-

tion of basic-level effects is compatible with the fine-to-coarse dissolution of conceptual knowledge observed in SD.

Implications for Spreading-Activation Theories of Semantic Cognition

In some respects, basic-level phenomena and spreading-activation models of semantic knowledge make strange bedfellows. Spreading-activation theories seem intuitive when they are applied to purely propositional knowledge—that is, when the nodes in the network are understood as corresponding to individual words, and the links to individual predicates, so that the entire system of knowledge may be accurately characterized as a system of propositions. Under such a scheme, there are few questions about which concepts—which nodes and links—should inhabit the network. Very simply, each node and link correspond to a word in the language so that the contents of the network are determined by the lexicon, and the structure of the network represents beliefs that can be explicitly stated by propositions (e.g., “All birds have feathers”). And, such a representational scheme seems most plausible when considering experiments of the kind conducted by Collins and Quillian (1969), where participants must make judgments about the truth of written propositions.

The simple appeal of this approach gets complicated when the stimuli to be comprehended are perceptual representations of objects rather than words because it is less clear which node in the network should be activated by a given stimulus. A particular robin belongs equally to the classes “robin,” “bird,” “animal,” and “living thing,” so which of these nodes should a visual depiction of the robin activate? Any answer to this question involves postulating some additional mechanism, outside of the processing hierarchy itself, that constrains the appropriateness of assigning a stimulus to the various nodes in the network and indeed specifies how the assignment itself is carried out. A common assumption is that activation of an initial node is accomplished by a perceptual categorization mechanism, separate from the semantic network itself and providing the means for visual stimuli to access the propositional information stored there. Basic-level or entry-level phenomena are thought to provide evidence supporting this dual-mechanism perspective: By this account, rapid retrieval of basic or entry-level information reflects the operation of the initial categorization mechanism, whereas the slower retrieval of information at other levels of specificity (especially more general levels) reflects the operation of the semantic retrieval process (Jolicoeur et al., 1984).

This conception of the relationship between perceptual and semantic processes has shaped theories of both visual object recognition and semantics. In many theories, the “goal” of the recognition process is understood to be the activation of an entry-level category representation (Biederman & Gerhardstein, 1993; DeCaro & Reeves, 2002; Humphreys et al., 1995; Kurbat, 1994; Murray, 1998). Little consideration is given to the retrieval of more general category information, presumably because this is assumed to be “semantic” and therefore retrieved through independent mechanisms once the stimulus has been categorized (see, e.g., the collection of articles in Tarr & Bulthoff, 1998).

The current results challenge the view that the end product of the visual recognition system is the activation of an entry-level category representation: They show that classification at more

general levels is achieved earlier in processing in healthy individuals. At the very least, this finding reopens the question of which category representations are activated first by visual recognition processes. More substantively, we believe that the question itself may be misguided, arising as it does from the notion that mechanisms of perceptual recognition/categorization and semantic processing are separate and independent. The PDP theory is appealing in part because it does away with the discrete category representations and the propositional links characteristic of spreading-activation theories and, in so doing, renders moot the question of which category nodes are directly activated by perception. Instead, perceptual and semantic processes are viewed as continually interacting: perceptual processes dynamically construct distributed patterns of activity across modality-specific regions of cortex; these then interact with associated patterns in other modalities, via distributed representations in the anterior temporal-lobe hub. The time course with which information at different levels of specificity becomes available depends on (a) the similarity structure of the distributed patterns in the hub, (b) the consistency with which the properties to be retrieved are common to items with similar representations, and (c) the frequency with which the item-property correspondence has been observed in previous learning episodes. Basic-level effects, then, do not indicate the operation of two separate mechanisms, but arise from the influence of these factors operating in a more homogeneous and interactive system.

Relationship to Differentiation Theory

Throughout this article, we have noted similarities between the PDP theory and the differentiation theory of basic-level effects. The two accounts are similar in many critical respects: both eschew the notion that basic-level categories are “privileged” in processing; both suggest that semantic information at all levels of specificity is probed, in some sense, in parallel; and both place the root cause of basic-level advantages in the attribute structure of the environment. We therefore view our theory as capturing the spirit and aims of the differentiation theory.

With that said, we believe that the PDP theory extends differentiation theory in two important respects, both of which derive from differences in the putative nature of semantic representations under the two theories. Under the differentiation theory as we understand it, the semantic system is construed as containing a vast set of discrete prototype-like category representations, with separate representations for categories at different levels of specificity. In contrast to spreading-activation theories, these representations are not organized within a processing hierarchy but may be probed in parallel by a given stimulus. The rapidity with which a given representation is activated depends partly on the extent of featural overlap with the probe and partly upon competition amongst the representations themselves. It is these dynamics that promote the rapid activation of basic-level representations, as discussed in the introduction.

Feature-based prototype theories, however, raise difficult questions about the hierarchical organization of concepts. If prototypes at different levels of specificity are assumed not to be linked to one another in some kind of processing hierarchy—for instance, if representations of *robin*, *bird*, and *animal* constitute separate and unrelated prototypes—then it is difficult to understand phenomena that seem to require hierarchical organization. For instance, to

determine whether a particular fruit is a lime or a lemon, one must take account of its color; but color is irrelevant when deciding whether a certain vehicle is a car or a boat. Some form of hierarchical organization seems essential to explain such domain-specific attribute weighting: One must take into account whether a given item is a fruit or a vehicle before one can determine which properties are relevant to categorizing it more specifically. But the linkage of structured category prototypes within a taxonomic processing hierarchy makes it unclear when retrieval of semantic information will be governed by similarity to the prototype and when it will be governed by the spread of activation through the hierarchy. At least, we are unaware of any model that has successfully combined structured prototypes (at all levels of specificity) with a processing hierarchy. Such challenges to feature-based theories of semantic representation were noted in seminal work by Murphy and Medin (1985) and have spurred a productive program of research designed to investigate how existing knowledge constrains concept formation (Lin & Murphy, 1997; Murphy & Kaplan, 2000; Murphy & Allopenna, 1994).

In the PDP theory, semantic representations do not correspond to summary descriptions of discrete categories at different levels of specificity but instead are instantiated as patterns of activity across a common set of processing units. Hierarchical relationships are not explicitly represented but are implicit in the similarity relations among distributed representations (as schematized in Figure 4). It turns out that the latent encoding of hierarchical relationships within distributed representations allows for apparent domain-specific attribute weighting. A full explanation of such effects is beyond the scope of this article but was provided by Rogers and McClelland (2004; see chapter 6). Moreover, because semantic representations are not cast as discrete entities but as patterns of activation, with different representations instantiated across the very same processing units, other questions that are critical to prototype theories—for instance, questions about which categories are represented at which levels, which information is coded with which representations, how these representations are connected in the processing hierarchy, and how similarity-based processes interact with the spreading activation mechanism—are rendered moot in the PDP theory. The first advantage of instantiating the differentiation account within a PDP framework, then, is that it sidesteps difficult challenges faced by feature-based prototype theories.

The second way in which the PDP theory extends the differentiation theory is that it makes specific claims about the neural network that supports semantic abilities in the brain. That is, the processing units that mediate the interactions among perceptual, motor, and linguistic representations in our theory are understood as representing populations of neurons in the anterior temporal cortex; the distributed representations that arise across these units are understood to correspond to actual patterns of neural activity across said neurons; the patterns of connections between representations in our theory are understood to correspond to neuroanatomical pathways in the brain (see Rogers, Lambon Ralph, & Garrard, et al., 2004); and so on. These correspondences, though clearly coarse abstractions at best of the real neural system, nevertheless provide a useful framework for understanding why anterior temporal-lobe pathology should produce the particular pattern of disruption to semantic abilities observed in SD. This understanding has allowed us to make new predictions about

patterns of functional activation in healthy individuals (Hauk et al., 2006; Rogers, Hocking, Mechelli, Patterson, & Price, 2005; Rogers, Hocking, et al., in press), about patterns of impairment in SD on other kinds of tests (Adlam et al., 2006; Patterson et al., 2006), and in other neuropsychological syndromes (Lambon Ralph, Lowe, & Rogers, in press). Thus the PDP framework links the cognitive-level explanation of basic-level advantages offered by the differentiation theory to a neuroanatomical account that makes contact with neuropsychological and neuroimaging data.

Other Accounts

Recent work by Collin and McMullen (2005) and others (Thorpe, Gegenfurtner, Fabre-Thorpe, & Bulthoff, 2001) suggests that, for some natural categories, information about superordinate category membership can be discerned from low spatial-frequency visual information, whereas information about more specific category membership may depend to a greater extent upon high spatial-frequency information. One can imagine an explanation of the current results on the basis of these observations. Perhaps low spatial-frequency information is available earlier than is high spatial-frequency information but is, for some reason, insufficient to support a confident decision in free-response categorization tasks. Under this view, the general over basic pattern observed in the fast condition of Experiment 3 is supported by rapidly available low spatial-frequency information, and the observed pattern of results reflects performance characteristics of the visual system rather than of the semantic system. Such an explanation might also explain why ERP signals appear to differentiate superordinate categories earlier than basic-level categories, as reported by Large et al. (2004). Variations on Experiment 3 in the current work may be useful for assessing this hypothesis—one could, for instance, conduct the same experiment by using specific object names rather than pictures or by using items for which low spatial-frequency information is not sufficient to discern superordinate category membership. If the general over basic pattern depends upon the rapid availability of low spatial-frequency visual information, then it should disappear under either of these manipulations.

References

- Adlam, A. L., Patterson, K., Rogers, T. T., Salmond, C. H., & Hodges, J. R. (2006). Semantic dementia and primary progressive aphasia: Two sides of the same coin? *Brain*, *129*, 3066–3080.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–426.
- Barsalou, L. W., Simmons, W. K., Barbey, A., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, *7*, 84–91.
- Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.
- 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. *Cognitive Neuropsychology*, *20*, 27–47.
- Bozeat, S., Lambon Ralph, M. A., Patterson, K., Garrard, P., & Hodges, J. R. (2000). Nonverbal semantic impairment in semantic dementia. *Neuropsychologia*, *38*, 1207–1215.
- Bozeat, S., Lambon Ralph, M. A., Patterson, K., & Hodges, J. R. (2002). When objects lose their meaning: What happens to their use? *Cognitive, Affective, and Behavioral Neuroscience*, *2*, 236–251.
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, *2*, 913–919.
- Collin, C. A., & McMullen, P. A. (2005). Subordinate-level categorization relies on high spatial frequencies to a greater degree than basic-level categorization. *Perception & Psychophysics*, *67*, 354–364.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*, 407–428.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning & Verbal Behavior*, *8*, 240–247.
- Corter, J. E., & Gluck, M. (1992). Explaining basic categories: Feature predictability and information. *Psychological Bulletin*, *111*, 291–303.
- Cree, G., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, *23*, 371–414.
- Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, *1*, 123–132.
- Damasio, H., Grabowski, T. J., Tranel, D., & Hichwa, R. D. (1996, April 11). A neural basis for lexical retrieval. *Nature*, *380*, 499–505.
- DeCaro, S. A., & Reeves, A. (2002). The use of word-picture verification to study entry-level object recognition: Further support for view-invariant mechanisms. *Journal of Experimental Psychology: Applied*, *8*, 811–821.
- Devlin, J. T., Gonnerman, L. M., Andersen, E. S., & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of Cognitive Neuroscience*, *10*, 77–94.
- Farah, M., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality-specificity and emergent category-specificity. *Journal of Experimental Psychology: General*, *120*, 339–357.
- Garrard, P., & Hodges, J. R. (2000). Semantic dementia: Clinical, radiological, and pathological perspectives. *Journal of Neurology*, *247*, 409–422.
- Gauthier, I., Anderson, A. W., Tarr, M. J., Skudlarski, P., & Gore, J. C. (1997). Levels of categorization in visual recognition studied with functional MRI. *Current Biology*, *7*, 645–651.
- Glenberg, A., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory & Language*, *43*, 379–401.
- Gloor, P. (1997). *The temporal lobe and limbic system*. New York: Oxford University Press.
- Gorno-Tempini, M., Wenman, R., Price, C., Rudge, P., & Ciolotti, L. (2001). Identification without naming: A functional neuroimaging study of an amnic patient. *Journal of Neurology, Neurosurgery and Psychiatry*, *70*, 397–400.
- Gotts, S., & Plaut, D. C. (2002). The impact of synaptic depression following brain damage: A connectionist account of “access/refractory” and “degraded-store” semantic impairments. *Cognitive, Affective and Behavioral Neuroscience*, *2*, 187–213.
- Hauk, O., Patterson, K., Woollams, A., Watling, L., Pulvermüller, F., & Rogers, T. T. (2006). [Q:] When would you prefer a SOSSAGE to a SAUSAGE? [A:] At about 100 msec. ERP correlates of orthographic typicality and lexicality in written word recognition. *Journal of Cognitive Neuroscience*, *18*, 818–832.
- Heaton, R. K., Grant, D. A., & Berg, E. A. (1995). *Wisconsin Card Sorting Task*. Lutz, FL: Psychological Assessment Resources.
- Hinton, G. E. (1981). Implementing semantic networks in parallel hardware. In J. A. Anderson (Ed.), *Parallel models of associative memory* (pp. 161–187). Hillsdale, NJ: Erlbaum.
- Hinton, G. E. (1989). Learning distributed representations of concepts. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for*

- psychology and neurobiology* (pp. 46–61). Oxford, England: Clarendon Press.
- Hinton, G. E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, *98*, 74–95.
- Hodges, J. R., Garrard, P., & Patterson, K. (1998). Semantic dementia and Pick complex. In A. Kertesz & D. Munoz (Eds.), *Pick's disease and Pick complex*. New York: Wiley.
- Hodges, J. R., Graham, N., & Patterson, K. (1995). Charting the progression in semantic dementia: Implications for the organisation of semantic memory. *Memory*, *3*, 463–495.
- 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., Spatt, J., & Patterson, K. (1999). What and how: Evidence for the dissociation of object knowledge and mechanical problem-solving skills in the human brain. *Proceedings of the National Academy of Sciences*, *96*, 9444–9448.
- Humphreys, G., Lamote, C., & Lloyd-Jones, T. J. (1995). An interactive activation approach to object processing: Effects of structural similarity, name frequency, and task in normality and pathology. *Memory*, *3*, 535–586.
- Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General*, *126*, 248–277.
- Jolicoeur, P., Gluck, M., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, *19*, 31–53.
- Kellenbach, M., Brett, M., & Patterson, K. (2001). Large, colorful or noisy? Attribute- and modality-specific activations during retrieval of perceptual attribute knowledge. *Cognitive, Affective and Behavioral Neuroscience*, *1*, 207–221.
- Kello, C. (2004). Control over the time-course of cognition in the tempo-naming task. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 942–955.
- Kurbat, M. A. (1994). Structural description theories: Is RBC/JIM a general-purpose theory of human entry-level object recognition? *Perception*, *23*, 1339–1368.
- Lambon Ralph, M. A., Lowe, C., & Rogers, T. T. (in press). Both type and distribution of damage are critical for category-specific semantic deficits: Evidence from semantic dementia, herpes simplex viral encephalitis and a neural network model of conceptual knowledge. *Brain*.
- 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. *Journal of Cognitive Neuroscience*, *13*, 341–356.
- Large, M., Kiss, I., & McMullen, P. A. (2004). Electrophysiological correlates of object categorization: Back to basics. *Cognitive Brain Research*, *20*, 415–426.
- Lin, E. L., & Murphy, G. (1997). Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 1153–1169.
- McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, *4*, 310–322.
- McRae, K., de Sa, V. D., & Seidenberg, M. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 9–130.
- Mervis, C. B. (1987a). Acquisition of a lexicon. *Contemporary Educational Psychology*, *8*, 210–236.
- Mervis, C. B. (1987b). Child basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge, England: Cambridge University Press.
- Mervis, C. B., & Crisafi, M. A. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development*, *53*, 258–266.
- Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, *32*, 89–115.
- Mummery, C. J., Patterson, K., Price, C. J., Ashburner, J., Frackowiak, R. S. J., & Hodges, J. (2000). A voxel-based morphometry study of semantic dementia: Relationship between temporal lobe atrophy and semantic memory. *Annals of Neurology*, *47*, 36–45.
- Murphy, G. L. (1991). Parts in object concepts: Experiments with artificial categories. *Memory & Cognition*, *19*, 423–438.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 904–919.
- Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 70–84.
- Murphy, G. L., & Kaplan, A. S. (2000). Feature distribution and background knowledge in category learning. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *53(A)*, 962–982.
- Murphy, G. L., & Lassaline, M. E. (1997). Hierarchical structure in concepts and the basic level of categorization. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, concepts, and categories: Studies in cognition* (pp. 93–131). Cambridge, MA: MIT Press.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.
- Murphy, G. L., & Smith, E. E. (1982). Basic-level superiority in picture categorization. *Journal of Verbal Learning & Verbal Behavior*, *21*, 1–20.
- Murray, J. E. (1998). Is entry-level recognition viewpoint invariant or viewpoint dependent? *Psychonomic Bulletin & Review*, *5*, 300–304.
- Nakamura, K., Kawashima, R., Sugiura, M., Kato, T., Nakamura, A., Hatano, K., et al. (2001). Neural substrates for recognition of familiar voices: A PET study. *Neuropsychologia*, *39*, 1047–1054.
- Nestor, P. J., Fryer, T. D., & Hodges, J. R. (2006). Declarative memory impairments in Alzheimer's disease and semantic dementia. *NeuroImage*, *30*, 1010–1020.
- Osterrieth, P. A. (1944). Le test de copie d'une figure complexe: Contribution à l'étude de la perception et de la mémoire [The complex figure copy test: Contribution to the study of perception of memory]. *Archive de Psychologie*, *30*, 256–286.
- Patterson, K., & Hodges, J. (2000). Semantic dementia: One window on the structure and organisation of semantic memory. In J. Cermak (Ed.), *Handbook of neuropsychology: Vol. 2. Memory and its disorders* (pp. 313–333). Amsterdam: Elsevier Science.
- Patterson, K., Lambon Ralph, M. A., Jeffries, E. A., Woollams, A., Jones, R., Hodges, J. R., et al. (2006). "Presemantic" cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience*, *18*, 169–183.
- Plaut, D. C. (1999). Systematicity and specialization in semantics. In D. Heinke, G. W. Humphreys, & A. Olson (Eds.), *Connectionist models in cognitive neuroscience: Proceedings of the Fifth Annual Neural Computation and Psychology Workshop*. New York: Springer.
- Plaut, D. C. (2002). Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cognitive Neuropsychology*, *19*, 603–639.
- Pulvermüller, F., Harle, M., & Hummel, F. (2001). Walking or talking? Behavioural and neurophysiological correlates of action verb processing. *Brain and Language*, *78*, 143–168.
- Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing* (pp. 227–270). Cambridge, MA: MIT Press.
- Rey, A. (1941). L'examen psychologique dans les cas d'encéphalopathie

- traumatique [The psychological examination of cases of traumatic brain disease]. *Archives de Psychologie*, 28, 286–340.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning & Verbal Behavior*, 12, 1–20.
- Rogers, T. T., Hocking, J., Mechelli, A., Patterson, K., & Price, C. (2005). Fusiform activation to animals is driven by the process, not the stimulus. *Journal of Cognitive Neuroscience*, 17(3), 434–445.
- Rogers, T. T., Hocking, J., Noppeney, U., Mechelli, A., Gorno-Tempini, M., Patterson, K., et al. (2006). The anterior temporal cortex and semantic memory: Reconciling findings from neuropsychology and functional imaging. *Cognitive Affective and Behavioral Neuroscience*, 6, 201–213.
- 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 computational and neuropsychological investigation. *Psychological Review*, 111, 205–235.
- Rogers, T. T., Lambon Ralph, M. A., Hodges, J. R., & Patterson, K. (2003). Object recognition under semantic impairment: The effects of conceptual regularities on perceptual decisions. *Language and Cognitive Processes*, 18, 625–662.
- 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. *Cognitive Neuropsychology*, 21, 331–352.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382–439.
- Rosch, E., Simpson, C., & Miller, R. S. (1976). Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 491–502.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986, October 9). Learning representations by back-propagating errors. *Nature*, 323, 533–536.
- 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.
- Saffran, E. M. (2000). The organization of semantic memory: In support of a distributed model. *Brain and Language*, 71, 204–212.
- Saffran, E. M., & Schwartz, F. F. (1994). Of cabbages and things: Semantic memory from a neuropsychological perspective. In C. Umlita & M. Moscovitch (Eds.), *Attention and performance XV* (pp. 507–536). Cambridge, MA: MIT Press.
- Schwartz, M. F., Marin, O. S., & Saffran, E. M. (1979). Dissociations of language function in dementia: A case study. *Brain and Language*, 7, 277–306.
- Schyns, P. G. (1991). A modular neural network model of concept acquisition. *Cognitive Science*, 15, 461–508.
- Scott, S., Leff, A., & Wise, R. J. (2003). Going beyond the information given: A neural system supporting semantic interpretation. *NeuroImage*, 19, 870–876.
- Snowden, J. S., Goulding, P. J., & Neary, D. (1989). Semantic dementia: A form of circumscribed temporal atrophy. *Behavioural Neurology*, 2, 167–182.
- Tanaka, J., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, 23, 457–482.
- Tarr, M. J., & Bulthoff, H. H. (1998). *Object recognition in man, monkey and machine*. Cambridge, MA: MIT Press.
- Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory: Inferring “how” from “where.” *Neuropsychologia*, 41, 280–292.
- Thorpe, S. J., Gegenfurtner, K. R., Fabre-Thorpe, M., & Bulthoff, H. H. (2001). Detection of animals in natural images using far peripheral vision. *European Journal of Neuroscience*, 14, 869–876.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113, 169–191.
- Van Rullen, R., & Thorpe, S. J. (2001). Is it a bird? Is it a plane? Ultra-rapid visual categorization of natural and artificial objects. *Perception*, 30, 655–668.
- Warrington, E. K. (1975). Selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, 27, 635–657.
- Warrington, E. K., & James, M. (1991). Visual Object and Space Perception Battery. Edmunds, UK: Thames Valley Test Company.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107, 829–854.

Received July 7, 2006

Revision received December 22, 2006

Accepted December 28, 2006 ■