

Graded Modality-Specific Specialization in Semantics: A Computational Account of Optic Aphasia

David C. Plaut

Departments of Psychology and Computer Science
and the Center for the Neural Basis of Cognition
Carnegie Mellon University

A distributed connectionist model of semantics is presented in which semantic representations develop under the pressure of learning to mediate between multiple input and output modalities in performing various tasks. The system has a topographic bias on learning that favors short connections, leading to a graded degree of modality-specific functional specialization within semantics. The approach thus constitutes a middle ground between unitary- versus multiple-semantics accounts. As a result of the topographic bias, as well as the relative degrees of systematicity among tasks, damage to connections from vision to regions of semantics near phonology impairs visual object naming far more than visual gesturing or tactile naming, as observed in optic aphasia. Moreover, as in optic aphasia, the system is better at generating the name of an action associated with an object than at generating the name of the object itself, because action naming receives interactive support from the activation of action representations. The system also exhibits modality-specific impairments for grammatical categories (nouns vs. verbs) following lesions to or from regions of semantics which are partially specialized for objects (nouns) or actions (verbs).

Introduction

A central issue in the study of language and cognition concerns the organization of semantic representations for words, objects and their associated actions. A natural perspective on this issue is what has been termed the *unitary-semantics* account (Caramazza, Hillis, Rapp, & Romani, 1990; Hillis & Caramazza, 1995a; Hillis, Rapp, Romani, & Caramazza, 1990; Riddoch, Humphreys, Coltheart, & Funnell, 1988): that the meanings of objects are stored in a central, amodal semantic system that can be accessed from various input modalities (e.g., vision, touch, spoken and written language) and can be used to direct behavior in various output modalities (e.g., physical action, writing, speaking). Such an organization is parsimonious in that it allows knowledge derived from one modality to generalize automatically to others.

Financial support for this research was provided by a FIRST award from the National Institute of Mental Health (Grant MH55628). The computational simulations were run using customized software written within the Xerion simulator (version 3.1) developed by Drew van Camp, Tony Plate, and Geoff Hinton at the University of Toronto. Some preliminary results of the simulations were reported in Plaut (in press). I thank Marlene Behrmann, Jay McClelland, Tim Shallice, and the CMU PDP research group for helpful comments and discussion. Correspondence concerning this article may be sent to David Plaut, Mellon Institute 115-CNBC, Carnegie Mellon University, 4400 Fifth Avenue, Pittsburgh PA 15213-2683; email: pplaut@cmu.edu.

There are, however, certain empirical findings that seem problematic for a unitary-semantics account. A number of these come from the study of patterns of cognitive impairments that result from brain damage. A major focus of the current work is on the modality-specific aphasias, in which patients have naming deficits specific to a particular input modality. An example is optic aphasia, which typically results from a lesion to left medial occipital cortex and the underlying white matter (see Davidoff & Bleser, 1993; Riddoch, 1999, for reviews). Optic aphasic patients exhibit an impairment in naming visually presented objects that is not reducible to visual agnosia nor to a more general anomia. Agnosia is ruled out because the patients can demonstrate that they recognize the objects they cannot name—for example, by gesturing their use appropriately.¹ Anomia is ruled out because the patients can name the same objects from verbal definition or when presented in another input modality (e.g., tactile or auditory). Analogous selective naming deficits have been documented in the auditory modality (Denes & Semenza, 1975) and in the tactile modality (Beauvois, Saillant, Meininger, & Lhermitte, 1978). Table 1 (based in part on Table 1 of Sitton, Mozer, & Farah, in press) shows the relative correct performance on various relevant tasks of a number of optic aphasic patients, ordered in terms of the severity of their visual naming impairment (from mild to se-

¹ Those optic aphasic patients who do not exhibit intact gesturing to visually presented objects (e.g., Assal & Regli, 1980; Endo, Makishita, Yanagisawa, & Sugishita, 1996; Gil et al., 1985; Goldenberg & Karlbauer, 1998; Peña-Casanova & Roig-Rovira, 1985) can demonstrate relatively intact comprehension via other means (e.g., picture matching, category sorting).

Table 1
Percent Correct Performance of Optic Aphasic Patients on Various Tasks

Study	Task				
	Visual Naming	Visual Gesturing	Tactile Naming	Auditory Naming	Action Naming
Lhermitte and Beauvois (1973)	73	100	91	96	—
Casanova and Roig-Rovira (1985)	72	— ^{a,b}	100	100	—
Goldenberg and Karlbauer (1998)	70	— ^{a,b}	95	86	100
Gil et al. (1985)	64	100	—	100	—
Teixeira Ferreira et al. (1997)	53	95	81	78	75
Schnider et al. (1994)	50	58 ^b	75	100	83 ^c
Riddoch and Humphreys (1987)	46	75	75	100	—
Manning (2000)	38	83	100	81	—
De Renzi and Saetti (1997)	37	42 ^b	100	97	50
Endo et al. (1996)	30	17 ^b	46	90	—
Manning and Campbell (1992)	27	75	90	100	67
Luzzatti et al. (1998)	23	79	77	92	—
Coslett and Saffran (1992)	21	100	68	68	—
Larrabee et al. (1985)	20	— ^b	81	85	—
Raymer et al. (1997)	15	46 ^b	45	63	65 ^c
Hillis and Caramazza (1995a)	10	30 ^b	94	95	—
Poeck (1984)	8	75	35	90	—
Coslett and Saffran (1989b)	0	50	92	73	—

Note: “—” indicates task not tested or data unavailable.

^aGesturing performance stated to be equivalent to visual naming performance.

^bRelatively preserved visual comprehension demonstrated in semantic matching tasks.

^cGeneration of object name from viewing a gesture of its use.

vere).

Interestingly, in those patients for whom it has been tested, the visual naming deficit appears to be less severe when generating the names of actions compared with objects (Campbell & Manning, 1996; Druks & Shallice, 1996, 2000; Ferro & Santos, 1984; Goldenberg & Karlbauer, 1998; Manning & Campbell, 1992; Teixeira Ferreira et al., 1997; Zingeser & Berndt, 1990) or when naming objects from a pantomime or demonstration of their use (Campbell & Manning, 1996; Goldenberg & Karlbauer, 1998; Raymer et al., 1997; Schnider et al., 1994; Teixeira Ferreira et al., 1997). For example, when shown a set of 30 pictures of objects, patient AG (Campbell & Manning, 1996; Manning & Campbell, 1992) was only 27% correct in answering “What is the name of this?” but 67% correct in answering “What can you do with this?” When shown pictures of people using the same objects, AG was 63% correct at naming the object but 97% correct at naming the action. Similarly, patient CN (Teixeira Ferreira et al., 1997; see also Chanoine, Teixeira Ferreira, Demonet, Nespoulous, & Poncet, 1998), when shown a set of 24 real objects, was 46% correct at naming the objects but 75% correct at generating the specific action associated with the object.

The pattern of impaired and preserved performance in optic aphasia is difficult to reconcile with standard forms of the unitary-semantics account if it is assumed that naming requires semantic mediation.² Damage prior to or within se-

mantics would be expected to impair comprehension; damage within semantics or between semantics and phonology would be expected to impair naming regardless of the modality of input.

Based on these and other considerations, some researchers (e.g., Beauvois, 1982; Luzzatti et al., 1998; Shallice, 1987, 1993; Warrington, 1975, 1981; Warrington & McCarthy, 1987, 1994; Warrington & Shallice, 1984) have challenged the existence of a unitary, amodal semantic system, and instead have proposed that the semantic system is divided into separate subsystems on the basis of modality of input (visual, verbal, tactile, etc.) and/or type of information (natural kinds vs. artifacts, objects vs. actions, concrete vs. abstract concepts). This *multiple-semantics* account would seem to provide a relatively straightforward explanation of modality-specific aphasias. Optic aphasia, for example, would result from a disconnection between visual and verbal semantics: The intact access of visual input to visual semantics would support effective recognition and gesturing, but only non-visual input could access the verbal semantic information

² Some researchers (Davidoff & Bleser, 1993; Rapcsak, Rothi, & Heilman, 1987; Ratcliff & Newcombe, 1982) have suggested that optic aphasia results from impairment to a non-semantic pathway that maps high-level visual representations directly onto the verbal/phonological representations involved in naming, but the existence of such a pathway lacks any independent motivation or support.

necessary for naming (Beauvois, 1982; Beauvois & Saillant, 1985).

There are, however, a number of reasons why the multiple-semantics account is less than satisfactory. For one, strict modality- and category-specific subdivisions within semantics are considered by many researchers to be unparsonious (Riddoch et al., 1988) if not theoretically incoherent (Caramazza et al., 1990). More specifically, a disconnection between visual and verbal semantics would appear to be inconsistent with the relative sparing of generating action names from visual stimuli: Action information in visual semantics should be subject to the same disconnection from verbal semantics that prevents objects from being named.

Coslett and Saffran (1989b, 1992; see also Endo et al., 1996; Luzzatti et al., 1998; McCormick & Levine, 1983) have proposed an account of optic aphasia that is closely related to the multiple-semantics perspective, in which semantics is divided not by modality but by hemisphere. They assume that the left and right hemispheres represents various (but not identical) types of semantic information (see Beeman & Chiarello, 1998), but that only left-hemisphere semantics can support naming. The left occipital lesion that produces optic aphasia is assumed to destroy the high-level visual representation in the left hemisphere and the transmission (via the splenium of the corpus collosum) of right-hemisphere semantics to left-hemisphere semantics. The patients can thus comprehend and gesture to objects based on right-hemisphere visual and semantic processing, but cannot access left-hemisphere semantics in order to name the objects. A more thorough discussion of the hemisphere-based account is taken up in the General Discussion; at this point it is sufficient to point out that the properties it ascribes to left-hemisphere semantics are exactly the same as those ascribed to verbal semantics on the standard multiple-semantics account: it is required for naming but disconnected from vision.

As it turns out, detailed empirical testing (Gil et al., 1985; Hillis & Caramazza, 1995a; Riddoch & Humphreys, 1987) suggests that semantic access from vision in optic aphasia may not be as fully intact as originally thought. For example, patient JB (Riddoch & Humphreys, 1987) was only 54% correct at identifying which two of three visually presented objects were functionally or associatively related (e.g., cup and saucer vs. colander), although his performance was perfect when presented with their spoken names. JB was also impaired at picture-word matching and at matching top- and bottom-halves of objects and animals when target and distractors were semantically and visually similar. In those cases in which optic aphasic patients have exhibited intact semantic matching (e.g. Coslett & Saffran, 1989b, 1992), the tests involved distractors from a different category (e.g., pencil and paper vs. knife) and so could have been performed on the basis of relatively coarse semantic distinctions (Hillis & Caramazza, 1995a). Indeed, Hillis and Caramazza showed that their patient, DHY, was unimpaired at semantic matching with between-category distractors but was only 58% correct with within-category distractors. In fact, for a wide variety of semantic tasks, DHY was unimpaired on less strin-

gent versions but was significantly impaired when the tasks required more precise semantic information.

The observation of mild-to-moderate visual comprehension impairments in optic aphasic patients raises the possibility that partial damage between vision and a unitary semantic system may be able to account for their pattern of performance (Hillis & Caramazza, 1995a; Riddoch & Humphreys, 1987; Plaut & Shallice, 1993b). This account proposes that the patients have partial damage between high-level visual representations (e.g., structural descriptions) and the amodal semantic system. The damage is sufficient to impair naming, but comprehension (and gesturing) are relatively preserved due to the “privileged access” of structural descriptions to certain semantic features (Hillis & Caramazza, 1995a; Plaut & Shallice, 1993b) and/or to action representations (Riddoch & Humphreys, 1987; Rumiati & Humphreys, 1998).

McGuire and Plaut (1997) reported the results of a computational simulation that is consistent with this proposal. They trained a distributed connectionist network to map a visual or tactile representation of an object onto its phonological and action representations via a common set of intermediate units (corresponding to a unitary semantics system). The critical distinction among tasks was their degree of *systematicity*: the extent to which similar inputs map to similar outputs. The simulation employed abstract representations that were designed so that visually similar objects had similar associated actions (Gibson, 1979; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) but unrelated names. The sensitivity of learning to the systematicity between vision and action provided the basis for the “privileged access” assumed in other accounts. McGuire and Plaut demonstrated that mild damage between vision and semantics did, in fact, impair naming relative to gesturing (and other tests of comprehension). The magnitude of the effects in the simulation were, however, relatively small compared with those observed in some patients (e.g., Coslett & Saffran, 1989b, 1992), even though the degree of systematicity between vision and action in the simulation was, if anything, unrealistically high. Although the simulation was far from definitive, its quantitative inadequacy raises doubts about the sufficiency of differences in task systematicity alone to explain the dissociation of visual naming and gesturing in optic aphasia. More generally, an account based solely on systematicity (or another form of privileged access) would not seem to generalize to the analogue of optic aphasia in the auditory domain (Denes & Semenza, 1975), given the relative lack of systematicity between object sounds and actions.

Farah (1990; see also Campbell & Manning, 1996; Raymer et al., 1997) proposed an alternative approach to providing a unitary-semantics account of optic aphasia. She suggested that, in addition to a lesion between vision and semantics, optic aphasic patients have a second lesion between semantics and phonology. Each lesion is sufficiently mild that tasks involving only one of the damaged pathways (e.g., visual gesturing, non-visual naming) are relatively unimpaired, but visual naming—which requires both damaged pathways—is disproportionately impaired due to the super-additive effects of the two lesions. Although a standard con-

nectionist implementation of a vision-semantic-phonology pathway (Plaut & Shallice, 1993a) failed to exhibit the proposed superadditivity, Sitton et al. (in press) have recently provided computational support for the proposal using a system composed of linked modules exhibiting attractor dynamics. Following a combination of vision-to-semantic and semantic-to-phonology lesions, the model did, in fact, exhibit a superadditive impairment on visual naming relative to visual gesturing and non-visual naming. However, the effect held only for relatively mild impairment—down to about 70% correct on visual naming—and thus the approach fails to account for all but the most mildly impaired patients (see Table 1). As with the hemisphere-based account, the superadditive account is considered more fully in the General Discussion; the relevant point here, though, is that there is as yet no quantitatively adequate implementation of a unitary-semantic account of optic aphasia.

The current work attempts to articulate and support a theory of semantic organization that constitutes a middle ground between the unitary- and multiple-semantic accounts. Instead of casting semantics as entirely amodal or entirely modality-specific, the semantic system is claimed to have a graded degree of functional specialization that is influenced both by modality of input and by the nature of the information being represented. Moreover, this graded specialization is not hard-wired but emerges naturally in the process of learning to map among multiple modalities in performing a wide range of cognitive tasks.

Specifically, in conjunction with differences in relative task systematicity, the current work investigates an additional pressure for functional specialization: a topographic bias on learning favoring short connections (Jacobs & Jordan, 1992). The basic idea is that brain organization must permit sufficient connectivity among neurons to carry out the necessary information processing, but the total axon volume must fit within the confines of the skull. This constraint is severe: as Jacobs and Jordan note, if the brain's 10^{11} neurons were placed on a sphere and fully interconnected with $0.1 \mu\text{m}$ radius axons, accommodating the axon volume would require a sphere over 12.5 miles in diameter (Nelson & Bower, 1990). Clearly, connectivity must be as local as possible.

A bias favoring short connections can be instantiated within a connectionist network by assigning spatial locations to units and reducing the effectiveness of weight changes for each connection as a function of the distance between the connected units. As a result, during learning the network uses short connections as much as possible and develops significant weights on longer connections only where necessary. Jacobs and Jordan (1992) demonstrated that a bias favoring short connections can induce varying degrees of functional specialization among hidden units in a network trained to derive both the identity and position of a visual object (see Jacobs, 1997, for further discussion and results).

In the current context, internal (hidden) units form semantic representations that mediate between multiple input and output modalities. Under a topographic bias, the degree to which the internal semantic units participate in a particular input-output mapping depends on their proximity to the rel-

evant modalities. Semantic regions that are equidistant from multiple modalities learn to function in a relatively amodal way, whereas regions near a particular modality serve more modality-specific functions. Within such a system, the degree of modality specificity is graded and subject to the demands of the relevant tasks. On this account, and as demonstrated by the simulation presented here, optic aphasia arises from damage to the connections from high-level visual representations to semantic regions that are partially specialized for naming. The relative sparing of naming visual actions in optic aphasia results from the preserved support of action representations in generating action names.

In addition, the learned differentiation in the representation and processing of objects and actions indirectly provides an account for another, seemingly unrelated phenomena: the otherwise puzzling observation of impairments to grammatical categories, such as nouns or verbs, that are restricted to particular input or output modalities (Caramazza & Hillis, 1991; De Renzi & Pellegrino, 1995; Hillis & Caramazza, 1995b; Rapp & Caramazza, 1997, 1998). For example, patient HW (Caramazza & Hillis, 1991) performed much worse on verbs than nouns (22% vs. 56% correct, respectively) in tasks requiring spoken output, but was unimpaired when tested using written output. An analogous dissociation in written output was exhibited by patient PW (Rapp & Caramazza, 1998). Even more remarkable is patient EBA (Hillis & Caramazza, 1995b), who suffered both left temporal and left fronto-parietal lesions. EBA was much worse on the written comprehension of verbs than nouns (47% vs. 98% correct, respectively) but much better on the spoken production of verbs than nouns (72% vs. 13% correct, respectively). Thus, performance on nouns versus verbs can be doubly dissociated both within the same modality (spoken output, considering EBA and HW together) as well as between input and output modalities (considering EBA alone). Caramazza and Hillis (1991; Hillis & Caramazza, 1995a) originally interpreted these findings as implying that grammatical categories like nouns and verbs were represented separately within each input and output modality (see also De Renzi & Pellegrino, 1995). Rapp and Caramazza (1997, 1998) later acknowledged that it might be possible to explain the findings by positing impairments to modality-specific input or output pathways which selectively affected only one grammatical category, but they provided no basis for how such impairments might arise. On the current account, such impairments can arise following lesions to connections coming into or out of semantic regions that are partially specialized for objects (nouns) or actions (verbs) (see also Damasio & Tranel, 1993; Daniele, Suistolisi, Silveri, Colosimo, & Gainotti, 1994; Gainotti, Silveri, Daniele, & Giustolisi, 1995; Teixeira Ferreira et al., 1997).

To be clear, the notion of graded functional specialization within semantics is not novel to the current work; it derives from a perspective first articulated by Allport (Allport, 1985) and Warrington and McCarthy (Warrington & McCarthy, 1987) and later elaborated by Shallice (Shallice, 1988, pp. 302–304):

It may be useful to think of it (i.e., the semantic system) as a giant distributed net in which regions tend to be more specialised for different types of process.... The basis on which differentiation between processing regions within semantics would develop would include the most favoured modality of input for the process. Modality-specific pre-semantic classification subsystems would, thus, come to be more closely linked with some of the processing regions within the overall semantic system. So “visual semantic” and “verbal semantic” could be thought of as partially specialized subregions.... However, for explanations of this sort to be more than a speculation, a simulation of the hypothetical semantic system would be required.

This article presents just such a simulation.

Simulation

The current work adopts the perspective that semantics is a learned, internal representation that develops under pressure of performing a variety of tasks involving various input and output modalities (see also Rogers, Lambon Ralph, Patterson, McClelland, & Hodges, 2000; Rogers & Plaut, in press). The nature of these learned representations is shaped both by differences in task systematicity—the extent to which similar inputs map to similar outputs—and by a topographic bias on learning favoring short connections. Task systematicity is important because, in general, systematic tasks (e.g., visual gesturing) are learned more quickly and are more robust to damage compared with unsystematic tasks (e.g., visual naming; see McGuire & Plaut, 1997; Plaut & Gonnerman, 2000; Plaut, McClelland, Seidenberg, & Patterson, 1996). The topographic bias is important because, in learning to perform a given task, the network will use internal (hidden) units that are nearby the relevant modalities, and will use more distant hidden units only when necessary. As a result, the network will develop a graded degree of modality-specific specialization within its internal semantic representations. Localized damage to these representations (or to their incoming or outgoing connections) should, thus, give rise to modality-specific impairments on tasks—particularly those that are unsystematic—that require semantics.

Method

Network Architecture.

A continuous recurrent attractor network was trained to map either visual or tactile input to action and/or phonological output. The architecture of the network is shown in Figure 1. It has two input groups (*Vision* and *Touch*) and two output groups (*Action* and *Phonology*), each of which contains 20 units. These groups are connected with 225 *Semantic* (hidden) units, organized in a 15×15 grid. The semantic units receive inputs from both vision and touch and are bidirectionally connected with both action and phonology (which

are each fully intraconnected). The network also has two *Task* units which project to Semantics, Action, and Phonology, and whose function is described below. Including bias connections for non-input units (equivalent to a connection from an additional unit whose state is always 1), the network has a total of 28,555 connections.

In order to impose a topographic bias, it is necessary that units be localized in some metric space. For simplicity, units were assigned functional positions in two dimensions, exactly as depicted in Figure 1. This spatial configuration is, of course, not intended as a serious claim about the functional proximity of the corresponding brain regions. Rather, it was designed so that the architecture of the network would not bias the comparisons among the various tasks that are relevant for demonstrating a pattern of performance corresponding to optic aphasia. Thus, Vision and Touch are equidistant from Phonology (allowing an unbiased comparison of visual vs. tactile naming) and Vision is equidistant from Action and Phonology (allowing an unbiased comparison of visual naming vs. gesturing).

The activations of units in the network range between zero and one and change continuously in time as a function of their summed input from other units. To simulate on a digital computer, this continuous process is approximated by finite difference equations (with a discretization of τ), such that the new activation of a unit is a sigmoid function of a weighted average of its old summed input and the new input it is currently receiving (where τ is the weighting factor). Specifically, if $n_j^{[t]}$ is the instantaneous net input of unit j at time t , and $a_j^{[t]}$ is its output activation, then

$$n_j^{[t]} = \tau \sum_i w_{ij} a_i^{[t-\tau]} + (1 - \tau) n_j^{[t-\tau]} \quad (1)$$

$$a_j^{[t]} = \frac{1}{1 + \exp(-n_j^{[t]})} \quad (2)$$

Representations.

No attempt was made to model the detailed structure of any of the input or output modalities. Rather, sets of more abstract representations were defined in such a way that the similarity structure both within and between modalities approximated the central theoretical claims about the relevant tasks—namely, that there is considerable systematicity among visual, tactile, and action representation, but no systematicity between any of these and phonology.

Within each modality other than Phonology, 100 representations were generated to form 5 categories of 20 exemplars each. In the current context, a “category” consists of a set of patterns whose mutual overlap is relatively high compared with their overlap with patterns from other categories. Sets of patterns with this property were generated by first creating 5 random prototype patterns, each with 10 of 20 features equal to one and the rest zero. Each prototype was then used to generate 20 exemplars. In generating an exemplar, each feature of the prototype had a probability of 0.1 of changing its original value (from zero to one or vice versa). The result of

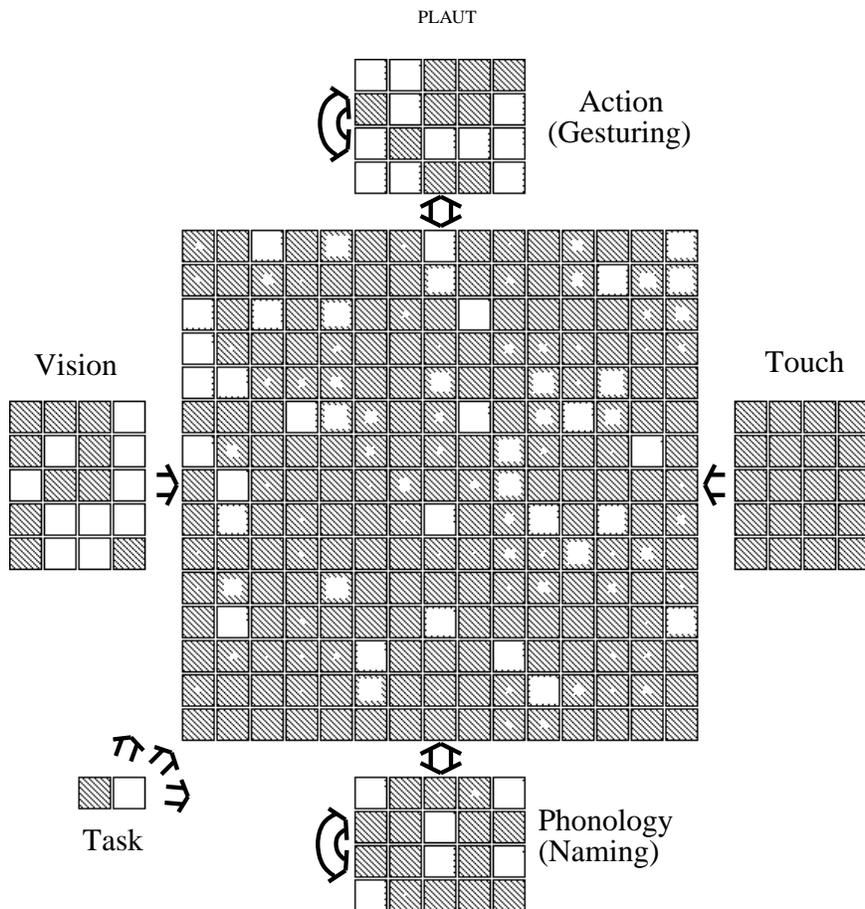


Figure 1. The architecture of the network. Each grey square constitutes a unit whose activity value is indicated by the size of the white region within the square. The activations shown are those generated by the fully trained network when presented with a visual object and instructed (with the Task units) to generate the action representation and name of the action associated with the object. Arrows indicate fully connectivity between the indicated unit groups; bidirectional arrows indicate two separate projections. The Task units are connected to all non-input groups. The positions of units in the figure correspond to their functional positions with respect to calculating connection lengths (as Euclidean distance).

the procedure is that the set of 20 exemplars generated from the same prototype have a high degree of overlap with that prototype (and, hence, with each other). Applying this procedure independently to Vision, Touch, and Action created 5 categories in each modality. Three-way partial systematicity among these modalities was enforced by assigning representations to objects in such a way that, if two objects were in the same category in one modality, they tended (with probability 0.8) to be in the same category in each of the other two modalities. In this way, similarity within either Vision, Touch, or Action was highly (but not perfectly) predictive of similarity within the other two domains.

Phonological representations, by contrast, were generated to form consonant-vowel-consonant strings over 20 features in three slots. The first slot of 7 features coded the onset consonant, the next slot of 6 features coded the vowel, and the last slot of 7 features coded the final consonant. Each of 16 possible consonants and 5 possible vowels were represented by a particular pair of units in the relevant slot being active. Out of the $16 \times 5 \times 16 = 1280$ possible names, 100 were chosen randomly and assigned as the names of objects, and another 100 were chosen randomly and assigned as the

names of the actions associated with the objects. The random assignment of names to objects and actions ensured that there was no systematicity between Phonology and any of the other domains. That is to say, phonological similarity is not at all predictive of visual, tactile, or action similarity.

Training Procedure.

The network was trained to perform two tasks, termed the *object* task and the *action* tasks. The task the network should perform for a given input was indicated by activating one of two Task units (see Figure 1). For the object task, the network was presented with either the Vision or Touch representation of an object (with the other modality set to all zeros) and trained to generate the name of the object over Phonology (*object naming*). In this case, no targets were specified for the Action units; the network was free to activate these units in any way. For the action task, the network was presented with either Visual or Touch input for an object and trained to generate both the name of the action associated with the object over Phonology (*action naming*), and the corresponding action representation over the Action

units (*gesturing*).³ In addition, as an approximation of direct experience with the actions, the network was presented with Action representations as input (with both Vision and Touch units set to zero) and trained to generate the name of the action.

Once the input was clamped on a particular modality, the remaining units in the network updated their states according to Equations 1 and 2 over a total of 5.0 units of time, with $\tau = 0.2$ (25 updates). The resulting unit activations were then compared with the appropriate targets for the presented object and the task being performed. Error in the network’s performance was defined as the cross-entropy (Hinton, 1989) between the generated activations, $a_j^{[t]}$, and the target activations, a_j^* , over the last unit of time (5 updates).

$$E = \sum_{4 < t \leq 5} \sum_j a_j^* \log(a_j^{[t]}) + (1 - a_j^*) \log(1 - a_j^{[t]}) \quad (3)$$

Error derivatives were then calculated using a version of back-propagation-through-time adapted for continuous units (Pearlmutter, 1989).

The topographic bias on learning was implemented by scaling the magnitude of the derivative on each connection by a Gaussian function ($SD = 10$) of its length.⁴ This produces a scaling of near 1.0 for the shortest connections, and near 0.1 for the longest. Thus, learning on the shortest connections was 10 times more effective than learning on the longest connections. For simplicity, this scaling was not applied to the connections from the Task units as none of the current theoretical issues relate to topographic influences on executive control.

Once the derivatives were scaled, the weights were updated using a learning rate of 0.01, momentum of 0.9, and weight decay of 0.00005. Following this, unit states were reinitialized, another object, task, and modality were selected, and the process was repeated. The network was trained on a total of 110,000 object presentations, corresponding to 220 presentations per condition (object \times modality \times task). At this point, all output activations generated by the network in all conditions were on the correct side of 0.5; 96% (672/700) were within 0.2 of their target (0 or 1).

Lesioning Procedure.

Lesions to the network were assumed to be topographically constrained by the spatial layout of the architecture. Specifically, lesions to the Semantic units (or to their incoming or outgoing connections) were defined to have a center located at a particular position within Semantics, and to fall off in severity with increasing distance from this center. Most of the results reported in this paper are for lesions administered to Vision-to-Semantics connections, on the hypothesis that optic aphasia arises from impaired semantic access from vision. Taking such a lesion as an example, the probability of removing a given connection was a 2D Gaussian function of the position of the receiving Semantic unit relative to the center of the lesion. Thus, connections to the unit at the

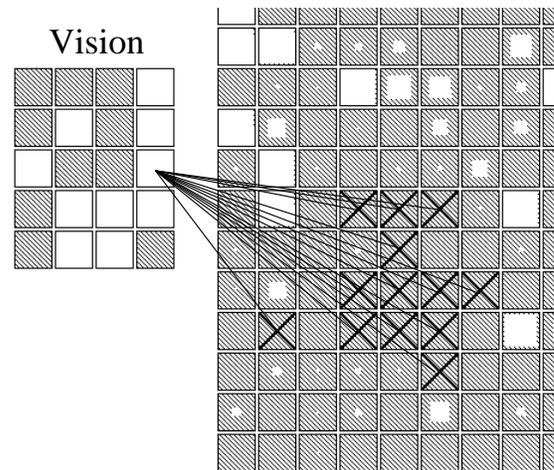


Figure 2. An illustration of a lesion to Vision-to-Semantics connections with $SD = 1.5$. Only the lesioned connections from a given Vision unit are shown, with the corresponding Semantic units marked with Xs. The portion of the network shown corresponds to a region in the lower left of Figure 1.

center of the lesion were lesioned with probability 1.0; connections to progressively more distant Semantic units were less and less likely to be lesioned. The severity of the lesion was controlled by the standard deviation (SD) of the Gaussian. Analogous topographic lesions could also be applied to outgoing connections from Semantics, or to the Semantic units themselves. Figure 2 illustrates a lesion to Vision-to-Semantics connections with $SD = 1.5$.

Results and Discussion

Semantic Similarity.

The most basic claim of the current work is that the semantic system corresponds to internal (hidden) representations that mediate between input and output modalities. Thus, before considering the effects of damage, it is important to verify that the learned hidden representations in the network do, in fact, exhibit some of the basic characteristics of semantic representations.

A full exploration of the extent to which the model exhibits all of the various behavioral phenomena that are relevant to semantics is, of course, beyond the scope of the current work. For present purposes, it was considered sufficient to examine the extent to which the degree of similarity

³ The motivation for activating both Action and Phonology in the action task is that the network is generating everything it knows about the action. In contexts where only one overt response is called for, it was assumed that the effects of activity in the inappropriate output modality would be suppressed by downstream inhibitory mechanisms.

⁴ This procedure differs slightly from the one used by Jacobs and Jordan (1992), who scaled the magnitude of weight decay rather than error derivatives. This change was purely for computational convenience and is immaterial to the results.

among hidden representation in the network in various conditions mirrored the relative levels of similarity among semantic representations (as suggested by empirical studies). Specifically, the network was run on all 100 objects under both visual and tactile presentation in the action task, and the correlation (over units) between each resulting hidden representation and every other was computed.⁵ Mean correlations were then computed across object pairs as a function of whether the objects were drawn from the same or different “categories” and presented in the same or different modality, as well as for the same object presented in different modalities. For this purpose, objects were considered to be in the same category if their visual representations were generated from the same prototype; an analysis based on tactile categories would produce essentially the same results as visual and tactile categories were generated in the same manner.

The results of this analysis are presented in Figure 3. First, the network shows a basic semantic relatedness effect: the similarity in hidden representations for objects in the same category is much greater than for objects in different categories. In addition, among related objects, there is an effect of modality: presentation within the same modality produces greater representational similarity than cross-modal presentation. Finally, and perhaps most important, the network is highly sensitive to the *identity* of an object regardless of its modality of presentation: a given object is more similar to itself across modalities than are objects from the same category presented in the same modality. It is interesting, though, that the network does not generate *identical* representations for an object regardless of modality (i.e., a correlation of 1.0); rather, the semantic representation of an object does retain some sensitivity to the modality in which the object was presented. If degree of pattern similarity is related to the degree of facilitation that would be expected of one stimulus followed the other in the course of processing (see Masson, 1995; Plaut & Booth, in press), then the overall pattern of results seems in accordance with the basic pattern of results from, for example, studies of repetition and semantic priming (see, e.g., Neely, 1991). Thus, at least at a general level, there is some justification for interpreting the hidden representations of the network as corresponding to semantics.

Nature of Semantic Specialization.

A central claim of the current work is that semantic representations exhibit a graded degree of modality specificity. One question, then, is whether learning has, in fact, produced internal representations with this structure. More specifically, it is important to demonstrate that the differences in task systematicity and the topographic bias on learning have not led the system to develop functionally separate modality-specific subregions within semantics. If this were true, the network might still be an interesting implementation of the multiple-semantics account (Shallice, 1987; Warrington, 1975; Warrington & McCarthy, 1987, 1994; Warrington & Shallice, 1984) but it would not constitute an alternative account.

If the network had developed modality-specific subregions, it would be natural to expect that information from

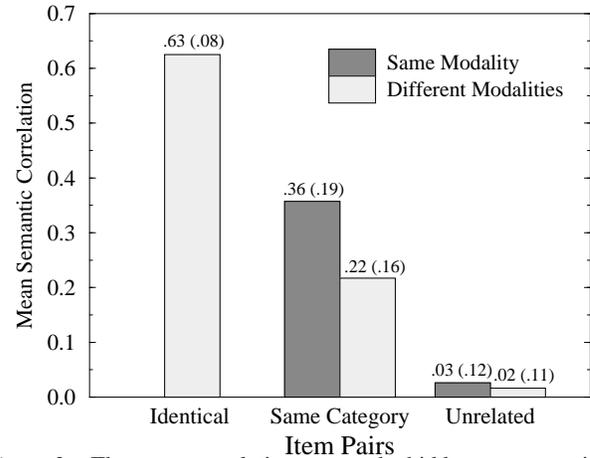


Figure 3. The mean correlation among the hidden representations generated by pairs of objects in the action task, as a function of their relation to each other (identical, from the same visual category, or from different categories) and whether they were presented in the same or different modalities. Standard deviations are shown in parentheses.

a given input (or output modality) would be represented primarily by the Semantic (hidden) units that are nearest that modality. For example, the left half of Semantics might represent only visual information whereas the right half might represent only tactile information. The results on similarities among semantic representations already provide some evidence that this is not the case: the representations generated by cross-modal presentation of the same object are highly correlated (i.e., show a considerable degree of overlap). Nonetheless, it seems important to explore the degree of modality-specific specialization in the system more directly.

To do this, the average activation of each Semantic unit was calculated when the network was presented with each object in each input modality. Figure 4a shows the activations for Vision input and 4b shows the activation for Touch input. Visual inspection of the figure suggests that there is little if any evidence of modality specificity—Semantic units that are distant from the input modality (i.e., on the right of Figure 4a for Vision input, and on the left of 4b for Touch input) are, on average, activated as strongly as units that are closer to the input modality. This impression is confirmed in Figure 4c, which plots the mean activation values as a function of the horizontal position of the Semantic unit, averaging across vertical positions. In an analysis of variance (ANOVA) of mean unit activations over all tasks, there was no reliable effect of the x-position of the unit, the modality of presentation, or the interaction of these factors ($ps > .29$). Thus, there is no difference in the extent to which the two modalities generate greater activation over closer Semantic units as compared with more distant units; the entire seman-

⁵ The action task was chosen because it involves activating both Action and Phonology representations; the object task produces qualitatively equivalent results, although the overall levels of similarity are somewhat lower.

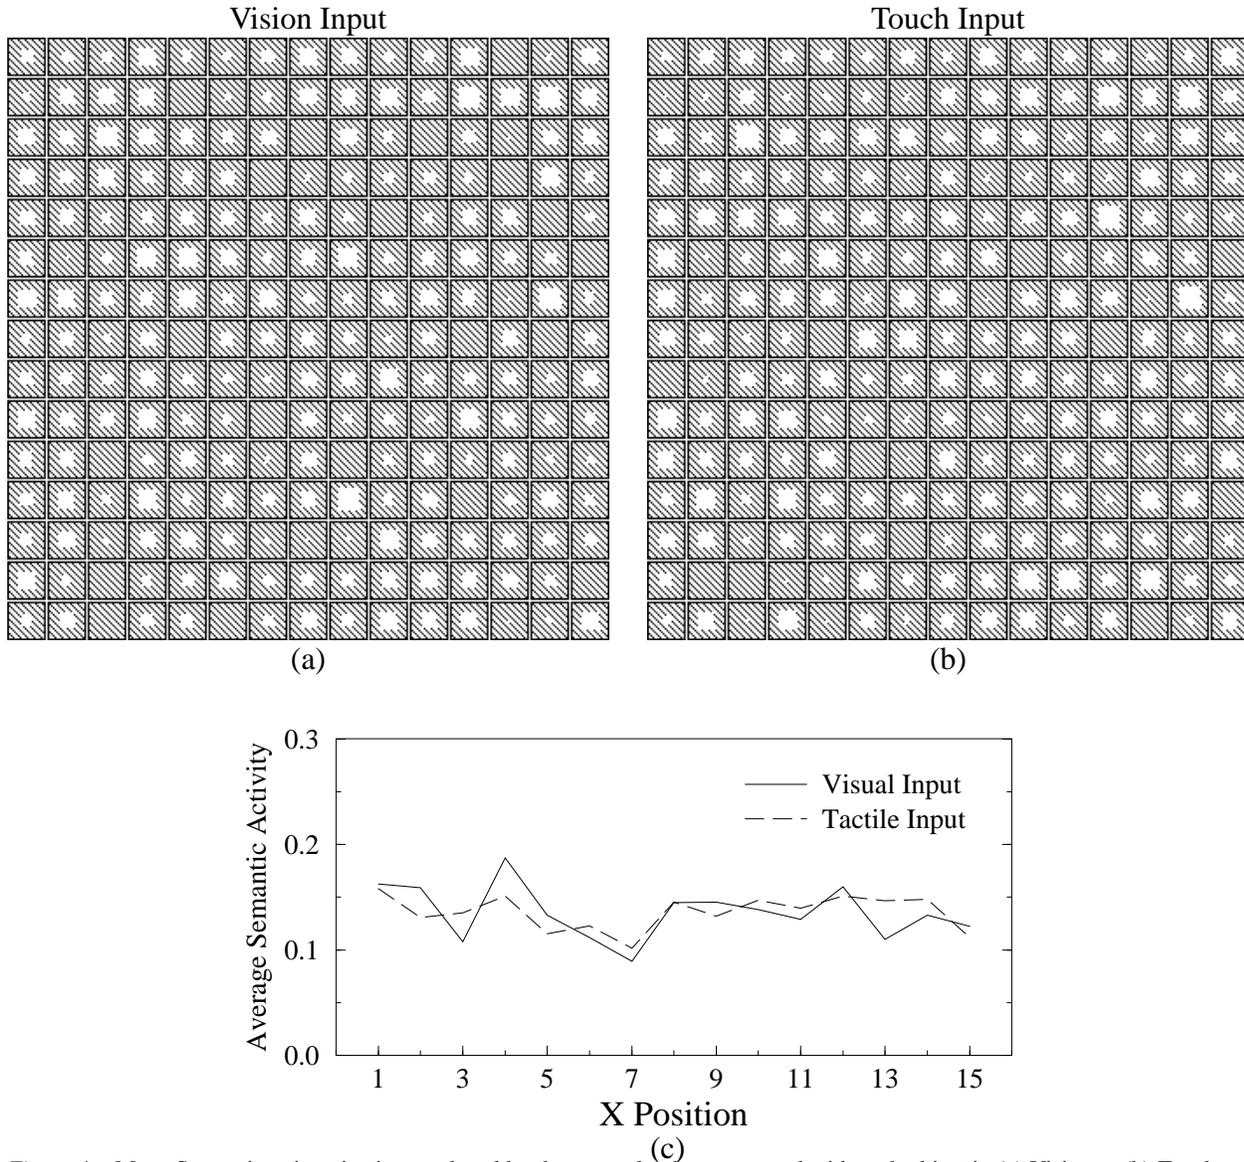


Figure 4. Mean Semantic unit activation produced by the network when presented with each object in (a) Vision or (b) Touch, as well as (c) the mean activations plotted as a function of the horizontal position of the Semantic unit (i.e., also averaged over vertical position). In (a) and (b), a full white square indicates a mean activation of 1.0; an entirely grey square indicates a mean of 0.0.

tic system is involved in representing both visual and tactile input.

It is possible that, in processing an input, the relative importance of a given Semantic unit is carried less by its overall level of activation than by the degree to which its activation *varies* as a function of the specific object and the modality of presentation. Figure 5 plots the mean variance in Semantic unit activations for visual and tactile input, as a function of the horizontal position of the unit. There is, in fact, a statistically reliable interaction of modality of presentation and spatial position, such that the activations of Semantic units vary slightly more when driven by input from a closer as compared with a more distant modality ($F_{14,870} = 2.66$, $p < .001$). The effect is clearly very weak, however, and nothing

like the magnitude that would be expected from a system with separate, modality-specific subsystems.

In some respects, the finding of very little spatial functional specialization by modality in the Semantic activations of the network is surprising in light of the topographic bias on learning that favored short over long connections. In fact, as would be expected, this bias did produce a graded modality specificity in the underlying weights in the network. This can be seen in Figure 6, which present data analogous to Figure 4 but for the absolute magnitude of incoming weights to each Semantic unit. Here, the interaction of modality by spatial position is very clear ($F_{14,420} = 28.16$, $p < .001$). Semantic units near each modality have much larger weights on the connections coming from that modality. Note that the rel-

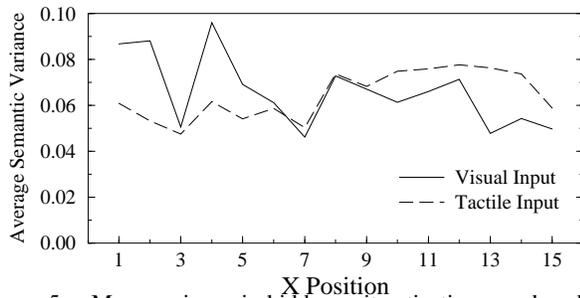


Figure 5. Mean variance in hidden unit activations produced by the network when presented with each object in each modality, as a function of the horizontal position of the unit.

ative magnitudes of the largest to smallest weights is about ten to one, which exactly mirrors the range of the Gaussian scaling factors implementing the topographic bias. The correspondence simply demonstrates that the relative magnitude of weights is closely related to the relative magnitudes of the error derivatives they experience during training.

Why, then, weren't these differences in weight magnitudes reflected more directly in differences in unit activations, given that unit activations ultimately depend on the weights (see Equation 1)? The reason is that the network is highly interactive. Consider the time course of the network's settling process in response to input in a given modality. The stronger weights from this modality to nearby Semantic units cause these units to become more active initially than more distant units. The nearby Semantic units then begin to activate units in each output modality, and these units, in turn, activate other Semantic units, some of which are more distant from the input modality. Gradually, unit interactions within the network activate the full hidden and output representations of the presented input. Because all of the knowledge in the network (i.e., the connection weights) ultimately contributes to generating unit activations, the final activations are not strongly biased by spatial position.

In summary, the effects of task systematicity and topographic bias did induce graded modality specificity in the underlying knowledge of the network, but this property does not translate directly into specificity in the representations derived by the network in processing different modalities of input. In this way, the network's Semantic representations are neither completely amodal (as in unitary-semantics accounts) nor completely modality specific (as in multiple-semantics accounts). Even so, the relative similarities among the representations capture at least broad aspects of what is required of semantic representations, in terms of having greater similarity within than between categories.

Effects of Lesion Location.

The central focus of the current work is on whether damage to a system with a graded degree of modality specificity within semantics can account for the overall pattern of performance of optic aphasic patients. The current work proposes that the damage giving rise to optic aphasia primarily affects the mapping from vision into regions of semantics that are partially specialized for generating phonological out-

put. The more specific assumption is that such graded specialization arises as a result of a topographic bias on learning. As will be demonstrated, this assumption is critical in producing the magnitude of the dissociation exhibited by optic aphasic patients in naming versus gesturing to visual stimuli.

To understand the effect of the spatial location of lesions, the network's performance was tested after probabilistic lesions to connections from Vision units to Semantic units, in which the likelihood of removing a given connection was a Gaussian function ($SD = 1.5$) of the location of the Semantic unit from the specified center of the lesion (see Figure 2). To avoid sampling artifacts, 10 instances of lesion centered at each of the 225 Semantic positions were administered. After each lesion, the network's performance was tested for each of the 100 objects on tactile object naming, visual object naming, visual action naming, and visual gesturing. The response of the network was considered correct if the activations of all of the output units in the relevant modality were on the correct side of 0.5.

Figure 7 shows the levels of correct performance on visual object naming and visual gesturing for Vision-to-Semantics lesions centered at each location of Semantics. Two general effects are apparent. The first is that the locations of damage that are most detrimental to the two tasks differ. Not surprisingly, a given task is most impaired by Vision-to-Semantics lesions that are centered at Semantic locations that are roughly equidistant between the input and output modalities involved in the task. Thus, visual object naming is impaired most by lesions in the lower-left corner of Semantics; whereas visual gesturing is impaired most by upper-left lesions. Second, and more striking, is that the overall levels of impairment are much greater for visual object naming than for visual gesturing. This difference arises directly from the relative degree of systematicity of the two tasks. Damage causes a much greater impairment on a unsystematic task like naming than on a systematic task like gesturing.

The relative difference in performance on visual gesturing versus object naming is clearer in a plot of the absolute differences in levels of correct performance (see Figure 7c). The majority of lesion locations produce an advantage for gesturing over naming, with the magnitude of the advantage dependent on how close the lesion is to Phonology (i.e., the bottom edge of Semantics). Interestingly, the network also exhibits the reverse pattern of performance for some lesions near Action: greater impairment on gesturing than on naming. Given that the lesions are to incoming connections from Vision, the behavior of the network under these conditions corresponds to an impairment in generating action that is restricted to (or at least most severe for) visual input: optic apraxia. Note that many optic aphasic patients also exhibit optic apraxia (Assal & Regli, 1980; Coslett & Saffran, 1989b; Endo et al., 1996; Casanova & Roig-Rovira, 1985); this co-occurrence can arise in the network following lesions to connections from Vision to regions of Semantics near both Action and Naming (while sparing much of the projection to the remainder of semantics). However, the model goes on to predict that it should be possible for brain damage to produce optic apraxia without optic aphasia, although the relatively

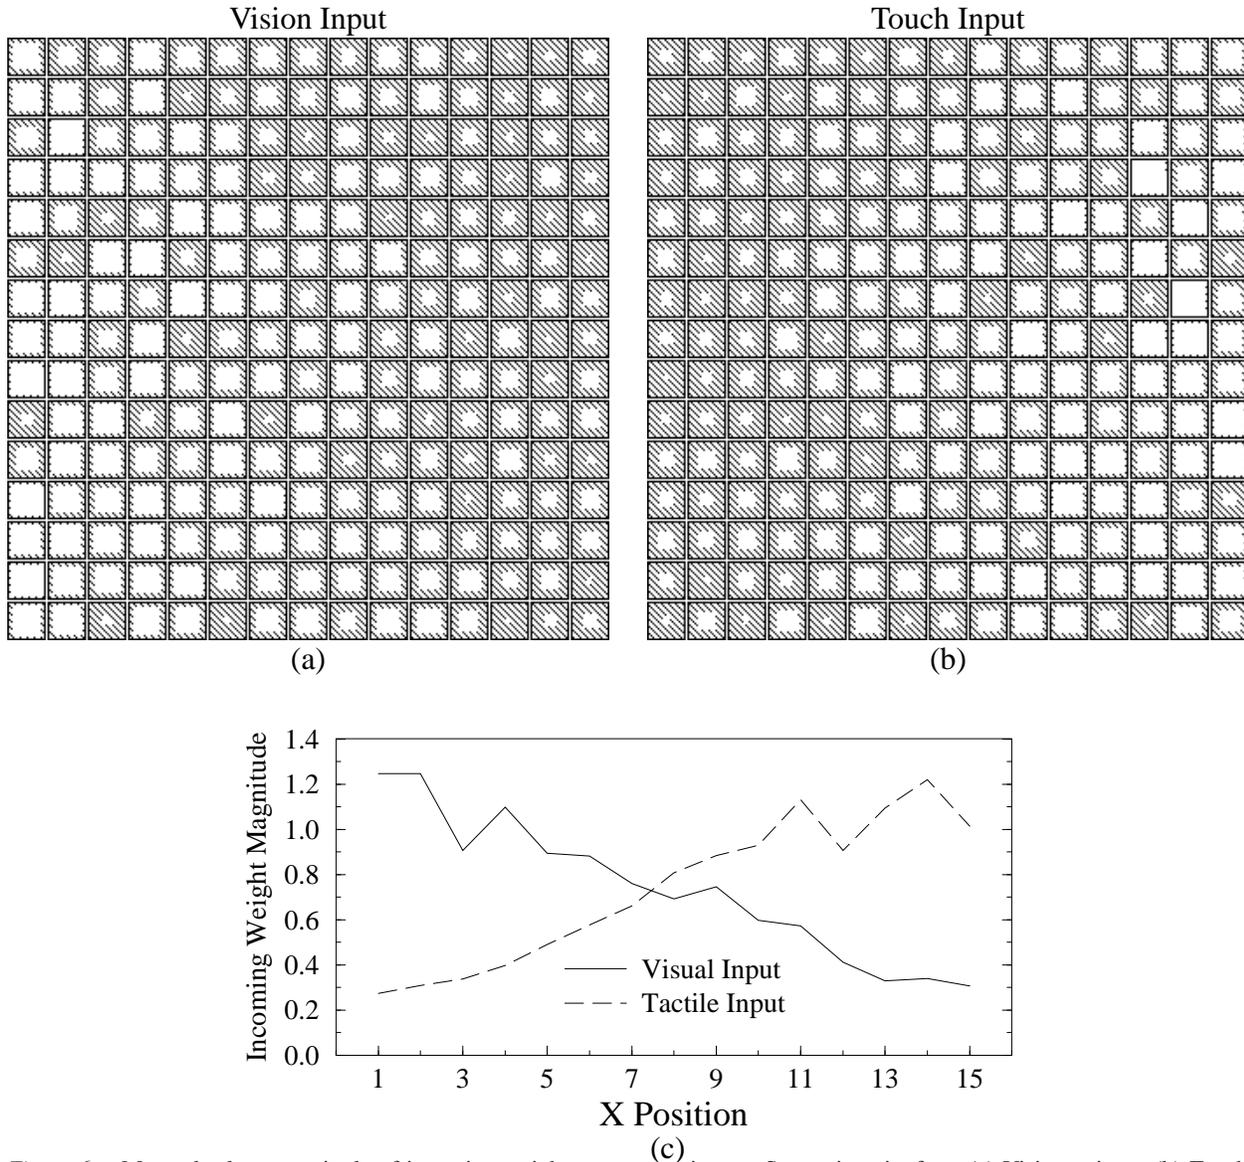
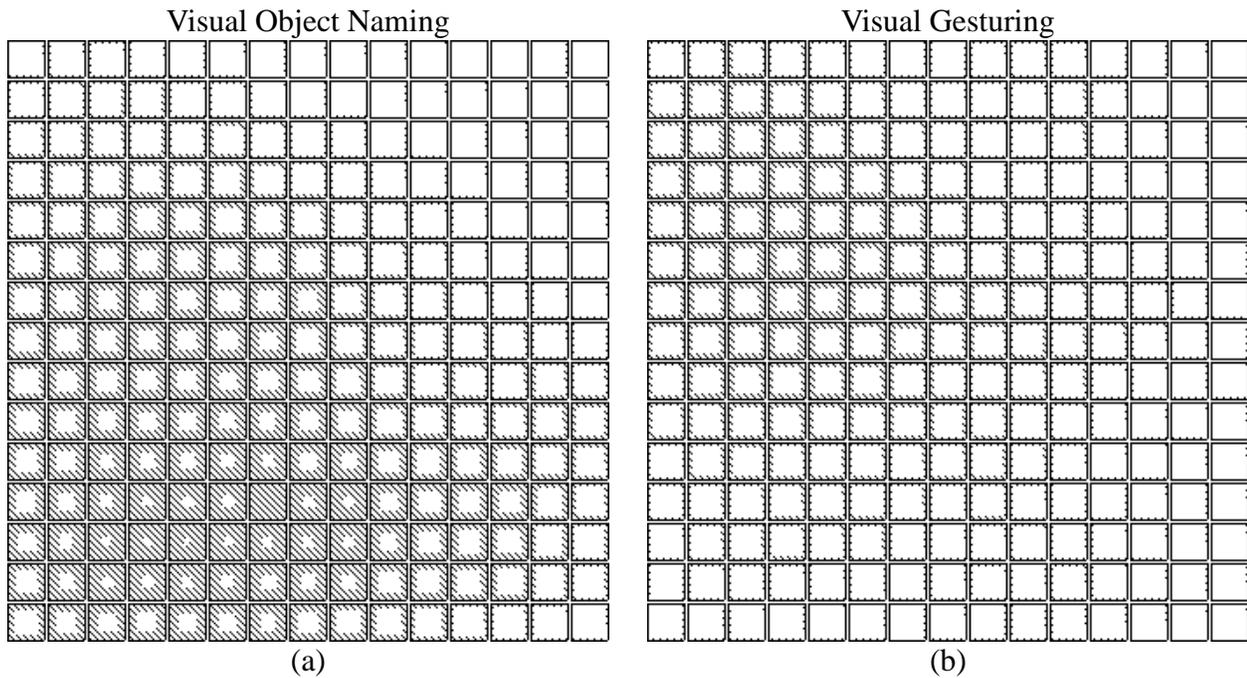


Figure 6. Mean absolute magnitude of incoming weights on connections to Semantic units from (a) Vision units or (b) Touch units, as well as (c) the mean weight magnitudes plotted as a function of the horizontal position of the Semantic unit (i.e., also averaged over vertical position). In (a) and (b), a full white square indicates a weight magnitude of 2.0.

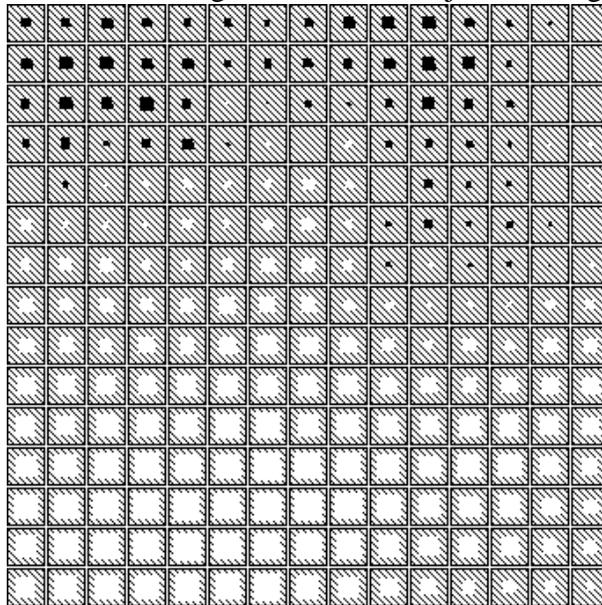
small region of lesion locations giving rise to this pattern in the model suggests such patients may be rare.

The existing empirical evidence for such a pattern is weak but suggestive. A number of researchers have reported cases in which hand posture during reaching for objects can be dissociated from object recognition (see Jeannerod, 1997; Milner & Goodale, 1995, for reviews). De Renzi, Faglioni, and Sorgato (1982) reported a number of cases of modality-specific impairments of object use but most exhibited selective preservation of gesturing to tactile input and the few who showed a relative impairment for visual versus verbal input did not have their visual recognition abilities tested carefully. Impaired gesturing to visual compared with verbal input, with intact visual recognition, has been documented in only

two cases. Patient CJ (Riddoch, Humphreys, & Price, 1989) was much better at generating right-handed gestures of the use of an object when given its spoken name (92% correct) than under combined visual and tactile presentation (68% correct). Patient GF (Pilgrim & Humphreys, 1991) showed a similar pattern of performance, although gesturing to spoken names was also somewhat impaired. Finally, some patients with intact visual recognition have exhibited impaired gesturing to visual compared with tactile input. Graham, Zeman, Young, Patterson, and Hodges (1999) report on a patient whose gestures of tool use were only 25% correct under visual presentation (and to spoken names) but 92% correct when also allowed to grasp the tool, and 75% correct under tactile presentation without visual input. Along similar



Visual Gesturing minus Visual Object Naming



(c)

Figure 7. Correct performance on (a) visual object naming and (b) visual gesturing, following Gaussian lesions to Vision-to-Semantics connections ($SD = 1.5$), and (c) the absolute difference between these levels of performance [(b) minus (a)], as a function of the center of the lesion within Semantics. The size of the square at each position corresponds to the mean performance level following lesions centered at that position. In (a) and (b) depicted values range from 100% correct (full white square) to 20% correct (no white square). In (c), white squares indicate better gesturing than naming; black squares indicate better naming than gesturing.

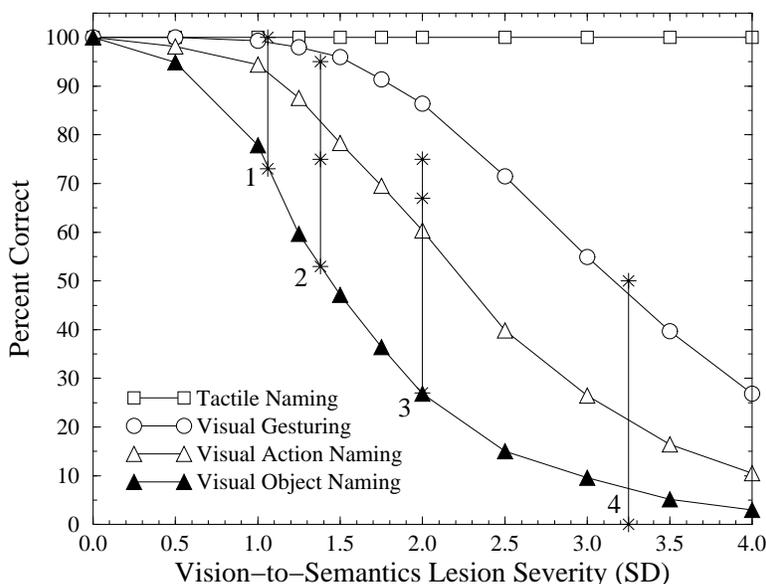


Figure 8. Correct performance of the network on various tasks as a function of lesion severity, and corresponding levels of performance of four optic aphasic patients, plotted on vertical lines with asterisks and labeled with the corresponding numbers from Table 2. For each line, the top asterisk indicates performance on visual gesturing and the bottom indicates performance on visual object naming. The middle asterisk, when present, indicates performance on visual action naming.

Table 2
Percent Correct Performance of Four Optic Aphasic Patients on Various Tasks

Study	Task			
	Visual Naming	Visual Gesturing	Tactile Naming	Action Naming
1. Lhermitte and Beauvois (1973)	73	100	91	
2. Teixeira Ferreira et al. (1997)	53	95	81	75
3. Manning and Campbell (1992)	27	75	90	67
4. Coslett and Saffran (1989b)	0	50	92	

Note. The numbers assigned to studies correspond to the labels on vertical lines in Figures 8 and 12 indicating the corresponding patient's performance

lines, Schnider, Hanlon, Alexander, and Benson (1997) report on a group study in which ideomotor apraxics with left-hemisphere damage made poorer gestures of tool use under visual presentation and to spoken names compared with tactile presentation or when holding the tool during the gesture.

Effects of Lesion Severity.

The results from the analysis of lesion location indicate that lesions to connections from Vision to regions of Semantics near Phonology produce the largest dissociation in performance on visual gesturing versus visual naming. To explore whether the network could provide a quantitative match to the performance levels of individual patients, a particular lesion location was chosen and the severity of the lesions (i.e., the *SD* of the Gaussian probability distribution for removing connections) was varied systematically. Specifically, the network's performance was measured on tactile naming, tactile object naming, visual object naming, visual action naming, and visual gesturing following Vision-to-Semantics

lesions centered on the seventh unit from the left in the bottom row of Semantics (see Figure 7c). Figure 8 shows the results, averaged over 40 instances of lesion at each level of severity. For comparison, the figure also shows (with vertical lines and asterisks) the levels of performance of four optic aphasic patients (taken from Table 1 and replotted in Table 2 for clarity) on visual naming, visual gesturing, and for two, action naming. These patients were selected to span the full range of severity in performance on visual naming.

As the figure shows, the lesioned network is far more impaired at visual object naming than at either visual gesturing or tactile naming (the latter is unaffected by Vision-to-Semantics lesions). Thus, the network is exhibiting the hallmark characteristics of optic aphasia. In fact, when compared with the levels of performance of the four patients, the network does a reasonable job at matching the magnitudes of the dissociation between visual naming versus gesturing across a range of severity. The only major discrepancy is that the network shows a *larger* dissociation than Manning

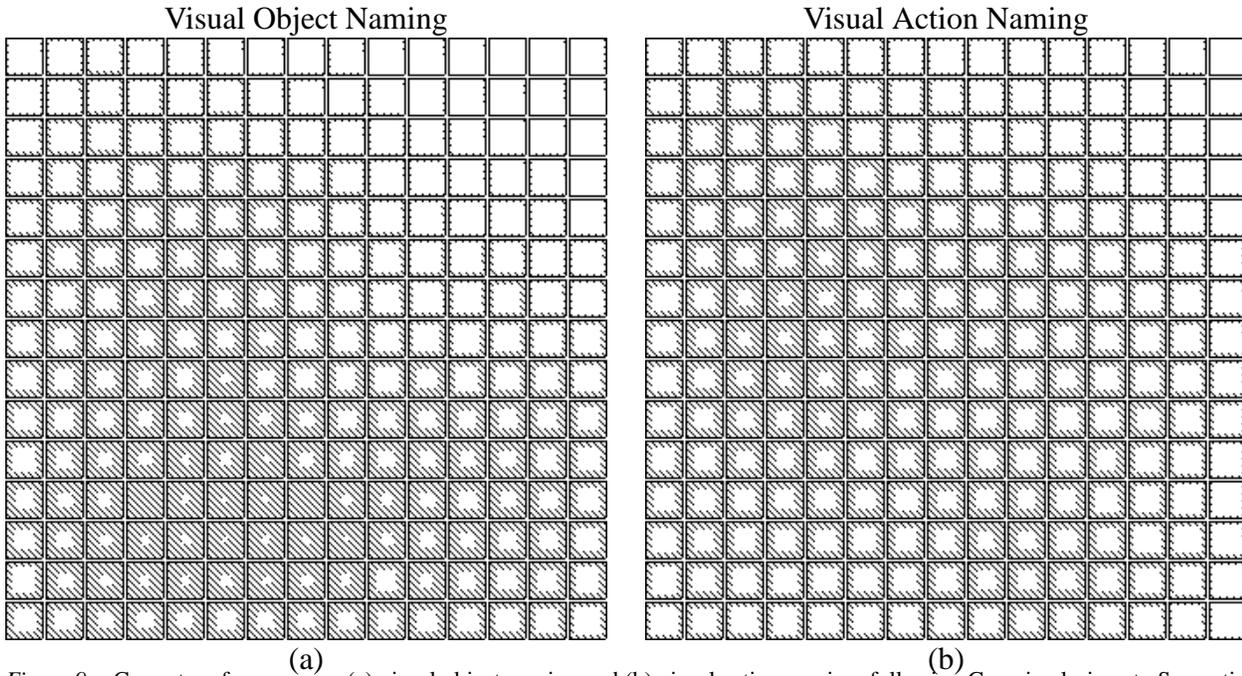


Figure 9. Correct performance on (a) visual object naming and (b) visual action naming, following Gaussian lesions to Semantic units ($SD = 2.0$), as a function of the center of the lesion within Semantics. The depicted values range from 100% correct (full white square) to 20% correct (no white square).

and Campbell's (1992) patient when matched on visual object naming performance.

Moreover, like the patients in whom it has been tested (e.g., Goldenberg & Karlbauer, 1998; Manning & Campbell, 1992; Teixeira Ferreira et al., 1997), the network is far better at naming the action associated with a visually presented object than at naming the object itself. Why is generating action names relatively preserved compared with generating object names? The reason is that the network learns to rely on support from the Action representation when generating an action name. It does this because a given visual input by itself is ambiguous with respect to the correct phonological output—it could be either the object name or the action name. Because the network is trained to generate the Action representation in conjunction with generating the action name over Phonology, it is natural for it to use the derived Action representation to resolve the ambiguity and override the object name. The damaged network generates the Action representation relatively successfully from visual input (as evidenced by the good visual gesturing performance), and thus this information is available to support relatively good (although far from perfect) naming of actions in the face of impaired object naming.

Lesions to Semantic Units.

To provide evidence for the above explanation, the network's performance was measured on visual object naming and visual action naming following lesions directly to the semantic units themselves, as a function of the location of lesion. On the hypothesis that action naming relies on interactions with Action representations, the lesion locations

producing the greatest impairment on visual action naming should be closer to the Action modality than the corresponding location for visual object naming, even though both tasks require mapping Vision to Phonology.

Figure 9 shows the levels of correct performance on naming objects versus actions from Vision following Gaussian lesions to the Semantic units ($SD = 2.0$). Indeed, the lesions which impair generating action names are located closer to Action representations than are lesions which impair generating object names. Thus, the results support the claim that action naming involves interaction with Action representations.

Modality-Specific Impairments of Grammatical Categories.

Given that, typically, object names are nouns and action names are verbs, another way to interpret the results in Figure 9 is in terms of a specialization within Semantics for the representation of nouns versus verbs (see also Damasio & Tranel, 1993; Daniele et al., 1994; Teixeira Ferreira et al., 1997). The learned specialization in the Semantic representations of nouns and verbs in the current model (see Figure 9) provides a basis for explaining the findings of modality-specific impairments to particular grammatical categories. Figure 10 shows the relative performance of the network on naming objects versus actions from Vision following lesions either to incoming Vision-to-Semantics connections or to outgoing Semantics-to-Phonology connections. The results show that both types of lesions can produce double dissociations in naming objects versus actions as a function of the location of lesion. For incoming connections, lesions

Visual Object Naming minus Visual Action Naming

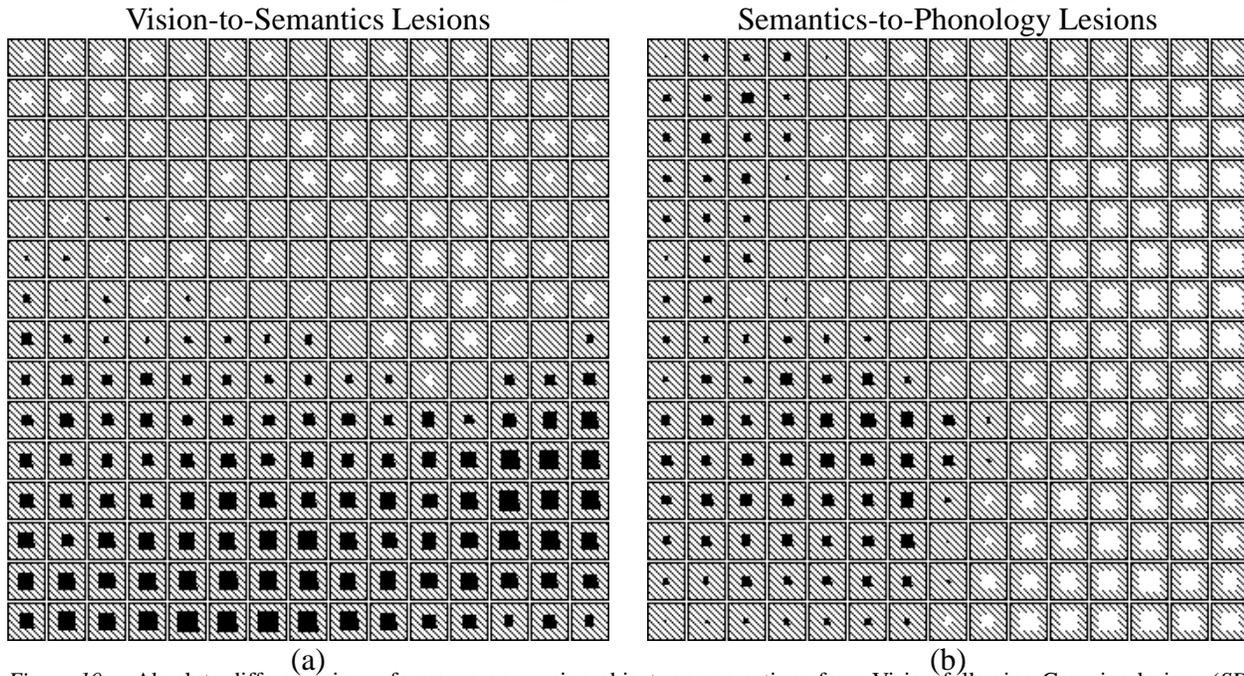


Figure 10. Absolute difference in performance on naming objects versus actions from Vision following Gaussian lesions ($SD = 2.0$) to (a) incoming Vision-to-Semantics connections and (b) outgoing Semantics-to-Phonology connections. White squares indicate better naming of objects than of actions; black squares indicate the reverse.

near Phonology impair object naming more than action naming, whereas lesions near Action produce the opposite pattern (consistent with the results in Figure 10 for lesions directly to Semantics). For outgoing connections, lesions near Vision (biased towards Phonology) impair object naming more than action naming, whereas lesions near Touch (biased towards Action) produce the reverse pattern.

This pattern can be seen more clearly in Figure 11, which shows the mean performance on nouns (object naming) and verbs (action naming) after both incoming and outgoing lesions centered at particular locations within semantics. Also note that, because the lesions are to connections to or from particular modalities, they leave performance involving other modalities essentially unimpaired. Thus, the model exhibits modality-specific impairments for nouns versus verbs following lesions to or from regions of semantics which are partially specialized for objects (nouns) or actions (verbs). In this way, it avoids the rather unpalatable alternative of proposing an anatomic separation of lexical representations by grammatical category within each input and output modality (Caramazza & Hillis, 1991).

Control Simulation: No Topographic Bias

The final issue to be considered is whether it is necessary, in the current context, to introduce a topographic bias on learning in order to account for optic aphasia. Recall from the Introduction that McGuire and Plaut (1997) demonstrated the basic optic aphasic pattern in a non-topographic network very similar to the current one, although the magnitude of

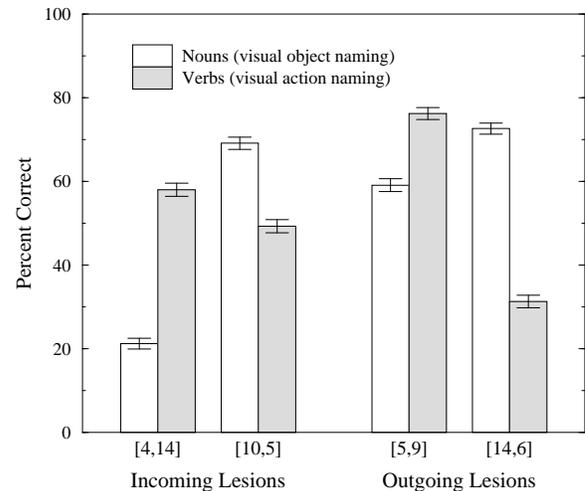


Figure 11. Mean correct performance of the network in generating nouns and verbs from visual input (i.e., visual object and action naming) following lesions to Vision-to-Semantics (Incoming) lesions and Semantics-to-Phonology (Outgoing) lesions centered at specific positions within semantics (indicated by x-y coordinates in brackets, where [0,0] is the upper-left-hand corner of Semantics).

the dissociation between visual object naming versus gesturing was not as large as in many patients (and the relative preservation of visual action naming was not investigated). Moreover, there are a number of related proposals (Hillis & Caramazza, 1995a; Plaut & Shallice, 1993b; Riddoch &

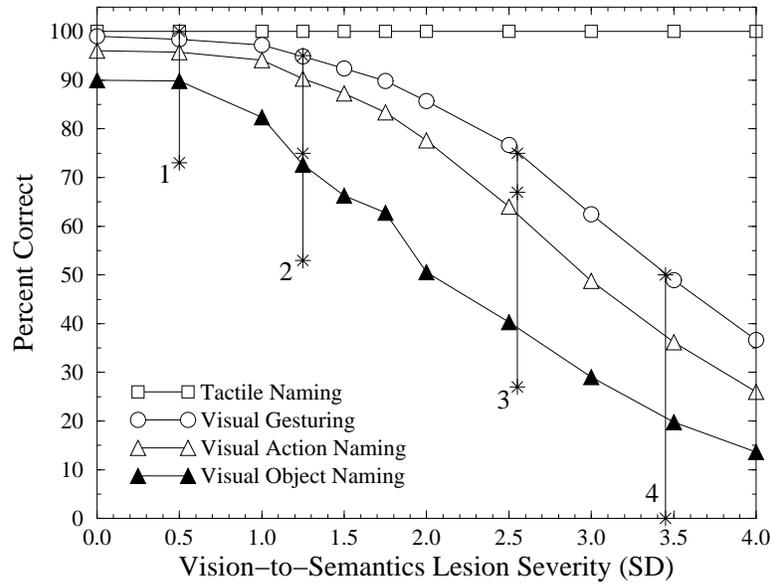


Figure 12. Correct performance of the replication simulation without topographic bias on various tasks as a function of lesion severity, and corresponding levels of performance of four optic aphasic patients, plotted on vertical lines with asterisks and labeled with the corresponding numbers from Table 2. For each line, the top asterisk indicates performance on visual gesturing and the bottom indicates performance on visual object naming. The middle asterisk, when present, indicates performance on visual action naming.

Humphreys, 1987) that make no reference to topographic distinctions within the network (except in the separation of semantics from input and output modalities). It is thus important to determine if a quantitatively adequate account of optic aphasia can be obtained without a topographic bias on learning.

A new simulation was carried out that was an exact replication of the current one, except that error derivatives were not scaled as a function of connection length. The replication was otherwise exact in that it used the identical network architecture, initial random weights, order of training presentations, and instances of lesion as the original simulation. Figure 12 presents the correct performance of the non-topographic network on tactile and visual object naming, and visual action naming and gesturing, as a function of lesion severity analogous to the results in Figure 8. Without the topographic bias, the magnitude of the dissociation in performance on visual object naming versus gesturing is not as large as observed in patients (again indicated by vertical lines with asterisks). These results essentially replicate in an attractor network those that McGuire and Plaut (1997) found in a feedforward network. Thus, at least some pressure for modalities to drive nearby Semantic units more strongly than more distant units seems important for producing a quantitative match to patient performance. In addition, without the topographic bias, there is no pressure for a spatial separation in the Semantic representations for nouns versus verbs, and thus no basis for explaining the modality-specific grammatical category impairments documented by Caramazza and colleagues (Caramazza & Hillis, 1991; Hillis & Caramazza, 1995b; Rapp & Caramazza, 1997, 1998).

General Discussion

The modality-specific aphasias, and optic aphasia in particular, have provided important empirical evidence bearing on the longstanding debate about whether semantic knowledge in the brain is organized as a unitary, amodal system (Caramazza et al., 1990; Riddoch et al., 1988) or whether it is divided into multiple, modality-specific subsystems (Shallice, 1987, 1993; Warrington & McCarthy, 1987; Warrington & Shallice, 1984). To date, neither position has formulated an entirely satisfactory account of the full range of relevant empirical phenomena. The unitary semantics account has yet to be shown capable of exhibiting the magnitude of dissociations in task performance observed among patients—for instance, the advantage of gesturing over naming visually presented objects across the full range of severity of impairment among optic aphasic patients. On the other hand, the multiple-semantics account not only suffers from a lack of parsimony but also has difficulty explaining more fine-grained aspects of patient performance—for example, the relative preservation in optic aphasia of naming visually presented actions compared with objects (e.g., Campbell & Manning, 1996; Teixeira Ferreira et al., 1997).

The current work argues for a theoretical perspective which constitutes a middle ground between the unitary- and multiple-semantics accounts. On this view, semantics consists of a learned, internal representation that develops graded modality-specific functional specialization under the pressure of mediating between multiple input and output modalities. The nature of this graded specialization is shaped by two general factors: differences in the relative degree of systematicity among tasks (i.e., combinations of input-output mappings); and 2) a topographic bias on learning that favors

short connections (Jacobs & Jordan, 1992), leading to the recruitment of regions of semantics that are anatomically close to the relevant input and output modalities.

A computational simulation instantiating these principles was applied to understanding the selective impairment of visual object naming in optic aphasia. A recurrent connectionist network was trained to map either visual or tactile input onto either action (gesturing) or phonological (naming) output via a common, topographically constrained internal (semantic) layer. Damage to connections from vision to regions of semantics near phonology produced the core pattern of performance in optic aphasia: The network was far worse at naming visually presented objects than at gesturing their use or at naming the objects under tactile presentation. This pattern held across the full range of severity of impairments observed among patients. Moreover, like the patients, the network also exhibited relatively spared naming of the actions associated with visually presented objects because action naming benefits from the preserved support of regions of semantics near action representations. Furthermore, the fact that regions of semantics near action representations become partially specialized for generating not only the actions themselves but also their associated names leads to an account of modality-specific dissociations in the processing of nouns versus verbs (e.g., Caramazza & Hillis, 1991; Hillis & Caramazza, 1995b), in terms of damage to the modality-specific incoming or outgoing connections of the semantic regions supporting object versus action naming.

The remainder of the article first considers limitations of the simulation. It then discusses how the current account might address additional characteristics of the performance of optic aphasic patients and how it is related to alternative accounts. Finally, it considers the implications of the current approach for understanding the relation of object and action knowledge, and how it might be extended to account for category-specific semantic impairments.

Limitations of the Simulation

The current simulation is a demonstration of how graded modality-specific specialization in semantic representations can address the core pattern of optic aphasia and some related phenomena. As a demonstration, though, it necessarily incorporates a number of simplifications and approximations, and it is important to consider whether these aspects undermine the relevance of the demonstration. Perhaps the most drastic simplification concerns the representations that were used. These bore no direct relationship to actual representations in the case of Vision, Action, and Touch, and only a very coarse correspondence (through the use of a consonant-vowel-consonant structure) in the case of Phonology.

In considering the implications of these design decisions, it is worth contrasting what might be called *realist* versus *fundamentalist* approaches to modeling (see Kello & Plaut, 2000, for discussion). The former (e.g., Plaut et al., 1996) tries to incorporate into a model as much detail as possible of what is known about the real system in the belief that complex interactions of these factors are necessary to capture the

relevant phenomena. The latter, fundamentalist approach—of which the current work is an example—holds that a model should, as much as possible, embody only those principles that are claimed to account for the relevant phenomenon and should abstract out extraneous details. Often the most effective modeling approach over the long term is to begin with fundamentalist models to isolate the key underlying principles, and then gradually move towards more realist models as the theoretical implications of additional details become understood.

Accordingly, the use of abstract representations in the current work was intended to isolate the implications of relative task systematicity from other idiosyncratic details specific to each modality. Specifically, the mappings among vision, touch and action were all designed to be relatively systematic (i.e., similarity preserving), whereas the mappings from each of these to phonology were designed to be unsystematic. There is no doubt, though, that incorporating more realistic representations for each modality will ultimately turn out to be critical to providing a full, detailed account of optic aphasia in particular, and the organization of semantic knowledge more generally. Even without such details, though, the current work represents an important first step in developing such an account.

Another simplification of the current work concerns the implementation of the learning bias favoring short connections. The literal implementation of this bias is implausible as it starts with full connectivity between groups of units and then influences the relative contribution of these connections. A more realistic implementation would be a distance-dependent constructive algorithm (see, e.g., Quartz & Sejnowski, 1997) that would start with relatively local and sparse connectivity and then grow new connections to nearby units based on the co-occurrence of pre- and post-synaptic activity, perhaps modulated by a global reinforcement signal (Mazzoni, Andersen, & Jordan, 1991). The problem is that such algorithms have yet to be shown to be capable of learning internal representations that support multiple mappings of varying degrees of systematicity, as required of semantics in the current work. On the other hand, error-correcting algorithms, such as back-propagation and more biologically plausible variants (O'Reilly, 1996), are capable of learning such representations but constructive variants of these (e.g., cascade-correlation; Fahlman & Lebiere, 1990) typically yield deeply layered rather than locally constrained connectivity. One possibility to explore in future work would be to employ error-correcting learning within an architecture in which new connections sprout randomly but with a locally biased distribution, and in which these connections are subject to a global cost function that prunes away the unnecessary ones (e.g., Weigand, Rumelhart, & Huberman, 1991). In any case, the main point in the current context is that these alternative implementations of a distance-dependent bias on learning would be expected to have the same the functional consequences as the implementation employed in the current work.

Pattern of Error Responses in Optic Aphasia

To this point, the current work has considered only the relative rates of correct performance of optic aphasic patients on various tasks. However, these patients exhibit a number of other characteristics in their behavior that call for explanation. Perhaps the most important of these is the pattern of error responses that optic aphasic patients make in attempting to name visually presented objects. Typically, they produce semantic errors (e.g., shoe → “hat”), response perseverations (e.g., → “wristwatch”; scissors → “wristwatch”), and unrelated errors (e.g., cat → “house”). Patients also make semantic perseverations—errors which are semantically related to previous responses (e.g., → “newspaper”; case → “two books”; Lhermitte & Beauvois, 1973). Errors sharing both visual and semantic similarity (e.g., orange ⇒ “lemon”) are also common, although purely visual errors (e.g., needle ⇒ “toothpick”) are rare. In fact, the lack of pure visual errors is one of the main ways in which optic aphasia contrasts with visual associative agnosia, where patients make predominantly visual errors (see Farah, 1990; Iorio, Falanga, Fragassi, & Grossi, 1992).

Plaut and Shallice (1993b) accounted for the optic aphasic error pattern using a recurrent connectionist network that was trained to map high-level visual representations to semantics (see Sitton et al., in press, for a related account). Following McClelland and Rumelhart (1985), Plaut and Shallice assumed that the network constantly adjusted its weights based on the pattern of activity generated by each object presentation. For convenience, these short-term weight changes were maintained in a separate set of weights, but the same results would obtain if the changes were applied to the standard long-term weights. Following damage to connections between vision and semantics—corresponding to the same locus of damage as assumed in the current work—the model exhibited semantic errors, mixed semantic-and-visual errors, response and semantic perseverations, and unrelated errors, but very few visual errors, in agreement with empirical observations in optic aphasia. Given that the Plaut and Shallice (1993b) simulation operated according to very similar principles as the current one, introducing short-term weight changes to the latter would also be expected to give rise to the appropriate error pattern under damage. This was not explored directly in the current work because semantic representations are learned rather than specified, making it more difficult to disentangle visual and semantic similarity.

In addition to the single-word error responses just discussed, optic aphasic patients occasionally also produce more extended descriptions of viewed objects. Interestingly, such descriptions tend to become progressively more accurate, often ultimately leading to the generation of the correct name. This type of response—termed “conduite d’approche” by Lhermitte and Beauvois (1973) and “homing in” by Sitton et al. (in press)—was not addressed by Plaut and Shallice (1993b). The following examples are from patient JF (translated from French; Lhermitte & Beauvois, 1973, p. 706).

a mussel: “it looks like two snails, two slugs, it is a shellfish, not an oyster, it should be mussels then.”

a bus: “a wagon...public transport since there is a back door...a stage coach...it would be...no...a city cab...not a cab but a city bus.”

a cup: (the preceding stimulus, a cork screw, had been named correctly) “the cork screw too...there is a porcelain handle...or a fancy cork...there is the reflection...then I should see only a cork unless it could be a cup.”

an aquarium: “a bird-cage, unless it is a pot for flowers, a container, a tank, the four aspects...the walls made of glass or wood...it could be an aquarium if it is made of glass.”

Sitton et al. (in press) interpret this behavior as reflecting a gradual increase in accuracy during the settling process of an attractor system. However, the relatively short duration of the settling process—which is also used to generate individual responses—seems at odds with the temporally extended, multi-response nature of these descriptions. An alternative and perhaps more satisfactory interpretation is that patients are progressively refining their semantic representations by successively processing their verbal responses using the relatively intact auditory-to-semantics mapping (cf. Lhermitte & Beauvois, 1973).

Other Accounts of Optic Aphasia

The current account of optic aphasia is a hybrid of unitary- and multiple-semantics accounts. It acknowledges that a common semantic system maps between multiple modalities, but goes on to claim that this system is not completely homogeneous but incorporates a graded degree of modality-specific functional specialization. Beyond this broad characterization, the current account bears important similarities and differences with a number of other explanations of optic aphasia, which are now discussed in turn.

Optic Aphasia as Mild Visual Associative Agnosia.

The relation between optic aphasia and visual associative agnosia has been a topic of considerable debate, going back to Freund’s (1889) characterization of the distinction. Lissauer’s (1890) original case of visual agnosia exhibited many of the characteristics of optic aphasia, including response and semantic perseverations and semantic errors (see Shallice & Jackson, 1988). Rubens (1979; see also Benke, 1988; De Renzi, Zambolin, & Crisi, 1987) described two cases of associative agnosia who resolved into a pattern more similar to optic aphasia, suggesting that the two syndromes fall on a continuum with optic aphasia simply a mild form of associative agnosia (see also Bauer & Rubens, 1979; Chanoine et al., 1998; De Renzi & Saetti, 1997). In fact, the proposal that optic aphasia arises from partial impairment in the mapping from high-level visual representations to semantics, while sparing both (Geschwind, 1965; Hillis & Caramazza, 1995a; Plaut & Shallice, 1993b; Riddoch et al., 1988) fits the classic definition of visual associative agnosia, whereas the proposal that visual associative agnosia results from damage to visual representations (Farah, 1990; Humphreys & Riddoch, 1987)

fits the classic definition of apperceptive agnosia (Lissauer, 1890; see Farah, 1990 for discussion).

However, in considering this issue, it is important to focus on whether there is empirical evidence for a qualitative distinction between the behavioral patterns commonly referred to as associative agnosia and optic aphasia, and put aside the relatively less interesting question of whether the classic definitions of these labels strictly apply in each case. In this respect, a number of distinctions seem to hold generally. The first is part of the definition of optic aphasia: such patients generally perform much better on tasks demonstrating comprehension of visually presented objects than do associative agnosic patients. The latter are also far more disrupted by manipulations of the visual quality of stimuli (Davidoff & Bleser, 1993). A further distinction was mentioned earlier: optic aphasic patients typically make few if any visual errors on visual object naming, whereas this error type predominates in associative agnosia. Finally, Endo et al. (1996) point out that the visual object naming performance of many optic aphasic patients (see Table 1) is far worse than that of most associative agnosic patients (cited as ranging from 30–80% correct), making it difficult to maintain that optic aphasia is a *mild* form of associative agnosia.

There also appears to be an anatomic distinction in the lesions that give rise to the two patterns of behavior. A review of patients classified as either associative agnosic or optic aphasic (Iorio et al., 1992) revealed that those with posterior unilateral left-hemisphere lesions exhibited symptoms associated with optic aphasia, whereas patients with bilateral lesions typically exhibited associated agnosia. The few cases of associative agnosia that have been reported following unilateral left-hemisphere lesions (e.g., Benke, 1988; De Renzi et al., 1987; Ferro & Santos, 1984; Larrabee et al., 1985; McCarthy & Warrington, 1986) resolved into optic aphasia after about a month (see Endo et al., 1996). Schnider et al. (1994) argued that optic aphasia can also be distinguished from associative agnosia by the presence of damage to the splenium of the corpus collosum, although splenial damage is absent in some optic aphasic patients (e.g., Teixeira Ferreira et al., 1997) and present in some associative agnosic patients (see De Renzi & Saetti, 1997).

Thus, there seems good reason to distinguish visual associative agnosia from optic aphasia—quite apart from the correspondence of these patterns of performance to the classic definitions of these labels—with the former resulting from damage to high-level visual representations and the latter resulting from damage in the mapping from these representations to semantics. Note that, on this account, it is not surprising that some patients (e.g., Chanoine et al., 1998; Iorio et al., 1992; Lissauer, 1890; Schnider et al., 1994) exhibit a mixture of the characteristics of associative agnosia and optic aphasia as, on anatomic grounds, damage to the mapping between vision and semantics might occasionally encroach on the visual representations themselves. Improvement of the latter impairment would then give rise to a resolution of associative agnosia into a pattern more typical of optic aphasia (e.g., Benke, 1988; De Renzi et al., 1987; Rubens, 1979).

Optic Aphasia From Superadditivity of Multiple Impairments.

Farah (1990) provided an alternative answer to how optic aphasics might differ from visual associative agnosics: they have a second mild impairment between semantics and phonology. On this account, a selective impairment in visual object naming arises because of the superadditive effects of utilizing two lesioned pathways in series, as demonstrated by Sitton et al. (in press) in recent computational simulations.

There are, however, a number of problems with this proposal. First, as noted in the Introduction, Sitton et al. (in press) obtained a superadditive effect only for lesions leading to a relatively mild visual object naming impairment, whereas patients span the full range of performance on this task. Second, the proposal of a lesion between semantics and phonology seems difficult to reconcile on anatomic grounds with the occipital locus of damage in optic aphasia. A third problem with the superadditive account is that it provides no explanation for the relative sparing of action naming (Campbell & Manning, 1996; Teixeira Ferreira et al., 1997). Finally, a clear implication of the proposal is that optic aphasics should present with some degree of word finding problems. While such problems have been documented in some cases (Campbell & Manning, 1996; Coslett & Saffran, 1992; Goldenberg & Karlbauer, 1998; Casanova & Roig-Rovira, 1985; Raymer et al., 1997) they are explicitly absent in others (Coslett & Saffran, 1989b; Hillis & Caramazza, 1995a; Iorio et al., 1992; Lhermitte & Beauvois, 1973; Riddoch et al., 1988; Schnider et al., 1994) including some of the most severely impaired patients (Coslett & Saffran, 1989b; Hillis & Caramazza, 1995a). Interestingly, many of the cases reporting word finding problems also exhibit impairments on visual imagery tasks (Campbell & Manning, 1996; Casanova & Roig-Rovira, 1985; Manning, 2000). One possibility, then, is that normal word finding involves some degree of interactive support from high-level visual representations and it is the lack of this support following damage between vision and semantics, rather than a post-semantic lesion, that causes mild word finding problems in some optic aphasic patients.

Optic Aphasia From Hemispheric Disconnection.

Coslett and Saffran (1989b, 1992) proposed an alternative account of optic aphasia that emphasizes the distinct roles that each hemisphere plays in visual, semantic, and language processing (see also De Renzi & Saetti, 1997; Endo et al., 1996; Luzzatti et al., 1998; McCormick & Levine, 1983; Raymer et al., 1997). On this account, the phonological representations required for object naming are localized only in the left hemisphere (LH) and can be activated only by left-hemisphere semantics. The left occipital lesion in optic aphasic patients gives rise to a right homonymous hemianopia, restricting initial visual and semantic processing to the right hemisphere (RH). RH semantics are sufficient to support relatively intact performance on visual comprehension tasks, including gesturing. However, an additional lesion to the splenium of the corpus collosum prevents activation of the LH semantics required for naming.

There are a number of positive aspects to this proposal in

addition to its direct contact with functional neuroanatomy. The putative semantic competence of the RH—sufficient for broad semantic tests but poorer than the LH at detailed distinctions—fits in a general way with a wide range of empirical findings from split-brain and hemispherectomy patients (see Baynes, 1990; Patterson, Vargha-Khadem, & Polkey, 1989), lateralized visual presentation in normal subjects (see Beeman & Chiarello, 1998; Chiarello & Richards, 1992), deep dyslexic patients (e.g., Coltheart, 1980, 2000; Saffran, Bogyo, Schwartz, & Marin, 1980) and pure alexic patients (e.g., Coslett & Saffran, 1989a; Saffran & Coslett, 1998). In fact, De Renzi and Saetti (1997) suggested that optic aphasic patients might differ from associative agnostic patients in having RH semantic systems with greater pre-morbid competence (see also Luzzatti et al., 1998).

The claim of a complete disconnection between RH semantics and LH semantics (and naming) runs into difficulties, however. Hillis and Caramazza (1995a) pointed out that a strict disconnection predicts that visual naming errors in optic aphasia should bear no systematic relationship to the stimulus. Although this was largely true for one severely impaired patient (Coslett & Saffran, 1989b), typically the error responses of optic aphasic patients are semantically related to the stimulus, as discussed earlier. This implies that at least partial information is being communicated from RH visual representations to LH semantic (and phonological) representations. If this is the case, then the performance of optic aphasic patients on semantic tasks does not reflect RH semantics alone but rather a combination of RH and (impoverished) LH semantics. Moreover, if RH semantics is simply a less precise version of LH semantics, it is difficult to see how, from a functional point of view, the account differs from a partial disconnection between visual representations and either a unitary semantic system or one with graded modality specificity, as in the current account.

Indeed, a version of the current simulation that was explicit in its distribution across hemispheres would bear considerable similarity to a graded version of the Coslett and Saffran (1989b, 1992) account. In particular, suppose the phonological representations were located in the LH and the internal, semantic layer were distributed across both hemispheres. Clearly, the semantic units in the LH would be closer to phonology than those in the RH, so the topographic bias would cause the former to play a greater role in generating phonological output. Damage that partially impairs activation of LH semantics from visual representations in both hemispheres would then lead to optic aphasia, as shown in the current work.

Relation of Object and Action Knowledge

A central claim of the current work is that modalities recruit nearby regions of semantics to become partially specialized in representing and processing information in that modality (see also Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995, for functional imaging data supporting this perspective). Thus, for example, semantic regions near action representations become particularly important in relating ac-

tion to information in other modalities. This learned specialization underlies the preserved generation of appropriate actions to visually presented objects following lesions to connections from vision to other semantic regions (e.g., those near phonology). It also explains why generating the names of actions associated with the objects is also relatively preserved following such lesions (Campbell & Manning, 1996; Teixeira Ferreira et al., 1997): action naming, unlike object naming, is supported by the relatively preserved activation of action-specialized regions of semantics. In other words, generating action names from visual input involves interactive support from action representations.

Teixeira Ferreira et al. (1997) offered a closely related account of the preserved action naming of their optic aphasic patient, drawing on Goodale and Milner's (1992) reformulation of the two cortical visual systems (Ungerleider & Mishkin, 1982). On this account, visual object recognition and naming depends primarily on the occipito-temporal "what" pathway (which is damaged in their patient), whereas generating and naming actions to visual input depends on the occipito-parieto-frontal "how" pathway (which is intact). Indeed, there are a broad range of behavioral, neuropsychological, and electrophysiological findings that support a dissociation between object knowledge in temporal cortex and action knowledge in parieto-frontal cortex (see Milner & Goodale, 1995, for review).

On close inspection, however, the evidence for a strict separation between object and action knowledge is less than compelling. For example, the reaching and grasping behavior of visual apperceptive agnostic patient DF (Goodale, Milner, Jakobson, & Carey, 1991) has been interpreted as reflecting the isolated operation of the dorsal "how" pathway (Goodale & Milner, 1992; Milner & Goodale, 1995). While DF exhibits sophisticated sensitivity to general affordances in how visually presented objects should be picked up and positioned for use (e.g., posted through a mail slot), she is unable to demonstrate knowledge of their use or function (A. D. Milner, personal communication, September 2000).

One source of evidence for a mapping from vision to action that is separate from semantic knowledge of objects comes from the error patterns exhibited by normal subjects in generating names or gestures to pictures under a response deadline (Rumiati & Humphreys, 1998). Whereas subjects produced more semantically related errors than pure visual errors in naming, they showed the opposite pattern in gesturing. Rumiati and Humphreys interpreted the higher proportion of visual errors in gesturing to reflect the operation of a direct pathway from vision to action that bypasses semantics.

Another interpretation, however, draws on properties of a distributed connectionist implementation of the mapping from visual object representations to semantics (Plaut & Shallice, 1993b). In that network, the semantic features that are more systematically related to visual form are activated earlier than semantic features that are less systematically related. More generally, as a distributed recurrent network settles in the course of performing a partially systematic mapping, more systematic aspects of the mapping are available earlier than less systematic aspects (e.g., in English word

reading, mapping orthography to phonology before semantics, Kawamoto, 1993, or activating consonants before vowels within phonology Plaut et al., 1996; cf. Berent & Perfetti, 1995). If it is assumed that responding under a deadline emphasizes early semantic activation, and if the aspects of semantics that are systematically related to visual form are also systematically related to action representations (but not to phonology), then gesturing would be expected to produce a higher proportion of visual errors under a deadline than naming.

A final line of evidence for the separation of object and action knowledge from the observation of patients with a selective impairment to one or the other of these. Sirigu, Duhamel, and Poncet (1991) reported a patient with bilateral temporal lobe damage due to herpes encephalitis who was severely impaired at recognizing objects from vision as well as other modalities, but who nonetheless was quite good at describing and demonstrating how to use objects he could not recognize. Similarly, patients with semantic dementia (Snowden, Goulding, & Neary, 1989), a deterioration of conceptual knowledge due to progressive temporal lobe atrophy, often exhibit relatively normal use of objects despite severe impairments in their conceptual knowledge (Buxbaum, Schwartz, & Carew, 1997; Hodges, Patterson, Oxbury, & Funnell, 1992; Lauro-Grotto, Piccini, & Shallice, 1997; Schwartz, Marin, & Saffran, 1979; Snowden, Griffiths, & Neary, 1995). By contrast, Buxbaum, Veramonti, and Schwartz (2000) recently reported two patients with ideomotor apraxia—one with a left occipito-parietal lesion and the other with a left fronto-parietal lesion—who had relatively intact semantic knowledge of the function of made-made objects but had severely impaired knowledge of how they are manipulated. Thus, it appears that knowledge of an objects' function—what it is for—can be doubly dissociated from knowledge of manipulation—how it is used (Buxbaum & Saffran, 2000). A natural interpretation of such a pattern is that object and action knowledge are subserved by separate neural mechanisms (although see Plaut, 1995; Shallice, 1988).

Recent findings by Hodges, Bozeat, Lambon Ralph, Patterson, and Spatt (2000) suggest, however, that the relative preservation of action knowledge in semantic dementia may not be as independent of object knowledge as first thought (see also Hodges, Spatt, & Patterson, 1999). Hodges and colleagues tested nine semantic dementia patients on their conceptual knowledge and their ability to use 20 everyday objects, as well as on their mechanical problem solving ability with novel tools (Goldenberg & Hagmann, 1998). The patients had normal problem solving ability but exhibited impaired object use that correlated strongly with the degree of impairment in their conceptual knowledge of the objects. The few instances in which appropriate object use co-occurred with chance performance on conceptual tasks were largely restricted to objects with strong visual affordances (i.e., a clear relationship between form and use; Gibson, 1979). The results suggest that general action knowledge, such as visual affordances and mechanical problem solving, makes important contributions to generating actions, but that

fully effective object use depends on object-specific conceptual knowledge.

Taken together, the available evidence supports a general anatomic and functional distinction between object and action knowledge but, as is true of the current simulation, the separation is graded rather than categorical.

Extension to Category-Specific Deficits

The modality-specific aphasias are but one source of evidence often interpreted as supporting the multiple-semantics account. Another relevant set of phenomena come from the observation that brain damage can produce selective deficits in knowledge of specific categories of stimuli—most commonly, natural kinds versus artifacts (e.g., De Renzi & Lucchelli, 1994; Farah, McMullen, & Meyer, 1991; Hillis & Caramazza, 1991; Laiacona, Barbarotto, & Capitani, 1993; Lambon Ralph, Howard, Nightingale, & Ellis, 1998; Moss, Tyler, & Jennings, 1997; Moss, Tyler, Durrant-Peatfield, & Bunn, 1998; Sacchett & Humphreys, 1992; Warrington, 1975; Warrington & McCarthy, 1983, 1987, 1994; Warrington & Shallice, 1984). In fact, the impairment to knowledge of natural kinds is even more specific in some patients—restricted, for example, to fruits and vegetables versus other foods (Hart, Berndt, & Caramazza, 1985), to animals versus other natural kinds (Hart & Gordon, 1992), or even to naming animals from auditory versus visual input (McCarthy & Warrington, 1988).

Initially it seemed that this double dissociation indicated that semantics includes anatomically separate subsystems for natural kinds and artifacts (e.g., Warrington, 1975; Warrington & McCarthy, 1983). However, Warrington and Shallice (1984) suggested that it could be more naturally explained if semantics were instead organized by modality, including visual semantics and functional semantics, under the assumption that natural kinds depend most heavily on visual semantics whereas artifacts depend most heavily on functional semantics. This *sensory-functional* account has received considerable empirical and computational support (see Saffran & Schwartz, 1994, for review). However, the account predicts that dissociations in knowledge of natural kinds versus artifacts should pattern in exactly the same way as dissociations in knowledge of the visual versus functional attributes of objects. Contrary to this prediction, Lambon Ralph et al. (1998) reported a patient with a category-specific impairment to natural kinds but no difference in visual versus functional knowledge, and a patient with worse visual compared to functional knowledge but no category specificity in performance. Partly on the basis of these cases and one of their own, Caramazza and Shelton (1998) argued for a return to the proposal that knowledge of natural kinds and artifacts are represented separately in the brain, and attempted to provide an evolutionary justification for this organization.

Recall, however, that the cases reported by Sirigu et al. (1991) and by Buxbaum et al. (2000), taken together, constitute a double dissociation between knowledge of an object's function and knowledge of how to use and manipulate the object. Perhaps the relevant distinction between natural kinds

and artifacts is not sensory versus functional, but conceptual knowledge (which includes both sensory and functional attributes) versus manipulation knowledge. Specifically, natural kinds have a distinct reliance on manipulation knowledge, whereas artifacts depend solely on conceptual knowledge (see also Buxbaum & Saffran, 2000).

Although this account is at an early stage of development, it fits well with the current perspective on semantic organization. Specifically, semantic damage near visual representations should produce selective deficits in natural kinds because, relative to artifacts, these items are more visually confusable (Lloyd-Jones & Humphreys, 1997) and do not have differentiating actions. Moreover, less severe damage should give rise to more selective deficits for stimuli that are the most visually similar in the least consistent tasks (e.g., naming animals; Hart et al., 1985). By contrast, damage near action representations should produce the opposite pattern because artifacts rely more heavily on the distinctiveness of their associated actions. A computational demonstration of the adequacy of this proposal is left to future work.

Conclusions

In 1973, Alan Newell—one of the founders of computational modeling of cognition (e.g., Simon & Newell, 1958)—wrote a paper entitled “You Can’t Play 20 Questions With Nature and Win.” In it, he argued that the simple dichotomies that had dominated the field of cognitive psychology (e.g., “Is attentional selection early or late?”, “Is visual search serial or parallel?”) were unlikely to be resolved in favor of one or the other alternative because the questions themselves belied the richness and complexity of human cognitive processing. Indeed, the fact that each debate had such a long and tortuous history in the field suggested that each side expressed important insights about the domain and a complete account would incorporate aspects of each. Newell argued forcefully for the value of developing explicit computational models in forging an integrative theoretical foundation that could do justice to the richness of human behavior.

In many ways, the question of whether semantics consists of a unitary or multiple systems constitutes a similar dichotomy, with its own tortuous, sometimes acrimonious history (cf. Caramazza et al., 1990; Shallice, 1993). Each side has considerable empirical and theoretical support but also faces considerable challenges. Sufficient evidence and counter-evidence has accumulated on each side that it seems difficult to imagine the issue being decided one way or the other.

The current work articulates a view of semantics that has important similarities with both the unitary- and multiple-semantics perspectives. As in unitary-semantics accounts, a common internal semantic representation mediates the mapping between multiple input and output modalities, and largely the same semantic representation is activated by an object regardless of the modality of presentation. As in multiple-semantics accounts, regions of semantics become partially specialized for some modalities and mappings over others, leading to more selective deficits following damage

than would otherwise be observed.

In the main case in point, optic aphasia, the graded semantic specialization leads to the possibility of a selective impairment in object naming, with relatively spared gesturing and action naming, in response to visual input. It also provides an indirect account of modality-specific impairments to nouns versus verbs (Caramazza & Hillis, 1991; Hillis & Caramazza, 1995b; Rapp & Caramazza, 1998), by association with object versus action knowledge, respectively. Although considerable work remains in extending the account to cover the full range of phenomena relevant to understanding semantic organization, the results to date suggest that a middle ground between the unitary- and multiple-semantics perspectives may be the most fruitful direction to pursue in future work.

References

- Allport, D. A. (1985). Distributed memory, modular systems and dysphasia. In S. K. Newman & R. Epstein (Eds.), *Current perspectives in dysphasia*. Edinburgh: Churchill Livingstone.
- Assal, G., & Regli, F. (1980). Syndrome de disconnection visuo-verbale et visuo-gestuelle. *Revue Neurologique*, *136*, 365-376.
- Bauer, R. M., & Rubens, A. B. (1979). Agnosia. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* (p. 187-242). Oxford: Oxford University Press.
- Baynes, K. A. (1990). Language and reading in the right hemisphere: Highways or byways of the brain? *Journal of Cognitive Neuroscience*, *2*(3), 159-179.
- Beauvois, M.-F. (1982). Optic aphasia: A process of interaction between vision and language. *Proceedings of the Royal Society of London, Series B*, *298*, 35-47.
- Beauvois, M.-F., & Saillant, B. (1985). Optic aphasia for colours and colour agnosia: A distinction between visual and visuo-verbal impairments in the processing of colours. *Cognitive Neuropsychology*, *2*(1), 1-48.
- Beauvois, M.-F., Saillant, B., Meininger, V., & Lhermitte, F. (1978). Bilateral tactile aphasia: A tacto-verbal dysfunction. *Brain*, *101*, 381-401.
- Beeman, M., & Chiarello, C. (Eds.). (1998). *Right hemisphere language comprehension: Perspectives from cognitive neuroscience*. Mahwah, NJ: Erlbaum.
- Benke, T. (1988). Visual agnosia and amnesia from a left unilateral lesion. *European Neurology*, *28*, 236-239.
- Berent, I., & Perfetti, C. A. (1995). A rose is a REEZ: The two-cycles model of phonology assembly in reading english. *Psychological Review*, *102*(1), 146-184.
- Buxbaum, L. J., & Saffran, E. M. (2000). *The semantics of object manipulation and object function: A double dissociation*. (Manuscript submitted for publication)
- Buxbaum, L. J., Schwartz, M. F., & Carew, T. G. (1997). The role of semantic memory in object use. *Cognitive Neuropsychology*, *14*(2), 219-254.
- Buxbaum, L. J., Veramonti, T., & Schwartz, M. F. (2000). Function and manipulation tool knowledge in apraxia: Knowing ‘what for’ but not ‘how’. *Neurocase*, *6*, 83-97.
- Campbell, R., & Manning, L. (1996). Optic aphasia: A case with spared action naming and associated disorders. *Brain and Language*, *53*, 183-221.

- Caramazza, A., & Hillis, A. E. (1991). Lexical organization of nouns and verbs in the brain. *Nature*, *349*, 788-790.
- Caramazza, A., Hillis, A. E., Rapp, B. C., & Romani, C. (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology*, *7*, 161-189.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience*, *10*(1), 1-34.
- Casanova, J. Peña, & Roig-Rovira, T. (1985). Optic aphasia, optic apraxia, and loss of dreaming. *Brain and Language*, *26*, 63-71.
- Chanoine, V., Teixeira Ferreira, C., Demonet, J. F., Nespoulous, J. L., & Poncet, M. (1998). Optic aphasia with pure alexia: A mild form of visual associative agnosia? a case study. *Cortex*, *34*(3), 437-448.
- Chiarello, C., & Richards, L. (1992). Another look at categorical priming in the cerebral hemispheres. *Neuropsychologia*, *30*(4), 381-392.
- Coltheart, M. (1980). Deep dyslexia: A right-hemisphere hypothesis. In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (p. 326-380). London: Routledge & Kegan Paul.
- Coltheart, M. (2000). Deep dyslexia is right-hemisphere reading. *Brain and Language*, *71*, 299-309.
- Coslett, H. B., & Saffran, E. M. (1989a). Evidence for preserved reading in "pure alexia". *Brain*, *112*, 327-359.
- Coslett, H. B., & Saffran, E. M. (1989b). Preserved object recognition and reading comprehension in optic aphasia. *Brain*, *112*, 1091-1110.
- Coslett, H. B., & Saffran, E. M. (1992). Optic aphasia and the right hemisphere: A replication and extension. *Brain and Language*, *43*, 148-161.
- Damasio, A. R., & Tranel, E. (1993). Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Science, USA*, *90*, 4957-4960.
- Daniele, A., Suistolisi, L., Silveri, M. C., Colosimo, C., & Gainotti, G. (1994). Evidence for a possible neuroanatomical basis for lexical processing of nouns and verbs. *Neuropsychologia*, *32*(11), 1325-1341.
- Davidoff, J., & Bleser, R. de. (1993). Optic aphasia: A review of past studies and reappraisal. *Aphasiology*, *7*(2), 135-154.
- Denes, G., & Semenza, C. (1975). Auditory modality-specific anomia: Evidence from a case of pure word deafness. *Cortex*, *11*, 401-411.
- De Renzi, E., Faglioni, P., & Sorgato, P. (1982). Modality specific and supramodal mechanisms of apraxia. *Brain*, *105*, 301-312.
- De Renzi, E., & Lucchelli, F. (1994). Are semantic systems separately represented in the brain? The case of living category impairment. *Cortex*, *30*(1), 3-25.
- De Renzi, E., & Pellegrino, G. di. (1995). Sparing of verbs and preserved, but ineffectual reading in a patient with impaired word production. *Cortex*, *31*(4), 619-636.
- De Renzi, E., & Saetti, M. C. (1997). Associative agnosia and optic aphasia: Qualitative or quantitative difference? *Cortex*, *33*(1), 115-130.
- De Renzi, E., Zambolin, A., & Crisi, G. (1987). The pattern of neuropsychological impairment associated with left posterior cerebral artery infarcts. *Brain*, *110*, 1099-1116.
- Druks, J., & Shallice, T. (1996). The production of verbs in an optic aphasic patient. *Brain and Cognition*, *32*, 246-248.
- Druks, J., & Shallice, T. (2000). Selective preservation of naming from description and the "restricted preverbal message". *Brain and Language*, *27*(2), 100-127.
- Endo, K., Makishita, H., Yanagisawa, N., & Sugishita, M. (1996). Modality specific naming and gesture disturbances: A case with optic aphasia, bilateral tactile aphasia, optic apraxia and tactile apraxia. *Cortex*, *32*(1), 3-28.
- Fahlman, S. E., & Lebiere, C. (1990). The Cascade-Correlation learning architecture. In D. S. Touretzky (Ed.), *Advances in Neural Information Processing Systems 2* (p. 524-534). San Mateo, CA: Morgan Kaufmann.
- Farah, M. J. (1990). *Visual agnosia: Disorders of object recognition and what they tell us about normal vision*. Cambridge, MA: MIT Press.
- Farah, M. J., McMullen, P. A., & Meyer, M. M. (1991). Can recognition of living things be selectively impaired? *Neuropsychologia*, *29*(2), 185-193.
- Ferro, J. M., & Santos, M. E. (1984). Associative visual agnosia: A case study. *Cortex*, *20*, 212-234.
- Freund, D. C. (1889). Ueber optische aphasie und seelenblindheit. *Arch. Psychiat Nervkrankh.*, *20*, 276-297. (Translated by Beaton, A., Davidoff, J., and Erstfeld, U., *Cognitive Neuropsychology*, 1991, *8*, 21-38.)
- Gainotti, G., Silveri, M. C., Daniele, A., & Giustolisi, L. (1995). Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory*, *3*, 247-265.
- Geshwind, N. (1965). Disconnexion syndromes in animals and man. *Brain*, *88*, 237-294.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.
- Gil, R., Pluchon, C., Toullat, G., Michenau, D., Rogez, R., & Levevre, J. P. (1985). Disconnexion visuo-verbale (aphasie optique) pour les objets, les images, les couleurs et les visages avec alexie "abstractive". *Neuropsychologia*, *23*, 333-349.
- Goldenberg, G., & Hagmann, S. (1998). Tool use and mechanical problem solving in apraxia. *Neuropsychologia*, *36*(7), 581-589.
- Goldenberg, G., & Karlbauer, F. (1998). The more you know the less you can tell: Inhibitory effects of visuo-semantic activation on modality specific visual misnaming. *Cortex*, *34*(4), 471-491.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20-25.
- Goodale, M. A., Milner, A. D., Jakobson, L. S., & Carey, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, *349*(6305), 154-156.
- Graham, N. L., Zeman, A., Young, A. W., Patterson, K., & Hodges, J. R. (1999). Dyspraxia in a patient with corticobasal degeneration: The role of visual and tactile inputs to action. *Journal of Neurology, Neurosurgery, and Psychiatry*, *67*, 334-344.
- Hart, J., Berndt, R. S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature*, *316*, 439-440.
- Hart, J., & Gordon, B. (1992). Neural subsystems for object knowledge. *Nature*, *359*, 60-64.
- Hillis, A. E., & Caramazza, A. (1991). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, *114*, 2081-2094.
- Hillis, A. E., & Caramazza, A. (1995a). Cognitive and neural mechanisms underlying visual and semantic processing: Implications from "optic aphasia". *Journal of Cognitive Neuroscience*, *7*(4), 457-478.

- Hillis, A. E., & Caramazza, A. (1995b). The representation of grammatical categories in the brain. *Journal of Cognitive Neuroscience*, 7(3), 457-478.
- Hillis, A. E., Rapp, B., Romani, C., & Caramazza, A. (1990). Selective impairments of semantics in lexical processing. *Cognitive Neuropsychology*, 7, 191-243.
- Hinton, G. E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40, 185-234.
- Hodges, J. R., Bozeat, S., Lambon Ralph, M. A., Patterson, K., & Spatt, J. (2000). The role of conceptual knowledge in object use: Evidence from semantic dementia. *Brain*, 123(9), 1913-1925.
- 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 Science, USA*, 96, 9444-9448.
- Humphreys, G. W., & Riddoch, M. J. (1987). *To see but not to see: A case-study of visual agnosia*. Hillsdale, NJ: Erlbaum.
- Iorio, L., Falanga, A., Fragassi, N. A., & Grossi, D. (1992). Visual associative agnosia and optic aphasia: A single case study and review of the syndromes. *Cortex*, 28, 23-37.
- Jacobs, R. A. (1997). Nature, nurture and the development of functional specializations: A computational approach. *Psychonomic Bulletin & Review*, 4(3), 299-309.
- Jacobs, R. A., & Jordan, M. I. (1992). Computational consequences of a bias toward short connections. *Journal of Cognitive Neuroscience*, 4(4), 323-336.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford, U.K.: Blackwell.
- Kawamoto, A. H. (1993). Nonlinear dynamics in the resolution of lexical ambiguity: A parallel distributed processing approach. *Journal of Memory and Language*, 32, 474-516.
- Kello, C. T., & Plaut, D. C. (2000). *Strategic control over rate of processing in word reading: A computational investigation of the tempo-naming task*. (Manuscript submitted for publication.)
- Laiacona, M., Barbarotto, R., & Capitani, E. (1993). Perceptual and associative knowledge in category specific impairment of semantic memory: A study of two cases. *Cortex*, 29, 727-740.
- Lambon Ralph, M. A., Howard, D., Nightingale, G., & Ellis, A. W. (1998). Are living and non-living category-specific deficits causally linked to impaired perceptual or associative knowledge? Evidence from a category-specific double dissociation. *Neurocase*, 4, 311-338.
- Larrabee, G. J., Levin, H. S., Huff, F. J., Kay, M. C., & Guinto, F. C. (1985). Visual agnosia contrasted with visual verbal disconnection. *Neuropsychologia*, 23(1), 1-12.
- Lauro-Grotto, R., Piccini, C., & Shallice, T. (1997). Modality-specific operations in semantic dementia. *Cortex*, 33(4), 593-622.
- Lhermitte, F., & Beauvois, M.-F. (1973). A visual-speech disconnection syndrome: Report of a case with optic aphasia, agnosic alexia and colour agnosia. *Brain*, 96, 695-714.
- Lissauer, H. (1890). Ein fall von seelenblindheit nebst einem beitrage zur theorie derselben. *Archiv Fur Psychaitrie und Nervenkrankheiten*, 21, 222-270.
- Lloyd-Jones, T. J., & Humphreys, G. W. (1997). Perceptual differentiation as a source of category effects in object processing: Evidence from naming and object decision. *Memory and Cognition*, 25(1), 18-35.
- Luzzatti, C., Rumiati, R. I., & Ghirardi, G. (1998). A functional model of visuo-verbal disconnection and the neuroanatomical constraints of optic aphasia. *Neurocase*, 4, 71-87.
- Manning, L. (2000). Loss of visual imagery and defective recognition of parts of wholes in optic aphasia. *Neurocase*, 6, 111-128.
- Manning, L., & Campbell, R. (1992). Optic aphasia with spared action naming: A description and possible loci of impairment [note]. *Neuropsychology*, 30(6), 587-592.
- Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, 270, 102-105.
- Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(1), 3-23.
- Mazzoni, P., Andersen, R. A., & Jordan, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Science, USA*, 88, 4433-4437.
- McCarthy, R., & Warrington, E. K. (1986). Visual associative agnosia: A clinico-anatomical study of a single case. *Journal of Neurology, Neurosurgery, and Psychiatry*, 49, 1233-1240.
- McCarthy, R., & Warrington, E. K. (1988). Evidence for modality-specific meaning systems in the brain. *Nature*, 334, 428-430.
- McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114(2), 159-188.
- McCormick, G. F., & Levine, D. A. (1983). Visual anomia: A unidirectional disconnection. *Neurology*, 33, 664-666.
- McGuire, S., & Plaut, D. C. (1997). Systematicity and specialization in semantics: A computational account of optic aphasia. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (p. 502-507). Hillsdale, NJ: Erlbaum.
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford, U.K.: Oxford University Press.
- Moss, H. E., Tyler, L. K., Durrant-Peatfield, M., & Bunn, E. M. (1998). "two eyes of a see-through": Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase*, 4(4/5), 291-310.
- Moss, H. E., Tyler, L. K., & Jennings, F. (1997). When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology*, 14(6), 901-950.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading* (p. 264-336). Hillsdale, NJ: Erlbaum.
- Nelson, M. E., & Bower, J. M. (1990). Brain maps and parallel computers. *Trends in Neurosciences*, 13, 403-408.
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers in this symposium. In W. E. Chase (Ed.), *Visual information processing* (p. 283-308). New York: Academic Press.
- O'Reilly, R. C. (1996). Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, 8(5), 895-938.
- Patterson, K., Vargha-Khadem, F., & Polkey, C. E. (1989). Reading with one hemisphere. *Brain*, 112, 39-63.
- Pearlmutter, B. A. (1989). Learning state space trajectories in recurrent neural networks. *Neural Computation*, 1(2), 263-269.

- Pilgrim, E., & Humphreys, G. W. (1991). Impairment of action to visual objects in a case of ideomotor apraxia. *Cognitive Neuropsychology*, 8(6), 459-473.
- Plaut, D. C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, 17(2), 291-321.
- Plaut, D. C. (in press). Systematicity and specialization in semantics. In G. W. Humphreys, D. Heinke, & A. Olson (Eds.), *Connectionist models in cognitive neuroscience: 5th Neural Computation and Psychology Workshop*. New York: Springer-Verlag.
- Plaut, D. C., & Booth, J. R. (in press). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*.
- Plaut, D. C., & Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*, 15, 445-485.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56-115.
- Plaut, D. C., & Shallice, T. (1993a). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10(5), 377-500.
- Plaut, D. C., & Shallice, T. (1993b). Perseverative and semantic influences on visual object naming errors in optic aphasia: A connectionist account. *Journal of Cognitive Neuroscience*, 5(1), 89-117.
- Poeck, K. (1984). Neuropsychological demonstration of splenial interhemispheric disconnection in a case of "optic anomia". *Neuropsychologia*, 22(6), 707-713.
- Quartz, S., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*, 20(4), 537-596.
- Rapcsak, S. Z., Rothi, L. J. G., & Heilman, K. M. (1987). Phonological alexia with optic and tactile anomia: A neuropsychological and anatomical study. *Brain and Language*, 31, 109-121.
- Rapp, B., & Caramazza, A. (1997). The modality-specific organization of grammatical categories: Evidence from impaired spoken and written sentence production. *Brain and Language*, 56(2), 248-286.
- Rapp, B., & Caramazza, A. (1998). A case of selective difficulty in writing verbs. *Neurocase*, 4, 127-139.
- Ratcliff, G., & Newcombe, F. A. (1982). Object recognition: Some deductions from the clinical evidence. In A. W. Ellis (Ed.), *Normality and pathology in cognitive functions* (p. 147-171). New York: Academic Press.
- Raymer, A. M., Greenwald, L., Richardson, M., Rothi, L. J. G., & Heilman, K. M. (1997). The right hemisphere and optic aphasia/optic apraxia. *Neurocase*, 3, 173-183.
- Riddoch, M. J. (1999). Optic aphasia: A review of some classic cases. In G. W. Humphreys (Ed.), *Case studies in the neuropsychology of vision* (p. 133-160). Hove, England: Psychology Press/Taylor & Francis.
- Riddoch, M. J., & Humphreys, G. W. (1987). Visual object processing in optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology*, 4(2), 131-185.
- Riddoch, M. J., Humphreys, G. W., Coltheart, M., & Funnell, E. (1988). Semantic systems or system? Neuropsychological evidence re-examined. *Cognitive Neuropsychology*, 5(1), 3-25.
- Riddoch, M. J., Humphreys, G. W., & Price, C. J. (1989). Routes to action: Evidence from apraxia. *Cognitive Neuropsychology*, 6(5), 437-454.
- Rogers, T. T., Lambon Ralph, M., Patterson, K., McClelland, J. L., & Hodges, J. R. (2000). *The structure and deterioration of semantic memory: A computational and neuropsychological investigation*. (Manuscript submitted for publication.)
- Rogers, T. T., & Plaut, D. C. (in press). Connectionist perspectives on category-specific deficits. In E. Forde & G. W. Humphreys (Eds.), *Category-specificity in brain and mind*. New York: Psychology Press.
- Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
- Rubens, A. R. (1979). Agnosia. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* (p. 233-267). Oxford: Oxford University Press.
- Rumiati, R. I., & Humphreys, G. W. (1998). Recognition by action: Dissociating visual and semantic routes to action in normal observers. *Journal of Experimental Psychology: Human Perception and Performance*, 24(2), 631-647.
- Sacchett, C., & Humphreys, G. W. (1992). Calling a squirrel a squirrel but a canoe a wigwam: A category-specific deficit for artefactual objects and body parts. *Cognitive Neuropsychology*, 9(1), 73-86.
- Saffran, E. M., Bogyo, L. C., Schwartz, M. F., & Marin, O. S. M. (1980). Does deep dyslexia reflect right-hemisphere reading? In M. Coltheart, K. Patterson, & J. C. Marshall (Eds.), *Deep dyslexia* (p. 381-406). London: Routledge & Kegan Paul.
- Saffran, E. M., & Coslett, H. B. (1998). Implicit versus letter-by-letter reading in pure alexia: A tale of two systems. *Cognitive Neuropsychology*, 15, 141-166.
- Saffran, E. M., & Schwartz, M. F. (1994). Of cabbages and things: Semantic memory from a neuropsychological perspective—A tutorial review. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance XV: Conscious and nonconscious aspects of cognitive processing* (p. 507-536). Hillsdale, NJ: Erlbaum.
- Schnider, A., Benson, D. F., & Scharre, D. W. (1994). Visual agnosia and optic aphasia: Are they anatomically distinct? *Cortex*, 30, 445-457.
- Schnider, A., Hanlon, R. E., Alexander, D. N., & Benson, D. F. (1997). Ideomotor apraxia: Behavioral dimensions and neuroanatomical basis. *Brain and Language*, 58(1), 125-136.
- Schwartz, M. F., Marin, O. S. M., & Saffran, E. M. (1979). Dissociations of language function in dementia: A case study. *Brain and Language*, 7, 277-306.
- Shallice, T. (1987). Impairments of semantic processing: Multiple dissociations. In M. Coltheart, G. Sartori, & R. Job (Eds.), *The cognitive neuropsychology of language* (p. 111-128). Hillsdale, NJ: Erlbaum.
- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
- Shallice, T. (1993). Multiple semantics: Whose confusions? *Cognitive Neuropsychology*, 10(3), 251-261.
- Shallice, T., & Jackson, M. (1988). Lissauer on agnosia. *Cognitive Neuropsychology*, 5(2), 153-192.
- Simon, H. A., & Newell, A. (1958). Heuristic problem solving: The next advance in operations research. *Operations Research*, 6, 1-10.

- Sirigu, A., Duhamel, J., & Poncet, M. (1991). The role of sensorimotor experience in object recognition: A case of multimodal agnosia. *Brain*, *114*, 727-741.
- Sitton, M., Mozer, M. C., & Farah, M. J. (in press). Superadditive effects of multiple lesions in a connectionist architecture: Implications for the neuropsychology of optic aphasia. *Psychological Review*.
- Snowden, J., Griffiths, H., & Neary, D. (1995). Autobiographical experience and word meaning. *Memory*, *3*, 225-246.
- Snowden, J. S., Goulding, P. J., & Neary, D. (1989). Semantic dementia: A form of circumscribed cerebral atrophy. *Behavioral Neurology*, *2*, 167-182.
- Teixeira Ferreira, C., Guisano, B., Ceccaldi, M., & Poncet, M. (1997). Optic aphasia: Evidence of the contribution of different neural systems to object and action naming. *Cortex*, *33*, 499-513.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (p. 549-586). Cambridge, MA: MIT Press.
- Warrington, E. K. (1975). The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology*, *27*, 635-657.
- Warrington, E. K. (1981). Concrete word dyslexia. *British Journal of Psychology*, *72*, 175-196.
- Warrington, E. K., & McCarthy, R. (1983). Category-specific access dysphasia. *Brain*, *106*, 859-878.
- Warrington, E. K., & McCarthy, R. (1987). Categories of knowledge: Further fractionation and an attempted integration. *Brain*, *110*, 1273-1296.
- Warrington, E. K., & McCarthy, R. A. (1994). Multiple meaning systems in the brain: A case for visual semantics. *Neuropsychologia*, *32*(12), 1465-1473.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*, 829-853.
- Weigand, A. S., Rumelhart, D. E., & Huberman, B. A. (1991). Generalization by weight-elimination with application to forecasting. In R. P. Lippmann, J. E. Moody, & D. S. Touretzky (Eds.), *Advances in Neural Information Processing Systems 3* (p. 875-882). San Mateo, CA: Morgan Kaufmann.
- Zingeser, L. B., & Berndt, R. S. (1990). Retrieval of nouns and verbs in agrammatism and anomia. *Brain and Language*, *39*(1), 14-32.