

# GRADED MODALITY-SPECIFIC SPECIALISATION IN SEMANTICS: A COMPUTATIONAL ACCOUNT OF OPTIC APHASIA

---

David C. Plaut

*Carnegie Mellon University, Pittsburgh, USA*

A long-standing debate regarding the representation of semantic knowledge is whether such knowledge is represented in a single, amodal system or whether it is organised into multiple subsystems based on modality of input or type of information. The current paper presents a distributed connectionist model of semantics that constitutes a middle ground between these unitary- versus multiple-semantics accounts. In the model, 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 favours short connections, leading to a graded degree of modality-specific functional specialisation within semantics. The model is applied to the specific empirical phenomena of optic aphasia—a neuropsychological disorder in which patients exhibit a selective deficit in naming visually presented objects that is not attributable to more generalised impairments in object recognition (visual agnosia) or naming (anomia). As a result of the topographic bias in the model, 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 ability of the model to account for the pattern of performance observed in optic aphasia across the full range of severity of impairment provides support for the claim that semantic representations exhibit graded functional specialisation rather than being entirely amodal or modality-specific.

## INTRODUCTION

A central issue in the study of language and cognition concerns the organisation 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, 1995; Hillis, Rapp, Romani, & Caramazza, 1990; Riddoch, Humphreys, Coltheart, & Funnell, 1988): The meanings of objects are stored in a central,

---

Requests for reprints should be addressed to Dr David Plaut, Department of Psychology, Carnegie Mellon University, Pittsburgh PA 15213–3890, USA (Email: [plaut@cmu.edu](mailto:plaut@cmu.edu)).

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 customised 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 (1999). I thank Marlene Behrmann, Jay McClelland, Tim Shallice, and the CMU PDP research group for helpful comments and discussion.

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 behaviour in various output modalities (e.g., physical action, writing, speaking). Such an organisation is parsimonious in that it allows knowledge derived from one modality to generalise automatically to others.

There are, however, certain empirical findings that seem problematic for a unitary-semantics account, at least in its most straightforward form. A number of these findings 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. For example, optic aphasic patients exhibit a selective impairment in naming visually presented objects, typically due to a lesion to left medial occipital cortex and the underlying white

matter (see Davidoff & De Bleser, 1993; Riddoch, 1999, for reviews). Critically, the visual naming impairment is reducible neither to visual agnosia nor to a more general anomia. Agnosia is ruled out because the patients can demonstrate that they recognise the objects they cannot name—for example, by gesturing their use appropriately.<sup>1</sup> 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, 2000) 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 severe).

**Table 1.** Percentage correct performance of optic aphasic patients on various tasks

Study	Task				
	Visual naming	Visual gesturing	Tactile naming	Naming to spoken definition	Naming action from object
a. Lhermitte and Beauvois (1973)	73	100	91	96	—
b. Gil et al. (1985)	64	100	—	100	—
c. Teixeira Ferreira et al. (1997)	53	95	81	78	75
d. Schnider et al. (1994)	50	58 <sup>b</sup>	75	100	83 <sup>a</sup>
e. Riddoch and Humphreys (1987)	46	75	75	100	—
f. Campbell and Manning (1996)	40	75	92	100	67
g. Manning (2000)	38	83	100	81	—
h. De Renzi and Saetti (1997)	37	42 <sup>b</sup>	100	97	50
i. Luzzatti et al. (1998)	23	79	77	92	—
j. Coslett and Saffran (1992)	21	100	68	68	—
k. Raymer et al. (1997)	15	46 <sup>b</sup>	45	63	65 <sup>a</sup>
l. Hillis and Caramazza (1995)	10	30 <sup>b</sup>	94	95	—
m. Poeck (1984)	8	75	35	90	—
n. Coslett and Saffran (1989b)	0	50	92	73	—

*Naming action from object* refers to naming the action associated with a visually presented object.

— indicates that a task was not tested or that data are unavailable.

<sup>a</sup> Generation of object name from viewing a gesture of its use.

<sup>b</sup> Relatively preserved visual comprehension demonstrated in semantic matching tasks.

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

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; Teixeira Ferreira, Guisano, Ceccaldi, & Poncet, 1997; Zingeser & Berndt, 1990) or when naming objects from a pantomime or demonstration of their use (Campbell & Manning, 1996; Goldenberg & Karlbauer, 1998; Raymer, Greenwald, Richardson, Rothi, & Heilman, 1997; Schnider, Benson, & Scharre, 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.<sup>2</sup> Damage prior to or within semantics 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 (and would not lead to relatively preserved action naming).

Based on these and other considerations, some researchers (e.g., Beauvois, 1982; Lauro-Grotto, Piccini, & Shallice, 1997; Luzzatti, Rumiati, & Ghirardi, 1998; Shallice, 1987, 1993; Warrington, 1975; 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 (e.g., visual, auditory, tactile) and/or type of information (e.g., visual vs. verbal, perceptual vs. functional).<sup>3</sup> A version of 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 nonvisual input could access the verbal semantic information necessary for naming (Beauvois, 1982; Beauvois & Saillant, 1985).

There are, however, a number of reasons why the multiple-semantics account is less than satisfactory. From an empirical perspective, a disconnection between visual and verbal semantics would appear to be inconsistent with the relative sparing in optic aphasia 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. From a theoretical perspective, strict modality- or domain-specific subdivisions within semantics are considered by many researchers to be unparsimonious (Riddoch et al., 1988) if not theoretically incoherent (Caramazza et al., 1990). In particular, Caramazza and colleagues argued that

<sup>2</sup> Some researchers (Davidoff & De Bleser, 1993; Rapcsak, Rothi, & Heilman, 1987; Ratcliff & Newcombe, 1982) have suggested that optic aphasia results from impairment to a nonsemantic 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.

<sup>3</sup> Proposals that semantics is organised by modality and/or by type of information should be distinguished from proposals that there are separate representations for different categories of entities, such as natural kinds versus artefacts (e.g., Caramazza & Shelton, 1998; Shelton & Caramazza, 1999) or concrete versus abstract concepts (e.g., Warrington, 1981). In the first set of proposals, the complete representation of a given entity is distributed across multiple subsystems, whereas in the second it is restricted to a single subsystem.

proponents of the multiple-semantics account have conflated independent distinctions in their characterisation of various semantic subsystems, regarding the *content* of what is represented (e.g., visual/perceptual vs. functional), the *format* in which it is represented (e.g., visual/pictorial vs. symbolic/propositional; Paivio, 1971), and the *context* in which it is acquired (e.g., from visual vs. spoken input). A critical case in this regard is the specification of verbal semantics, in that “verbal” is not a modality but rather some poorly unspecified combination of content, format, and context. Shallice (1993) responded to these criticisms by defending the traditional information-processing approach of specifying a set of components and their individual functions within the overall system prior to giving a precise characterisation of their internal representations and processes. He also argued that Caramazza et al.’s (1990) alternative, unitary-semantics account—the Organised Unitary Content Hypothesis (OUCH)—faced a number of empirical challenges and that the necessary modifications would render it equivalent to a multiple-semantics account (but see Rapp, Hillis, & Caramazza, 1993, for counterarguments).

Other researchers have attempted to articulate more specific versions of a multiple-semantics account of optic aphasia. For example, Coslett and Saffran (1989b, 1992; see also Endo, Makishita, Yanagisawa, & Sugishita, 1996; Luzzatti et al., 1998; McCormick & Levine, 1983) have proposed a theory in which semantics is divided not by modality but by hemisphere. They assume that the left and right hemispheres represent 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 callosum) 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, 1995; 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, 1995). 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 stringent 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, 1995; Plaut & Shallice, 1993b; Riddoch & Humphreys, 1987). 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, 1995; 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, nonvisual naming) are relatively unimpaired, but visual naming—which requires both damaged pathways—is disproportionately impaired due to the superadditive effects of the two lesions. Although a standard connectionist implementation of a vision-semantics-phonology pathway (Plaut & Shallice, 1993a) failed to exhibit the proposed superadditivity, Sitton et al. (2000) have recently provided computational support for the proposal using a system composed of linked modules exhibiting attractor dynamics. Following a combination of vision-to-semantics and semantics-to-phonology lesions, the model did, in fact, exhibit a superadditive impairment on visual naming relative to visual gesturing and nonvisual 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-semantics account of optic aphasia.

The current work attempts to articulate and support a theory of semantic organisation that constitutes a middle ground between the unitary- and multiple-semantics accounts. It adopts the perspective that semantics is a learned, internal representation that develops under the pressure of performing a variety of tasks involving various input and output modalities (see also Caramazza et al., 1990; Rogers et al., 2002; Rogers & Plaut, 2002). Instead of casting semantics as entirely amodal or entirely modality-specific, the semantic system is claimed to have a graded degree of functional specialisation that is influenced both by modality of input and by the nature of the information being represented.

Specifically, in conjunction with differences in task systematicity, the current work investigates an additional pressure for functional specialization: a topographic bias on learning favouring short connections (Jacobs & Jordan, 1992). The basic idea is that brain organisation 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 m radius axons, accommodating the axon volume would require a sphere over 20 km in diameter (Nelson & Bower, 1990). Clearly, connectivity must be as local as possible.

One way of instantiating a bias favouring short connections within a connectionist network is to assign spatial locations to units and to scale the learning rate on each connection as a function of the proximity of 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 favouring short connections can induce varying degrees of functional specialisation 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 relevant 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. Localised damage to semantic representations (or to their incoming or outgoing connections) should, thus, give rise to modality-specific impairments on tasks that require semantics, particularly those that are unsystematic.

Thus, 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 specialised for naming. The relative sparing of naming visual actions in optic aphasia results from the preserved support of action representations in generating action names.

To be clear, the notion of graded functional specialisation within semantics is not novel to the current work; it derives from a perspective first articulated by Allport (1985) and Warrington and McCarthy (1987) and later elaborated by 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 specialised 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

### 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, organised in a  $15 \times 15$  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 later.



**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 full connectivity between the indicated unit groups; bidirectional arrows indicate two separate projections. The Task units are connected to all noninput groups. The positions of units in the figure correspond to their functional positions with respect to calculating connection lengths (as Euclidean distance).

Including bias connections for noninput 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 localised 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 to be 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 discretisation of  $\tau$ ), 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  $\tau$  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

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

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

### 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 equal to zero. Each prototype was then used to generate 20 exemplars. In generating an exemplar, each feature of the prototype had a probability of .1 of changing its original value (from zero to one or vice versa), with the additional constraint that any two exemplars had to differ by at least two features. The result of 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 five 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 .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.

The high degree of visual-tactile systematicity is straightforward to justify given that they both

depend directly on the same three-dimensional structure of objects (see Amedi, Malach, Hendler, Peled, & Zohary, 2001; James et al., 2002, for recent relevant functional imaging results). However, it could be argued that the degree of systematicity of each of these domains with action, although substantial, is somewhat reduced. Following the main simulation, a control simulation is presented in which the probability that category coordinates in one modality were also category coordinates in another remained .8 for Vision-Touch but was reduced to .6 for Vision-Action and Touch-Action.

Phonological representations were generated to form consonant-vowel-consonant strings over 20 features in three slots. The first slot of seven features coded the onset consonant, the next slot of six features coded the vowel, and the last slot of seven features coded the final consonant. Each of 16 possible consonants and 5 possible vowels were represented by a particular pattern of activity in the relevant slot (with two active features). 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* task. 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*). 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.

The motivation for activating both Action and Phonology in the action task is that the network is generating everything it knows about the action associated with the presented object. 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. However, to test whether the current findings depend on this assumption, the control simulation with reduced action-related systematicity also used separate training of gesturing and action naming.

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 (five 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 non-normalised Gaussian function ( $SD = 10$ ) of the inverse of its length.<sup>4</sup> This produces a scaling of near 1.0 for the shortest connections, and near .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 relates to topographic influences on executive control.

Once the derivatives were scaled, the weights were updated according to

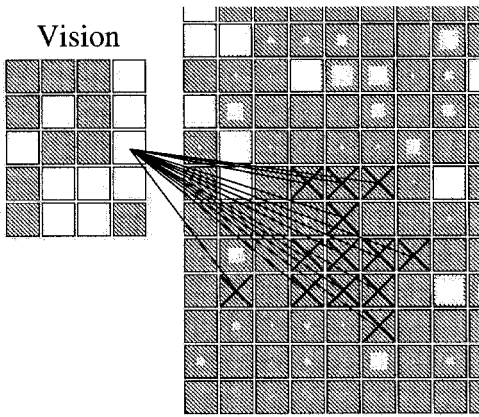
$$\Delta w_{ij}^{[t]} = \epsilon \left( \frac{\partial E}{\partial w_{ij}} = \alpha \Delta w_{ij}^{[t-1]} - \delta w_{ij}^{[t-1]} \right),$$

where  $\epsilon = .01$  is the learning rate,  $\alpha = .9$  is the momentum (i.e., proportion of previous weight change added into the current weight change), and  $\delta = .00005$  is the weight decay (i.e., the proportion of the current weight value that is subtracted from the weight, moving it toward zero). Following this, unit states were reinitialised, 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 .5; 96% (672/700) were within .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 unit (or to their incoming or outgoing connections) were defined to have a centre located at a particular position within Semantics, and to fall off in severity with increasing distance from this centre. 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 non-normalised Gaussian function of the position of the receiving

<sup>4</sup> 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.



**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.

Semantic unit relative to the centre of the lesion. Thus, connections to the unit at the centre 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 extent 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

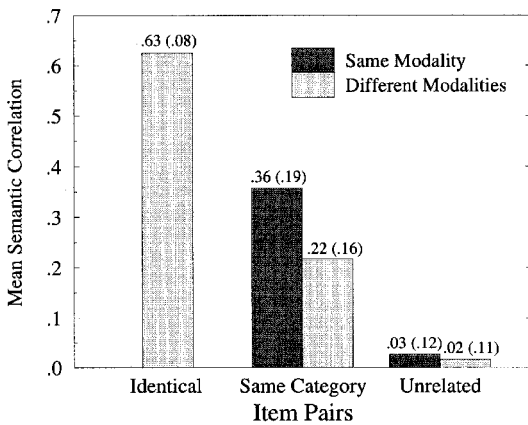
### *Semantic similarity*

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 behavioural 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 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.<sup>5</sup> 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

<sup>5</sup> 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 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 ( $n = 100$  for Identity, 1900 for same-category conditions, and 8000 for different-category conditions).*

some sensitivity to the modality in which the object was presented.

Numerous connectionist simulations have shown that the degree of similarity in the patterns generated by two stimuli is closely related to the magnitude of facilitation in settling time to one of the stimuli when preceded by the other (see, e.g., Cree, McRae, & McNorgan, 1999; Masson, 1995; McRae, De Sa, & Seidenberg, 1997; Plaut, 1995b; Plaut & Booth, 2000; Plaut & Gonnerman, 2000; Sharkey & Sharkey, 1992). The similarity of two patterns of activity corresponds to their distance when represented as points in multi-dimensional “state” space (with a dimension for each unit in the network). When the system starts from the representation of the first stimulus (prime), it has less distance to travel in state space—and, hence, requires less time—to reach the representation of the second stimulus (target) when the two representations are similar compared with when they are dissimilar. Thus, the relative similarities of hidden representations in the current network—strongest

cross-modality identity similarity, and greater within- versus cross-modality category coordinate similarity (compared with unrelated primes)—can be interpreted as corresponding to the basic pattern of results from studies of repetition and semantic priming (see, e.g., Neely, 1991).<sup>6</sup> 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 specialisation*

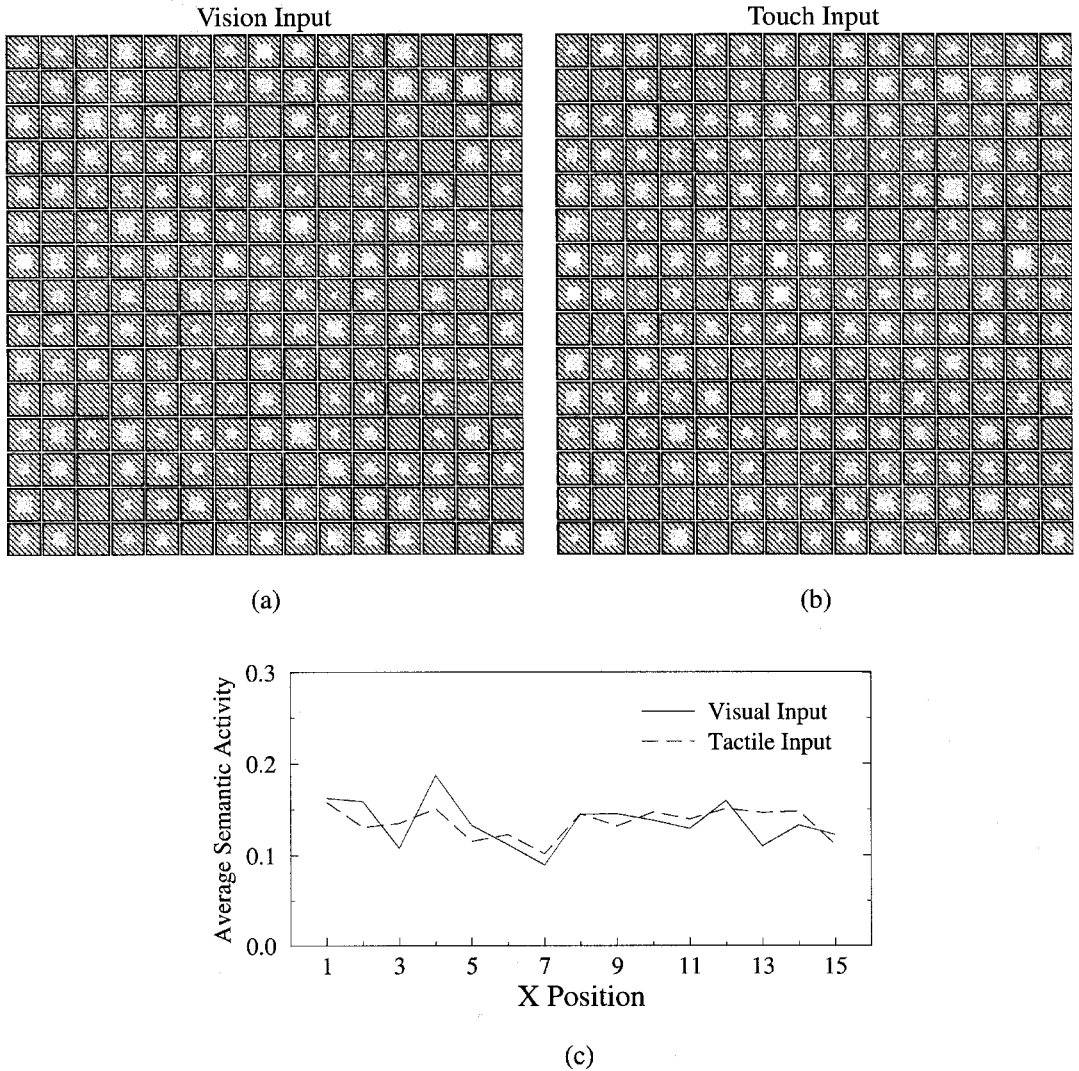
A central claim of the current work is that semantic representations exhibit a graded degree of modality specificity. One question, then, is whether learning has, in fact, produced internal representations with this structure. More specifically, it is important to demonstrate that the differences in 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 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 specialisation in the system more directly.

<sup>6</sup> The reason that such priming is not demonstrated directly in the current simulation is simply because the network was not trained to process objects when starting from the representation produced by the previous object; for convenience, the network always started from the same neutral state.

To do this, the activation of each Semantic unit was calculated when the network was presented with each object in each input modality and task. Figure 4a shows the activations for Vision input and 4b shows the activation for Touch input, averaged over object and task. 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 unit activations with modality of input (Vision, Touch) as a within-unit factor and




**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.

horizontal position (1–15) as a between-unit factor (and treating object, task, and vertical position as repeated measures), there was no reliable effect of the horizontal position of the unit, the modality of presentation, or the interaction of these factors ( $p$ s > .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 semantic 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 relatively 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 specialisation by modality in the Semantic activations of the network is surprising in light of the topographic bias on learning that favoured short over long connections. In fact, as would be

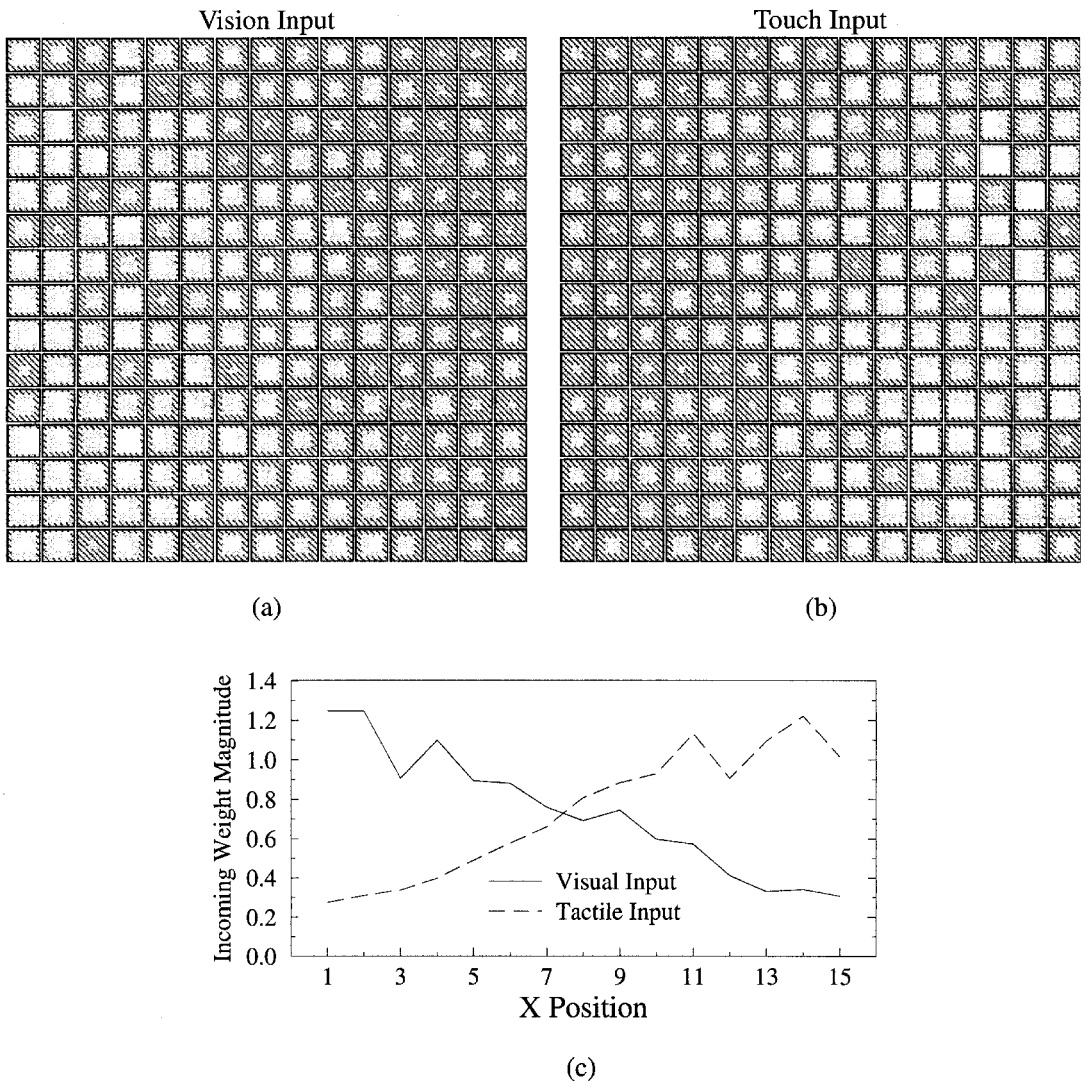


**Figure 5.** Mean variance in hidden unit activations produced by the network when presented with each object in each modality, as a function of the horizontal position of the unit.

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. (The outgoing weights from Semantics to Action or Phonology have exactly the same characteristics, but as a function of vertical distance.) Note that the relative magnitudes of the largest to smallest weights is about 10 to 1, 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. Figure 7 illustrates the time course of the network's settling process in response to a particular visual input. The stronger weights from Vision to nearby Semantic units cause these units initially to become more active 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 Vision. 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 modality of input.

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 repre-

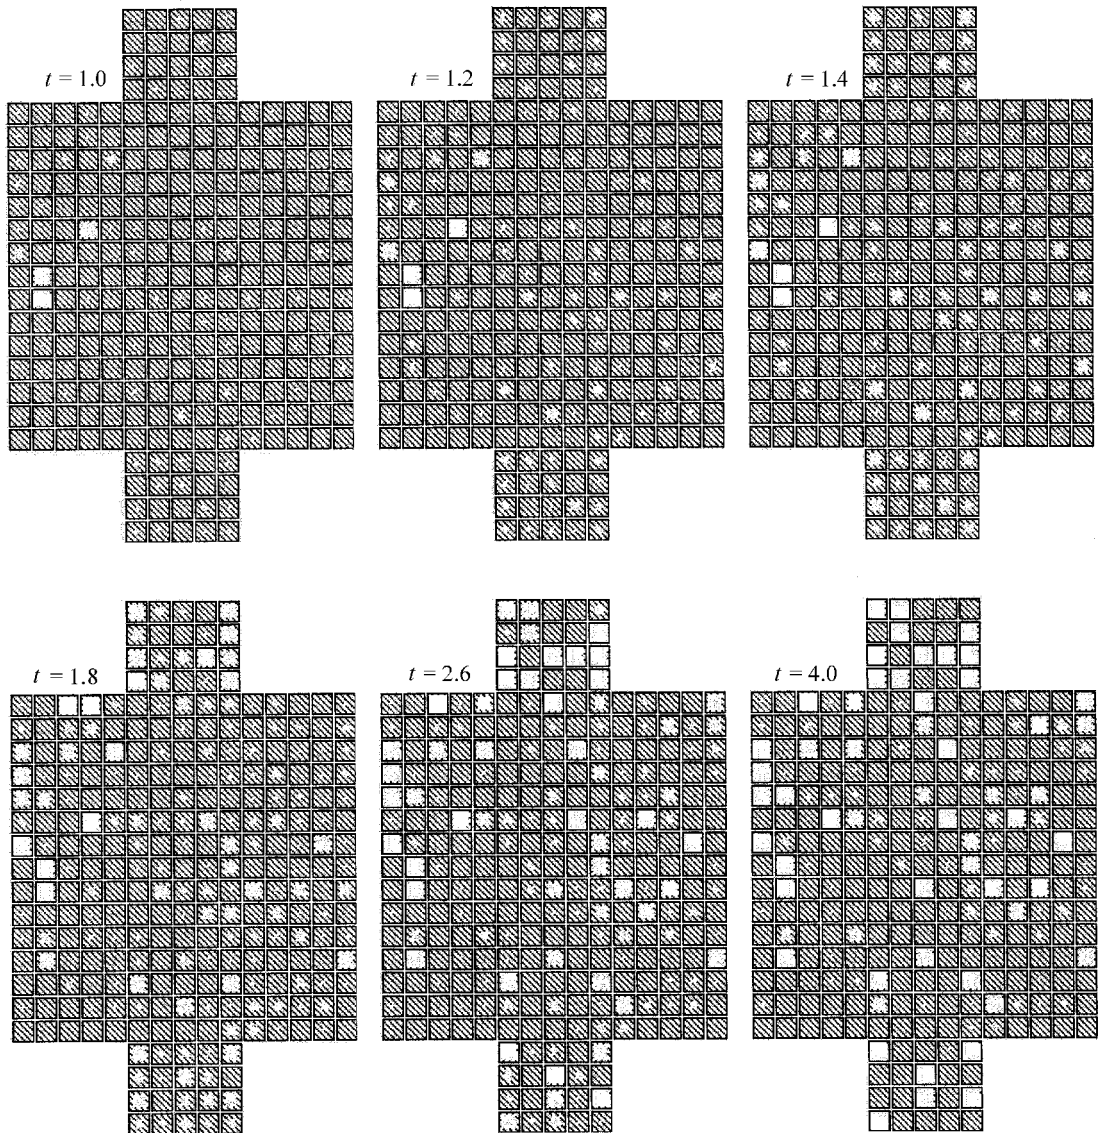


**Figure 6.** Mean absolute magnitude of incoming weights on connections to Semantic units from (a) Vision units or (b) Touch units, as well as (c) the mean weight magnitudes plotted as a function of the horizontal position of the Semantic unit (i.e., also averaged over vertical position). In (a) and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

sentations 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 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



**Figure 7.** Unit activations in the network at six points in time (labelled by  $t$ ) during the settling process in response to the visual input shown in Figure 1 (where activations are depicted in the same manner as in that figure). Each display shows activations for Semantics (middle  $15 \times 15$  block of units) with Action above and Phonology below (each as smaller  $5 \times 4$  blocks); Vision and Touch units are omitted because their activations remain fixed during settling.

partially specialised for generating phonological output. The more specific assumption is that such graded specialisation 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 likeli-

hood of removing a given connection was a Gaussian function ( $SD = 1.5$ ) of the location of the Semantic unit from the specified centre of the lesion (see Figure 2). To avoid sampling artefacts, 10 instances of lesion centred 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 .5.

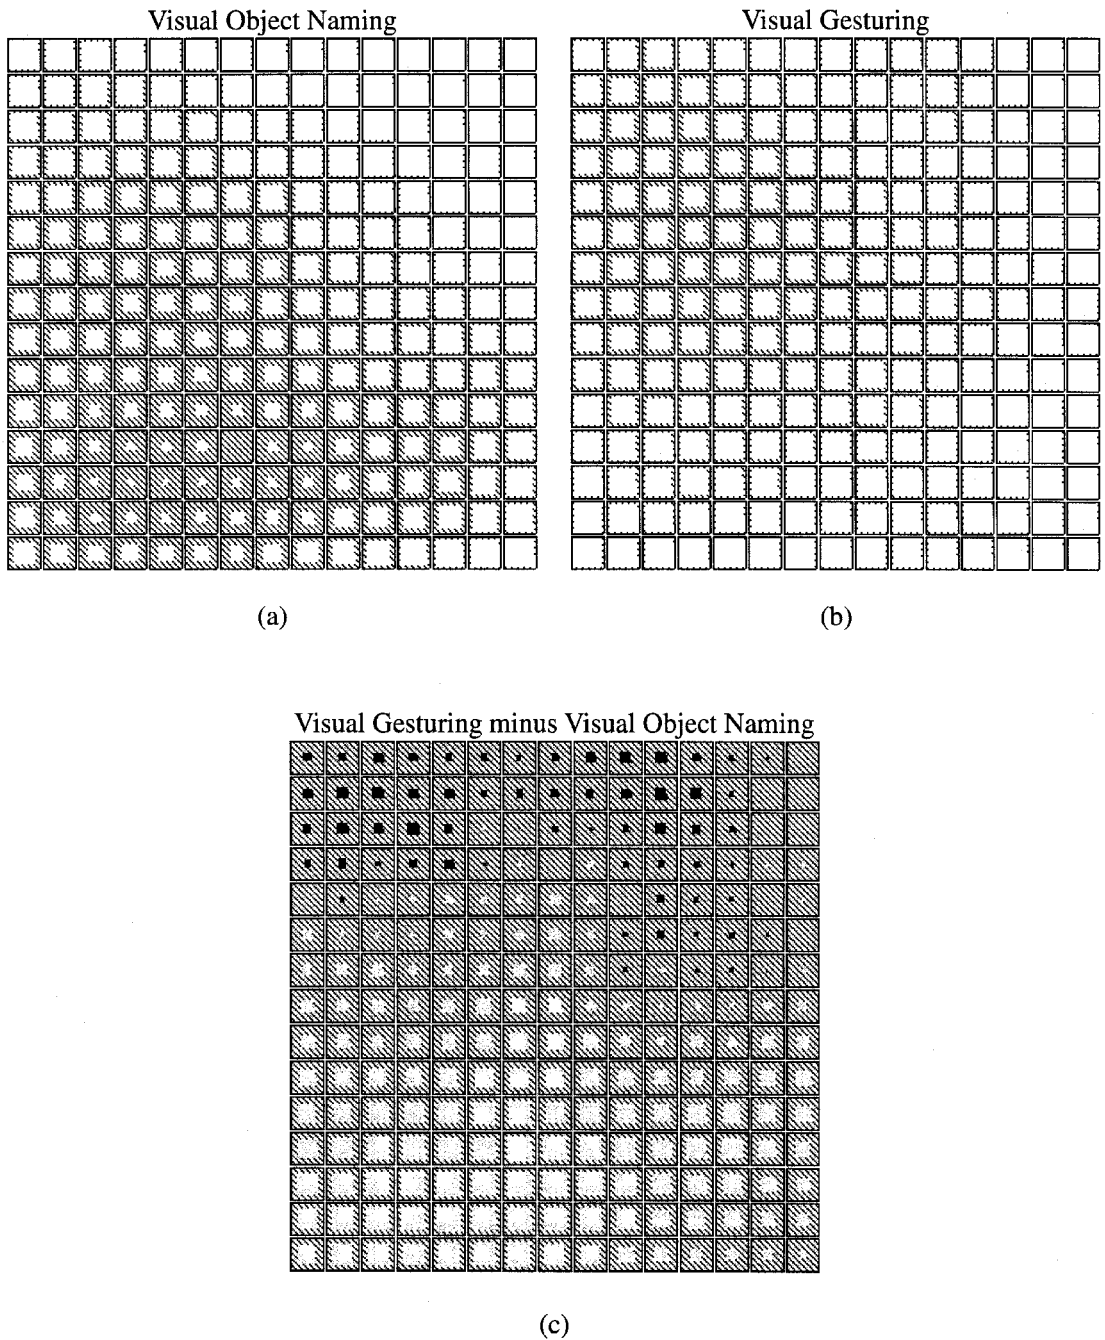
Figure 8 shows the levels of correct performance on visual object naming and visual gesturing for Vision-to-Semantics lesions centred at each location of Semantics. Two general effects are apparent. The first is that the locations of damage that are most detrimental to the two tasks differ. Not surprisingly, a given task is most impaired by Vision-to-Semantics lesions that are centred 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 an 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 differences in levels of correct performance (see Figure 8c). 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 behaviour 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 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 lines, Schnider, Hanlon, Alexander,





**Figure 8.** 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 difference between these levels of performance [(b) minus (a)], as a function of the centre of the lesion within Semantics. The size of the square at each position corresponds to the mean performance level following lesions centred 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.


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 object naming, visual object naming, visual action naming, and visual gesturing following Vision-to-Semantics

lesions centred on the seventh unit from the left in the bottom row of Semantics (see Figure 8c). Figure 9 shows the results, averaged over 40 instances of lesion at each level of severity. For comparison, the figure also shows (with vertical lines and, for some, small tick marks) the levels of performance of the 14 optic aphasic patients listed in Table 1.

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 patients, the network does a good job of matching the magnitudes of the dissociation between visual naming versus gesturing across a range of severity. Although a number of patients exhibit dissociations that are smaller than that of the network, this is to be expected given that the patients may well have lesions that are less than optimal for producing optic aphasia (e.g., lesions to Vision-to-Semantics



**Figure 9.** Correct performance of the network on various tasks as a function of lesion severity, and corresponding levels of performance of the optic aphasic patients listed in Table 1 (labelled with the corresponding letters from that table). The top of each line indicates performance on visual gesturing whereas the bottom indicates performance on visual object naming. Each line is placed horizontally to match model performance on one of the tasks, with the choice made so as to minimise discrepancy with model performance on the other. A small horizontal tick mark, where present, indicates performance on naming actions associated with objects. Note that, for some patients, this performance exceeds that for visual gesturing.

connections that project further from Phonology; see Figure 8c). The more pressing challenges are those patients who exhibit dissociations that are larger than that produced by the network.

There are only three major discrepancies in this regard. The largest is the patient reported by Coslett and Saffran (1992) (case j), who was 21% correct at visual naming but perfect (20/20) at visual gesturing, whereas the model predicts gesturing performance of just under 80% correct. First, note that the model's performance falls just outside the 95% confidence interval for perfect performance with 20 observations (83–100%). Also, the patient was only 68% correct at tactile naming and at naming to spoken definition. Thus, the poor visual naming performance appears to be exaggerated by a more generalised anomia. Moreover, although Coslett and Saffran report that the patient “produced well-formed, readily identified, specific gestures” (p. 153), correct performance was apparently evaluated with knowledge of the viewed object. Thus, no attempt was made to ensure, as was true for the model, that equivalently stringent scoring criteria were applied to visual naming and gesturing (see Riddoch & Humphreys, 1987). Similar concerns regarding the scoring of visual gesturing apply in the other two discrepant cases (case m: Poeck, 1984; case n: Coslett & Saffran, 1989b). In addition, the model's performance is well with the 95% confidence intervals for these patients' performance levels. Thus, the model is capable of producing dissociations between visual object naming and visual gesturing that are quantitatively as large as those of the patients (considering tests of reliability) across the full range of severity of impairment.

The network is also far better at naming the action associated with a visually presented object than at naming the object itself, as is true of patients for whom it has been tested (e.g., Campbell & Manning, 1996; Goldenberg & Karlbauer, 1998; Teixeira Ferreira et al., 1997). In fact, the quantitative match is reasonably good with regard to the two cases in which it has been examined most thoroughly (case f: Campbell & Manning, 1996; case c: Teixeira Ferreira et al., 1997). Why is generating action names relatively preserved compared with generating object names? The reason is that the

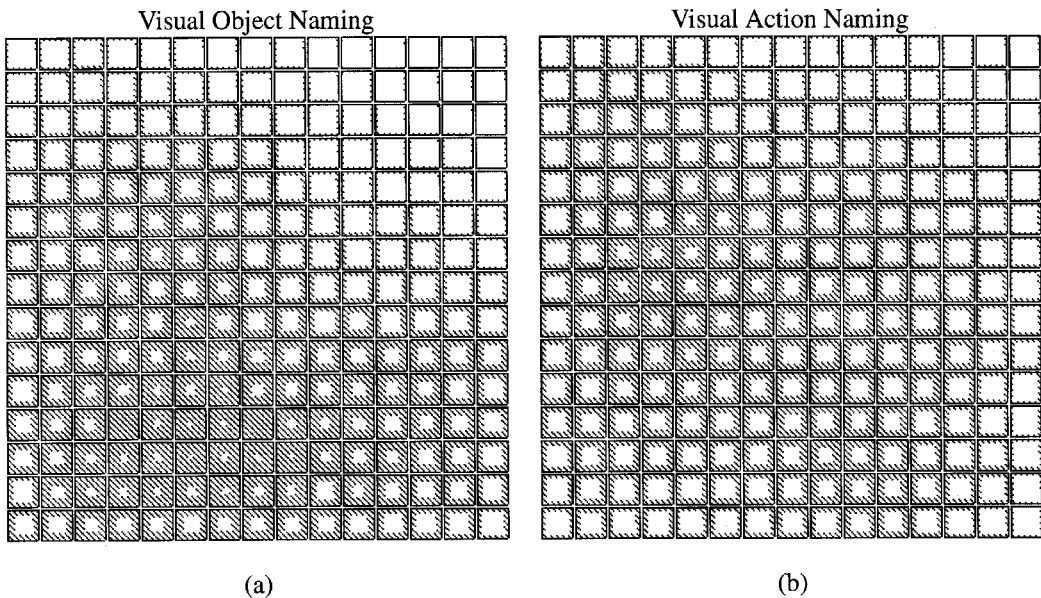
network learns to rely on support from the Action representation when generating an action name. It does this because a given visual input by itself is ambiguous with respect to the correct phonological output—it could be either the object name or the action name. Because the network is trained to generate 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 and to Action-to-Semantics connections*

To provide evidence for this 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 10 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 that impair generating action names are located closer to Action representations than are lesions that impair generating object names. Thus, the results support the claim that action naming involves interaction with Action representations.

In addition, a comparison of Figures 10a and 10b indicates that lesions to the upper-right quadrant of Semantics impair action naming more than object naming, even though the former is more robust to damage overall. This effect also holds following lesions to Action-to-Semantics connections



**Figure 10.** 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 centre of the lesion within Semantics. The depicted values range from 100% correct (full white square) to 20% correct (no white square).

coming into this region. For example, following lesions with  $SD = 2.5$  centred at position  $[11, 5]$  (where  $[x, y] = [0, 0]$  is the upper-left corner of Semantics), the network was 80.4% correct on visual object naming but only 67.2% correct on visual action naming,  $F(1, 99) = 12.33$ ,  $p < .001$ .

### Control simulation: Reduced action-related systematicity and separate training of gesturing and action naming

Earlier it was acknowledged that, in the simulation just reported, the degree of systematicity between Vision and Action and between Touch and Action may have been unduly high compared with the systematicity between Vision and Touch. The possibility was also raised that training gesturing and action naming simultaneously may have contributed to their relative preservation compared with object naming. To verify that these concerns do not undermine the implications of the previous simulation, a control simulation was carried out in which

the following five changes to the original simulation were made:

1. The assignments of Action representations to objects were adjusted so that object pairs in the same Vision category had a probability of .6 (reduced from .8) of being in the same Action category, and similarly for object pairs in the same Touch category.

2. To compensate for this reduced action-related systematicity, the strength of the topographic bias on learning was increased slightly, by reducing the standard deviation of the Gaussian used to scale weight derivatives from 10 to 9.

3. The action task for each object in each modality was split into two tasks: gesturing and action naming. For gesturing, targets were specified only for Action units, whereas for action naming, targets were specified only for Phonology units (just as in object naming).

4. Three task units were used instead of two in order to instruct the network to perform either object naming, action naming, or gesturing.

5. The network was trained for a total of 148,000 sweeps through the training corpus, or approximately 211 presentations per condition (cf., 220 for the original simulation).


The control simulation was identical to the original simulation in all other respects, including network architecture, initial random weights, learning parameters, and instances of lesions.

Figure 11 presents the correct performance of the control simulation on tactile and visual object naming, and visual action naming and gesturing, as a function of lesion severity analogous to the results in Figure 9. Although the magnitude of the dissociation between visual gesturing and visual object naming is slightly reduced compared with the original simulation, the control simulation still provides a reasonably good quantitative match to the four patients, including the relative preservation of visual naming of actions compared with objects. Thus, neither the potentially exaggerated amount of Vision-Action and Touch-Action systematicity in the original simulation, nor its use of a single task

to train gesturing and action naming, are fundamental to the model's ability to account for core characteristics of optic aphasia.

### 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 nontopographic network very similar to the current one, although the magnitude of 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, 1995; Plaut & Shallice, 1993b; Riddoch & Humphreys, 1987) that make no reference to topographic distinctions within the network (except in the separation of semantics from



**Figure 11.** Correct performance of the control network on various tasks as a function of lesion severity, and corresponding levels of performance of the optic aphasic patients listed in Table 1 (labelled with the corresponding letters from that table). The top of each line indicates performance on visual gesturing whereas the bottom indicates performance on visual object naming. Each line is placed horizontally to match model performance on one of the tasks, with the choice made so as to minimise discrepancy with model performance on the other. A small horizontal tick mark, where present, indicates performance on naming actions associated with objects. Note that, for some patients, this performance exceeds that for visual gesturing.

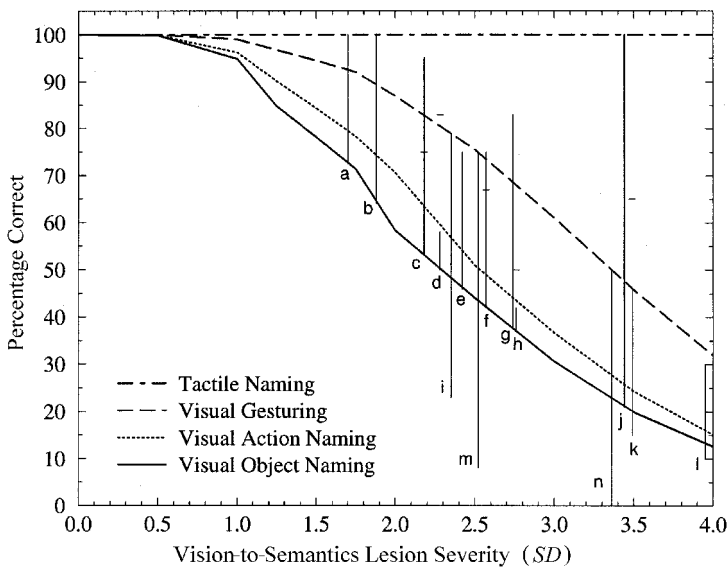
input and output modalities). It is thus important to determine if a quantitatively adequate account of optic aphasia can be obtained in the current context without a topographic bias on learning.

A second control simulation was carried out that was an exact replication of the initial one, except that error derivatives were not scaled as a function of connection length. The replication was otherwise identical in that it used the same 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 9. 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. 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.

## 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



**Figure 12.** Correct performance of the control simulation without topographic bias on various tasks as a function of lesion severity, and corresponding levels of performance of the optic aphasic patients listed in Table 1 (labelled with the corresponding letters from that table). The top of each line indicates performance on visual gesturing whereas the bottom indicates performance on visual object naming. Each line is placed horizontally to match model performance on one of the tasks, with the choice made so as to minimise discrepancy with model performance on the other. A small horizontal tick mark, where present, indicates performance on naming actions with objects. Note that, for some patients, this performance exceeds that for visual gesturing.

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 that 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 specialisation under the pressure of mediating between multiple input and output modalities. The nature of this graded specialisation is shaped by two general factors: (1) differences in the relative degree of systematicity among tasks (i.e., combinations of input-output mappings); and (2) a topographic bias on learning that favours 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.

The remainder of the article first considers limitations of the simulation and its relation to other models of semantics. 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.

### Limitations of the simulation

The current simulation is a demonstration of how graded modality-specific specialisation 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 modelling (see Kello & Plaut, 2002, for discussion). The former (e.g., Plaut, McClelland, Seidenberg, & Patterson, 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 (see also Plaut & Booth, 2000; Plaut & Gonnerman, 2000)—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 modelling 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 organisation of semantic knowledge more generally. Even without such details, though, the current work represents an important first step in developing such an account.

Another important limitation of the simulation is that it lacks conceptual knowledge that is not directly grounded in perceptual or action-based properties, including functional, encyclopaedic, and associative attributes. On the current view, such knowledge is largely acquired through language and could be incorporated into the current simulation through substantial elaboration of phonological inputs and outputs. Along these lines, Rogers and McClelland (2002b) have recently applied variants of the Rumelhart and Todd (1993) network to conceptual development, and interpret the propositional outputs of the network (e.g., *swims*, *has-feathers*) as collectively reflecting the kinds of things people say about objects. A fully adequate treatment of the impact of language on conceptual knowledge would, of course, be a major undertaking and is well beyond the scope of the current work. However, insofar as language can be introduced in the form of phonological input and output (and, for written language, visual input and output), its inclusion would not undermine the central claims of the current theory.

A final simplification of the current work concerns the implementation of the learning bias favouring 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), which 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 functional consequences as the implementation employed in the current work.

### **Relation to other connectionist models of semantics**

The current model has some similarities but also important differences with previous connectionist models incorporating semantic representations. In most of these models, semantic representations are assigned to entities by the modeller, in terms of particular patterns of activity over a set of units corresponding to semantic "features". In some cases (e.g., Harm, 1998, 2002; Harm & Seidenberg, 2001; Hinton, 1981; Hinton & Shallice, 1991; McRae et al., 1997; Plaut & Shallice, 1993a, 1993b; Small, Hart, Nguyen, & Gordon, 1995), these



features are verbalisable properties of the entity (e.g., *is-red*, *has-wings*), either assigned by hand (e.g., Hinton & Shallice, 1991) or derived from independent sources (e.g., feature norms; McRae et al., 1997). In other cases (e.g., Becker, Moscovitch, Behrmann, & Joordens, 1997; Devlin, Gonnerman, Andersen, & Seidenberg, 1998; Farah & McClelland, 1991; Gotts & Plaut, in press; Joordens & Becker, 1997; Lambon Ralph, McClelland, Patterson, Galton, & Hodges, 2001; Masson, 1995; McRae & Boisvert, 1998; Plaut, 1997; Plaut & Booth, 2000; Plaut & Gonnerman, 2000; Rogers et al., 2002; Tippet, McAuliffe, & Farah, 1995; Tyler, Moss, Durrant-Peateld, & Levy, 2000), individual units do not have specific interpretations; rather, more abstract patterns are generated and assigned to entities such that their collective similarities instantiate the putative properties of actual semantic representations (e.g., clustering into “categories”). In a few models, including the current one, semantic representations are not stipulated in advance but are internal, “hidden” representations that are learned in the service of performing various tasks (McClelland, St John, & Taraban, 1989; Rogers & McClelland, 2002a; Rumelhart & Todd, 1993; St John & McClelland, 1990). Although the resulting patterns of activity do not typically correspond to binary features that are either present or absent, the approach retains the core claim that what is relevant is the relative similarity structure of patterns and how this relates to the similarity structure over inputs and outputs. In this way, learning-based approaches can be viewed as elaborations of feature-based approaches by providing a basis for how such patterns (or, rather, ones with equivalent similarity structure) might emerge over the course of development.

One class of learning-based semantic models (Ritter & Kohonen, 1989; Miiikkulainen, 1993, 1997; Zorzi, Perry, Ziegler, & Coltheart, 1999) bears considerable superficial similarity to the current approach but, on close inspection, is actually quite different. This class consists of self-organising feature maps, also known as Kohonen networks (Kohonen, 1984, 1995). In such networks, units have a topological organisation, typically in the form of a two-dimensional grid. When

an input is presented to the units, the maximally responding unit, as well as those in a spatially defined neighbourhood around it, update their weights so as to respond to that input even more strongly (and less strongly to others). As a result, neighbouring units come to have similar response properties, and the grid as a whole comes to reflect the high-dimensional similarity structure among the input patterns. Although such networks, like the current one, have an underlying topology, the dimensionality of the resulting representations are far more constrained, being essentially limited to the dimensionality of the spatial neighbourhood that defines weight updates (which is typically two-dimensional). By contrast, in the current model, the topographic bias on learning simply influences the relative sensitivity of units to different inputs and outputs; semantic representations retain their full dimensionality (equal to the number of units). This is a critical distinction if a model is to capture the rich, high-dimensional similarity structure of actual semantic representations.

### Pattern of error responses in optic aphasia

To this point we have 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 behaviour 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 that 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., 2000, 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 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. (2000)—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 shell-fish, 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 corkscrew, 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. (2000) interpret this behaviour 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 specialisation. Beyond this broad characterisation, 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) characterisation 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 being simply a mild form of Graded Specialization in Semantics 32 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 (Geshwind, 1965; Hillis & Caramazza, 1995; 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 behavioural 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: These 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 & De 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 behaviour. 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, Levin, Huff, Kay, & Guinto, 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.<sup>7</sup> 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).

<sup>7</sup> Note that, if a topographic bias on learning were applied throughout the system, high-level visual representations would also be expected to exhibit graded functional specialisation depending on type of information being represented (see, e.g., Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999).

### *Optic aphasia from superadditivity of multiple impairments*

Farah (1990) provided an alternative answer to the way in which 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 utilising two lesioned pathways in series, as demonstrated by Sitton et al. (2000) in recent computational simulations.

There are, however, a number of problems with this proposal. First, as noted in the Introduction, Sitton et al. (2000) 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 single, 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. Although such problems have been documented in some cases (Campbell & Manning, 1996; Casanova & Roig-Rovira, 1985; Coslett & Saffran, 1992; Goldenberg & Karlbauer, 1998; Raymer et al., 1997), they are explicitly absent in others (Coslett & Saffran, 1989b; Hillis & Caramazza, 1995; 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, 1995). 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 emphasises 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 localised only in the left hemisphere (LH) and can be activated only by LH 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 callosum 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), lateralised 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 agnosic patients in having RH semantic systems with greater premorbid 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 (1995) 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 specialised in representing and processing information in that modality (see also Chao & Martin, 2000; Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995; Martin, Wiggs, Ungerleider, & Haxby, 1996, for functional imaging data supporting this perspective). Thus, for example,

semantic regions near action representations become particularly important in relating action to information in other modalities. This learned specialisation 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-specialised 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 behavioural, 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 behaviour of visual apperceptive agnosic 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). Whereas 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 (Carey, Harvey, & Milner, 1996).

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, Kawano, 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 emphasises 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 comes 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 recognising 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 recognise. Similarly, patients with semantic dementia (Snowden, Goulding, & Neary, 1989), a deteriora-

tion of conceptual knowledge due to progressive temporal lobe atrophy, often exhibit relatively normal use of objects despite severe impairments in their conceptual knowledge (Bozeat, Lambon Ralph, Patterson, & Hodges, 2002; Buxbaum, Schwartz, & Carew, 1997; Hodges, Patterson, Oxbury, & Funnell, 1992; Lauro-Grotto et al., 1997; Schwartz, Marin, & Saffran, 1979; Snowden, Griffiths, & Neary, 1995). Riddoch, Humphreys, Heslop, and Castermans (2002) reported a similar pattern of performance in a patient with impaired visual comprehension due to Alzheimer's Type Dementia. 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 man-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, 2002). A natural interpretation of such a pattern is that object and action knowledge are subserved by separate neural mechanisms (although see Plaut, 1995a; Shallice, 1988).

Recent findings (Bozeat, Lambon Ralph, Patterson, & Hodges, in press; Hodges, Bozeat, Lambon Ralph, Patterson, & Spatt, 2000; Hodges, Spatt, & Patterson, 1999) suggest, however, that the relative preservation of action knowledge in semantic dementia may not be as independent of object knowledge as first thought. For example, Hodges et al. (2000) 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.

## CONCLUSIONS

In 1973, Alan Newell—one of the founders of computational modelling 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 favour 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 behaviour.

In many ways, the question of whether semantics consists of a unitary or of multiple systems constitutes a similar dichotomy, with its own tortuous history (cf. Caramazza et al., 1990; Shallice, 1993). Each side constitutes a valid and plausible theoretical position with substantial empirical support, but each side also faces considerable challenges. So much 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 specialised 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. 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.

Manuscript received 16 October 2000

Revised manuscript received 29 May 2002

Revised manuscript accepted 11 June 2002

## 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.
- Amedi, A., Malach, R., Hendler, T., Peled, S., & Zohary, E. (2001). Visuo-haptic object-related activation in the ventral visual pathway. *Nature Neuroscience*, 4, 324–330.
- 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* (pp. 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, 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—Biological Sciences, 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, 1–48.
- Beauvois, M.-F., Saillant, B., Meininger, V., & Lhermitte, F. (1978). Bilateral tactile aphasia: A tacto-verbal dysfunction. *Brain*, 101, 381–401.
- Becker, S., Moscovitch, M., Behrmann, M., & Joordens, S. (1997). Long-term semantic priming: A computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1059–1082.
- Beeman, M., & Chiarello, C. (Eds.). (1998). *Right hemisphere language comprehension: Perspectives from cognitive neuroscience*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- 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, 146–184.
- Bozeat, S., Lambon Ralph, M.A., Patterson, K., & Hodges, J.R. (2002). The influence of personal familiarity and context on object use in semantic dementia. *Neurocase*, 8, 127–134.
- Bozeat, S., Lambon Ralph, M.A., Patterson, K., & Hodges, J.R. (in press). When objects lose their meanings: What happens to their use? *Cognitive, Affective and Behavioral Neuroscience*.
- Buxbaum, L.J., & Saffran, E.M. (2002). *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, 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., 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, 1–34.
- Carey, D.P., Harvey, M., & Milner, A.D. (1996). Visuomotor sensitivity for shape and orientation in a patient with visual form agnosia. *Neuropsychologia*, 34, 329–337.
- Casanova, J.P., & 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, 437–448.
- Chao, L.L., & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage*, 12, 478–484.
- Chiarello, C., & Richards, L. (1992). Another look at categorical priming in the cerebral hemispheres. *Neuropsychologia*, 30, 381–392.
- Coltheart, M. (1980). Deep dyslexia: A right-hemisphere hypothesis. In M. Coltheart, K. Patterson, & J.C. Marshall (Eds.), *Deep dyslexia* (pp. 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.
- Cree, G.S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, 23, 371–414.
- Davidoff, J., & De Bleser, R. (1993). Optic aphasia: A review of past studies and reappraisal. *Aphasiology*, 7, 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., & Saetti, M.C. (1997). Associative agnosia and optic aphasia: Qualitative or quantitative difference? *Cortex*, 33, 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.
- Devlin, J.T., Gonnerman, L.M., Andersen, E.S., & Seidenberg, M.S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of Cognitive Neuroscience*, 10, 77–94.
- 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, 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, 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* (pp. 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., & McClelland, J.L. (1991). A computational model of semantic memory impairment: Modality-specificity and emergent category-specificity. *Journal of Experimental Psychology: General*, 120, 339–357.
- Ferro, J.M., & Santos, M.E. (1984). Associative visual agnosia: A case study. *Cortex*, 20, 212–234.
- Freund, D.C. (1889). Ueber optische aphasia und seelenblindheit. *Archive für Psychiatrie und Nervenkrankheiten*, 20, 276–297. (Translated by Beaton, A., Davidoff, J., and Erstfeld, U., 1991, *Cognitive Neuropsychology*, 8, 21–38.)
- 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, 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, 471–491.
- Goodale, M.A., & Milner, A.D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15, 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, 154–156.
- Gotts, S., & Plaut, D.C. (in press). The impact of synaptic depression following brain damage: A connectionist account of “access/refractory” and “degraded-store” semantic impairments. *Cognitive, Affective and Behavioral Neuroscience*.
- 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.
- Harm, M.W. (1998). *Division of labor in a computational model of visual word recognition*. Unpublished doctoral dissertation, Department of Computer Science, University of Southern California, Los Angeles, CA.
- Harm, M.W. (2002). *Building large-scale distributed semantic feature sets with WordNet* (Tech. Rep. No. PDP.CNS.02.01). Pittsburgh, PA: Carnegie Mellon University, Center for the Neural Basis of Cognition.
- Harm, M.W., & Seidenberg, M.S. (2001). Are there orthographic impairments in phonological dyslexia? *Cognitive Neuropsychology*, 18, 71–92.
- Hillis, A.E., & Caramazza, A. (1995). Cognitive and neural mechanisms underlying visual and semantic processing: Implications from “optic aphasia.” *Journal of Cognitive Neuroscience*, 7, 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. (1981). Implementing semantic networks in parallel hardware. In G.E. Hinton & J.A. Anderson (Eds.), *Parallel models of associative memory* (pp. 161–188). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Hinton, G.E. (1989). Connectionist learning procedures. *Artificial Intelligence*, 40, 185–234.
- Hinton, G.E., & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review*, 98, 74–95.
- Hodges, J.R., Bozeat, S., Lambon Ralph, M.A., Patterson, K., & Spatt, J. (2000). The role of concep-

- tual knowledge in object use: Evidence from semantic dementia. *Brain*, *123*, 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*. Hove, UK: Lawrence Erlbaum Associates Ltd.
- 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.
- Ishai, A., Ungerleider, L.G., Martin, A., Schouten, J.L., & Haxby, J.V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Science, USA*, *96*, 9379–9384.
- Jacobs, R.A. (1997). Nature, nurture and the development of functional specializations: A computational approach. *Psychonomic Bulletin and Review*, *4*, 299–309.
- Jacobs, R.A., & Jordan, M.I. (1992). Computational consequences of a bias toward short connections. *Journal of Cognitive Neuroscience*, *4*, 323–336.
- James, T.W., Humphrey, G.K., Gati, J.S., Servos, P., Menon, R.S., & Goodale, M.A. (2002). Haptic study of three-dimensional objects activates extrastriate visual areas. *Neuropsychologia*, *40*, 1706–1714.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford: Blackwell.
- Joordens, S., & Becker, S. (1997). The long and short of semantic priming effects in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 1083–1105.
- 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. (2002). *Strategic control over rate of processing in word reading: A computational investigation of the tempo-naming task*. Manuscript submitted for publication.
- Kohonen, T. (1984). *Self-organization and associative memory*. New York: Springer-Verlag.
- Kohonen, T. (1995). *Self-organization maps*. New York: Springer-Verlag.
- Lambon Ralph, M.A., McClelland, J.L., Patterson, K., Galton, C.J., & Hodges, J.R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *Journal of Cognitive Neuroscience*, *13*, 341–356.
- 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*, 1–12.
- Lauro-Grotto, R., Piccini, C., & Shallice, T. (1997). Modality-specific operations in semantic dementia. *Cortex*, *33*, 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 für Psychiatrie und Nervenkrankheiten*, *21*, 222–270.
- 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*, 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.
- Martin, A., Wiggs, C.L., Ungerleider, L.G., & Haxby, J.H. (1996). Neural correlates of category-specific knowledge. *Nature*, *379*, 649–652.
- Masson, M.E.J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 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.

- McClelland, J.L., & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*, 159–188.
- McClelland, J.L., St John, M., & Taraban, R. (1989). Sentence comprehension: A parallel distributed processing approach. *Language and Cognitive Processes*, *4*, 287–335.
- 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* (pp. 502–507). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 558–572.
- McRae, K., De Sa, V.R., & Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 99–130.
- Miikkulainen, R. (1993). *Subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory*. Cambridge, MA: MIT Press.
- Miikkulainen, R. (1997). Dyslexic and category-specific aphasic impairments in a self-organizing feature map model of the lexicon. *Brain and Language*, *59*, 334–366.
- Milner, A.D., & Goodale, M.A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- 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* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- 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* (pp. 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*, 895–938.
- Paivio, A. (1971). *Imagery and verbal processes*. London: Holt, Rinehart & Winston.
- 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*, 263–269.
- Pilgrim, E., & Humphreys, G.W. (1991). Impairment of action to visual objects in a case of ideomotor apraxia. *Cognitive Neuropsychology*, *8*, 459–473.
- Plaut, D.C. (1995a). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*, *17*, 291–321.
- Plaut, D.C. (1995b). Semantic and associative priming in a distributed attractor network. In *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 37–42). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Plaut, D.C. (1997). Structure and function in the lexical system: Insights from distributed models of naming and lexical decision. *Language and Cognitive Processes*, *12*, 767–808.
- Plaut, D.C. (1999). Systematicity and specialization in semantics. In D. Heinke, G.W. Humphreys, & A. Olson (Eds.), *Connectionist models in cognitive neuroscience: 5th Neural Computation and Psychology Workshop*. London: Springer-Verlag.
- Plaut, D.C., & Booth, J.R. (2000). Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review*, *107*, 786–823.
- Plaut, D.C., & Gonnerman, L.M. (2000). Are nonsemantic 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*, 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*, 89–117.

- Poeck, K. (1984). Neuropsychological demonstration of splenial interhemispheric disconnection in a case of "optic anomia." *Neuropsychologia*, *22*, 707–713.
- Quartz, S., & Sejnowski, T.J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*, *20*, 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., Hillis, A.E., & Caramazza, A. (1993). The role of representations in cognitive theory: More on multiple semantics and the agnosias. *Cognitive Neuropsychology*, *10*, 235–249.
- 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* (pp. 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* (pp. 133–160). Hove, UK: Psychology Press.
- Riddoch, M.J., & Humphreys, G.W. (1987). Visual object processing in optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology*, *4*, 131–185.
- Riddoch, M.J., Humphreys, G.W., Coltheart, M., & Funnell, E. (1988). Semantic systems or system? Neuropsychological evidence re-examined. *Cognitive Neuropsychology*, *5*, 3–25.
- Riddoch, M.J., Humphreys, G.W., Heslop, J., & Castermans, E. (2002). Dissociations between object knowledge and everyday action. *Neurocase*, *8*, 100–110.
- Riddoch, M.J., Humphreys, G.W., & Price, C.J. (1989). Routes to action: Evidence from apraxia. *Cognitive Neuropsychology*, *6*, 437–454.
- Ritter, H., & Kohonen, T. (1989). Self-organizing semantic maps. *Biological Cybernetics*, *61*, 241–254.
- Rogers, T.T., Lambon Ralph, M.A., Garrard, P., Bozeat, S., McClelland, J.L., Hodges, J.R., & Patterson, K. (2002). *The structure and deterioration of semantic memory: A computational and neuropsychological investigation*. Manuscript submitted for publication.
- Rogers, T.T., & McClelland, J.L. (2002a). *Semantics without categorization: A parallel distributed processing approach to the acquisition and use of natural semantic knowledge*. Manuscript in preparation.
- Rogers, T.T., & McClelland, J.L. (2002b). *Semantics without categorization: A parallel distributed processing approach to the acquisition and use of natural semantic knowledge*. Manuscript submitted for publication.
- Rogers, T.T., & Plaut, D.C. (2002). 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* (pp. 233–267). Oxford: Oxford University Press.
- Rumelhart, D.E., & Todd, P.M. (1993). Learning and connectionist representations. In D.E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 3–30). Cambridge, MA: MIT Press.
- 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*, 631–647.
- Saffran, E.M., Bogvo, 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* (pp. 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.
- 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.N., & 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* (pp. 111–128). Hove, UK: Lawrence Erlbaum Associates Ltd.

- Shallice, T. (1988). *From neuropsychology to mental structure*. Cambridge: Cambridge University Press.
- Shallice, T. (1993). Multiple semantics: Whose confusions? *Cognitive Neuropsychology*, *10*, 251–261.
- Shallice, T., & Jackson, M. (1988). Lissauer on agnosia. *Cognitive Neuropsychology*, *5*, 153–192.
- Sharkey, A.J., & Sharkey, N.E. (1992). Weak contextual constraints in text and word priming. *Journal of Memory and Language*, *31*, 543–572.
- Shelton, J.R., & Caramazza, A. (1999). Deficits in lexical and semantic processing: Implications for models of normal language. *Psychonomic Bulletin and Review*, *6*, 5–27.
- 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. (2000). Superadditive effects of multiple lesions in a connectionist architecture: Implications for the neuropsychology of optic aphasia. *Psychological Review*, *107*, 709–734.
- Small, S.L., Hart, J., Nguyen, T., & Gordon, B. (1995). Distributed representations of semantic knowledge in the brain. *Brain*, *118*, 441–453.
- Snowden, J.S., Goulding, P.J., & Neary, D. (1989). Semantic dementia: A form of circumscribed cerebral atrophy. *Behavioral Neurology*, *2*, 167–182.
- Snowden, J.S., Griffiths, H., & Neary, D. (1995). Autobiographical experience and word meaning. *Memory*, *3*, 225–246.
- St John, M.F., & McClelland, J.L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, *46*, 217–257.
- 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.
- Tippett, L.J., McAuliffe, S., & Farah, M.J. (1995). Preservation of categorical knowledge in Alzheimer's disease: A computational account. *Memory*, *3*, 519–553.
- Tyler, L.K., Moss, H.E., Durrant-Peateld, M.R., & Levy, J.P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and Language*, *75*, 195–231.
- 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* (pp. 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. (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*, 1465–1473.
- Warrington, E.K., & Shallice, T. (1984). Category-specific semantic impairments. *Brain*, *107*, 829–854.
- 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* (pp. 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*, 14–32.
- Zorzi, M., Perry, C., Ziegler, J., & Coltheart, M. (1999). Category-specific deficits in a self-organizing model of the lexical-semantic system. In D. Heinke, G.W. Humphreys, & A. Olson (Eds.), *Connectionist models in cognitive neuroscience: 5th Neural Computation and Psychology Workshop* (pp. 137–148). London: Springer-Verlag.